

Working paper series

**The unequal gains from product innovations:
Evidence from the US retail sector**

Xavier Jaravel

March 2017

<http://equitablegrowth.org/working-papers/unequal-gains-from-product-innovations/>

The Unequal Gains from Product Innovations: Evidence from the US Retail Sector

Xavier Jaravel

March 2017

Abstract

Using detailed barcode-level data in the US retail sector, I find that from 2004 to 2013 higher-income households systematically experienced a larger increase in product variety and a lower inflation rate for continuing products. Annual inflation was 0.65 percentage points lower for households earning above \$100,000 a year, relative to households making less than \$30,000 a year. I explain this finding by the equilibrium response of firms to market size effects: (A) the relative demand for products consumed by high-income households increased because of growth and rising inequality; (B) in response, firms introduced more new products catering to such households; (C) as a result, continuing products in these market segments lowered their price due to increased competitive pressure. I use changes in demand plausibly exogenous to supply factors from shifts in the national income and age distributions over time to provide causal evidence that increasing relative demand leads to more new products and lower inflation for continuing products, implying that the long-term supply curve is downward-sloping. Based on this channel, I develop a model predicting a secular trend of faster-increasing product variety and lower inflation for higher-income households, which I test and validate using Consumer Price Index and Consumer Expenditure Survey data on the full consumption basket going back to 1953.

Xavier Jaravel

Stanford University

Stanford Institute for Economic Policy Research

xjaravel@stanford.edu

I am particularly indebted to Daron Acemoglu, Philippe Aghion, Raj Chetty, Ed Glaeser, Larry Katz, Pete Klenow, Josh Lerner, Marc Melitz and Andrei Shleifer for extensive advice and comments. For thoughtful discussions and suggestions, I thank Ufuk Akcigit, Adrien Auclert, David Autor, Matt Basilico, Shai Bernstein, Michael Best, Barbara Biasi, Laura Blattner, Nick Bloom, Alexis Bonnardon, Kirill Borusyak, Michael Boskin, Emanuele Colonnelli, Diego Comin, Gary Chamberlain, Varanya Chaubey, Gabriel Chodorow-Reich, John Coglianesi, Tilman Dette, Mark Duggan, Max Eber, Anastassia Fedyk, Josh Feng, Peter Ganong, Andy Garin, Matt Gentzkow, Paul Goldsmith-Pinkham, Robert Hall, Ben Hebert, Nathan Hendren, Leander Heldring, Nathan Hipsman, Giacomo Hodson, Erik Hurst, Simon Jaeger, Chad Jones, Max Kasy, Bill Kerr, David Koretz, David Laibson, Danial Lashkari, Jing Li, Yuanjian Li, Hannes Malmberg, Leslie McGranahan, Sendhil Mullainathan, Emi Nakamura, Arash Nekoei, Paul Oyer, Anna Paulson, Thomas Piketty, Dan Pollman, Benjamin Pugsley, Ananth Raman, Alexandra Roulet, Ishuwar Seetharam, Isaac Sorkin, Jann Spiess, Stefanie Stancheva, Jeremy Stein, Larry Summers, Colin Tomkins-Bergh, Chris Tonetti, Gabriel Unger, Michael Webb, Heidi Williams, David Yang, Alwyn Young, my discussants Katharine Abraham and Jessie Handbury and seminar participants at the 2016 NBER Summer Institute, as well as seminar participants at Harvard, Stanford, the College de France, the Kauffman Foundation, the 2016 Econometric Society North American Summer Meeting and the 2016 Society for Economic Dynamics conference. I also thank Max Morales for research assistance and I gratefully acknowledge financial support from the HBS doctoral office, the Kauffman Foundation, the Stanford Institute for Economic Policy Research and the Washington Center for Equitable Growth.

“The capitalist achievement does not typically consist in providing more silk stockings for queens but in bringing them within the reach of factory girls in return for steadily decreasing amounts of effort... The capitalist process, not by coincidence but by virtue of its mechanism, progressively raises the standard of life of the masses.”

Joseph Schumpeter, *Capitalism, Socialism, and Democracy*

1 Introduction

Do changes in the product market affect purchasing-power inequality? Who benefits from product innovations? Various studies have investigated how skill-biased technical change affected the relative price of skills in the labor market and resulted in higher nominal income inequality (Acemoglu, 1998). Much less attention has been paid to how price changes in the product market and the introduction of new products may differentially affect households across the income distribution. Yet households of different income levels consume very different kinds of goods and services. Due to price changes in the product market over time, as well as changes in product variety, trends in purchasing-power inequality may therefore differ from trends in nominal income inequality. Product innovations may greatly affect purchasing-power inequality by increasing the variety and quality of goods available in specific consumer segments, as well as by driving down the price of existing products in these segments due to increased competitive pressure. This paper shows the relevance of this hypothesis in the US retail sector.

I investigate this question in a series of steps. First, I show that in the retail sector between 2004 and 2013 the quality-adjusted price index of high-income households rose substantially *more slowly* than that of low-income households, which amplified inequality.¹ To establish this result, I build income-group-specific quality-adjusted price indices using detailed barcode-level scanner data, allowing me to observe consumption patterns across income groups, price changes for all products available in consecutive years (inflation) and changes in product variety (product entry and exit). Second, I find that firms’ equilibrium response to changes in demand across income groups explains why the quality-adjusted price index of high-income consumers rose more slowly than that of the low-income. Specifically, this analysis shows that because demand from high-income households grew faster during this period, firms strategically introduced more new products catering to these consumers, which in turn drove down the price of existing products in these segments due to competitive dynamics. The retail sector is ideal to conduct this investigation because it accounts for a sizable share of US GDP, rich data is available, and the notion of product (barcode) is well-defined and consistent over time. Finally, motivated by these findings, I develop a model predicting a secular increase in product variety and lower inflation for higher-income households. I test and validate this prediction, in retail and in other sectors, by using data from the Consumer Price Index (CPI) and the Consumer Expenditure Survey (CEX) on the full consumption basket of American households going back to 1953.

¹This finding stands in contrast with the existing literature, which has found no systematic difference in inflation rates for households across the income distribution. See for instance Hobijn and Lagakos (2003), McGranahan and Paulson (2005) and Chiru (2005). I reconcile my results with the existing literature by showing that these studies suffer from aggregation bias. The study closest to my paper is Argente and Lee (2016), who document lower inflation for higher-income households in the retail sector during the Great Recession. In contrast, I show that lower inflation for higher-income households is a long-term trend in retail. Prior work documenting long-term inflation trends in scanner data is inconclusive (Broda and Romalis, 2009).

In the first part of the paper, I establish that in the US retail sector from 2004 to 2013 higher-income households experienced lower inflation and a faster increase in product variety. The magnitude of these effects is large: over the sample period, the average annual inflation rate was 0.65 percentage points lower for households making more than \$100,000 a year, compared with households making less than \$30,000.² These results are very stable for a wide variety of price indices and hold before, during and after the Great Recession, both across and within product categories. They are based on detailed product-level data from the Nielsen Homescan Consumer Panel and Retail Scanner datasets, which are representative of a large subset of the retail sector, accounting for approximately 40% of household expenditures on goods and 15% of total household expenditures.

The analysis delivers a general methodological lesson for the measurement of inflation by statistical agencies: the difference in inflation rates across income groups can be accurately measured only with product-level data. A large share of the inflation difference between income groups occurs within detailed product categories, which cannot be captured by price series based on data aggregated at a level similar to what the Bureau of Labor Statistics (BLS) and other statistical agencies currently use. These findings challenge the result from the existing literature that inflation is similar across the income distribution (McGranahan and Paulson, 2005) and suggest that trends in purchasing-power inequality may be diverging from trends in nominal income inequality. Collecting product-level data is key to accurately measure this divergence.

Moreover, the results in the first part of the paper have direct implications for the indexation of various government transfers that are indexed on food CPI, such as the Food Stamp and Child Nutrition programs. Between 2004 and 2013, food CPI indexation implied an increase in nominal food stamp benefits of 24.8%. In contrast, indexation on the non-homothetic food price index for eligible households implies a 35.5% increase.

In the second part of the paper, I show that the equilibrium response of supply to faster growth of demand from high-income consumers explains the patterns of differential inflation and increase in product variety across the income distribution in the retail sector. This hypothesis appears natural to investigate because it is well-documented (e.g. Song et al., 2016) that in recent decades the share of national income accruing to high-income consumers has steadily increased - both because more and more households enter high-income brackets as the economy grows and because of rising inequality. Intuitively, firms can respond to changes in relative market size by skewing product introductions toward market segments that are growing faster. This process can lead to a decrease in the price of existing products in the fast-growing market segments because increased competitive pressure from new products pushes markups down. A variety of patterns in the data support this theory: product categories that grow faster feature a greater increase in product variety, lower inflation, and disproportionately cater to higher-income households.

To test the causal claim that increases in demand lead to an increase in product variety and a fall in

²As discussed in Section 3, increasing product variety is valuable on its own, but empirically most of the welfare difference between households across the income distribution is captured by price changes in the basket of products that are available across years. The annual inflation difference of 65 basis points in retail is large relative to the null hypothesis from the literature of no sustained inflation difference, but also relative to the average inflation rate (1.9% in CPI data) and to the increase in income inequality in this range of the income distribution (93 basis point per year in the Census public-use microdata, comparing households making more than \$100,000 a year to those earning below \$30,000, in 2004 dollars).

inflation, I use anticipated shifts in the national age and income distributions between 2004 and 2013 to estimate the causal effect of changes in market size on product innovations and inflation. I estimate the age-by-income spending profile of products in the base period to build predictors of demand in future periods that vary only due to changes in the age and income distributions (rather than due to changes in spending patterns that are endogenous to supply factors). This research design is similar in spirit to Acemoglu and Linn (2004).³ I find large effects: a 1 percentage point increase in the annualized growth of predicted demand leads to a 2.73 percentage point increase in the share of spending on new products and a 0.43 percentage point decline in inflation.⁴ This finding implies that the long-term supply curve is downward-sloping, which rules out a broad class of supply models and provides support for others.⁵ Moreover, this result shows that the equilibrium response of firms alters the cost-benefit analysis of any policy that affects relative market size, such as the Food Stamp program: transfers are more effective in general equilibrium than in partial equilibrium because they induce lower prices for the recipients through the supply response.

Using the point estimates for the effect of demand on supply, I show that historical changes in the US income distribution imply substantial inflation inequality through the supply response to market size effects. Shifts in the income distribution over time generate changes in demand across the product space, to which I apply the point estimates from the causal research design to infer the implications for product innovations and inflation. The predicted patterns closely approximate the actual relationship between consumer income, products innovation and inflation across product categories. In other words, absent changes in the income distribution and the induced supply response, retail inflation would not have been much lower for higher-income households.

In the last part of the paper, I develop a micro-founded general equilibrium model consistent with the various aspects of the data and featuring a secular trend of faster-increasing product variety and lower inflation for higher-income households. Using translog preferences nested in CES preferences, the model flexibly accommodates non-homotheticities with an arbitrary number of consumer groups and sectors, endogenous product variety, and endogenous markups. This unified framework brings together the various results of the

³Identification requires that socio-demographic groups that grow faster should not source their consumption from parts of the product space where innovation or inflation systematically differ due to unobserved supply factors. For instance, considering households in their thirties, the numbers of low-income and high-income households grew faster than the number of middle-income households during the sample. 30-year-olds are the main market for baby diapers and higher-income groups tend to purchase higher-quality diapers. Therefore, changes in the income distribution increased demand for both low-end and high-end diapers, relative to middle-range diapers. This identifying variation for demand shocks across the quality ladder within baby diapers appears unlikely to be correlated with supply shocks, which would need to vary non-monotonically across quality ladder for baby diapers. Section 4.2.1 provides a formal derivation and a complete discussion of the identification conditions in this research design.

⁴I also introduce a research design exploiting variation in food stamp policy across US states between 2000 and 2007 to trace out the causal impact of changes in *per capita spending* on product innovations and inflation. States changed eligibility requirements, which had a large impact on the food-stamp take-up rate (Ganong and Liebman, 2016) and generated variation in market size for products with local brand capital (Bronnenberg et al. 2012). I find point estimates consistent with those obtained with the first research design, based on changes in the national age and income distributions over time. In principle, changes in market size induced by changes in the number of consumers (as in the first research design) or by changes in per capita spending (as in the food-stamp design) could have different effects on the equilibrium. I use this evidence to discipline the model.

⁵The evidence that product innovations endogenously follow changes in market size both across and within detailed product categories is in line with endogenous growth models. See for instance Aghion and Howitt (1992), Jones (1995) and Acemoglu (2002).

paper by providing a tractable non-homothetic price index (part 1) and comparative statics for the study of the market size channel (part 2). Due to non-homotheticities and the supply response to market size effects induced by long-run trends of growth and rising inequality, the model predicts a secular pattern of decreasing price index for higher-income households relative to lower-income households. I test and validate this prediction using CPI and CEX data on the full consumption basket going back to 1953. Finally, I use the model for welfare calculations: from 2004 to 2013, inflation inequality in the retail sector alone led to a large increase in purchasing-power inequality between the top and bottom income quintiles, equal to 0.22 percentage points per year, about one fourth of the effect of increasing income inequality.

Quantity, price and innovation dynamics in the food industry in recent years illustrate particularly well the core ideas developed in this paper. Organic food sales have grown at an average annualized rate of 11.2% between 2004 and 2013, compared with 2.8% for total food sales, in the context of increasing demand from higher-income households. The price premium for organic products shrunk significantly: for instance, organic spinach cost 60% more than nonorganic spinach in 2004, compared with only 7% more today (Appendix Figure A1). Low inflation for organic products brought down the food CPI, implying that it reduced the rate of increase in food stamps through indexation, although most food-stamp recipients do not purchase organic products.⁶ Bell et al. (2015) show how innovations and increased competition led to the fall in organic food prices.⁷

Overall, this paper provides new evidence challenging the existing literature primarily in two respects. First, the literature suggests that households across the income distribution experience similar inflation rates (McGranahan and Paulson, 2005), except in some specific periods in the short-run (Chiru, 2005, and Argente and Lee, 2016). Second, theoretical work has focused on the “product cycle”, the idea that innovation is driven by economies of scale and allows for a trickle-down process bringing to the mass market the new products that were initially enjoyed by a select few at the top of the income distribution. In other words, innovations should lower all consumers’ price indices at approximately the same rate as they diffuse across the income distribution. My findings show that market size effects and endogenously-increasing product variety is an important force, distinct from the product cycle and contributing to lower quality-adjusted inflation for higher-income households when market size grows faster for premium products. More generally, this paper contributes to various strands of literature studying income inequality, structural change, price indices, directed technical change and monopolistic competition dynamics.⁸

⁶Handbury et al. (2015) show that higher-income households have stronger preferences for organic food products. Therefore, they disproportionately benefit from the falling price premium for these products. The implied welfare difference between higher- and lower-income households can be accurately captured only with detailed micro data: the important divergence in price dynamics occurs between organic versus nonorganic spinach/granola/coffee/carrot/milk/etc., i.e. within detailed product categories, rather than across broader categories like fruit versus vegetables.

⁷Due to growing market demand, farmers undertook investments to obtain organic label certifications: certified organic pasture, rangeland, cropland and livestock have been expanding at double-digit rates since 2004. To reduce cost, organic producers harnessed innovative techniques like integrated pest management and relied on innovations to product formulations. More recently, conventional consumer packaged goods companies, such as Hormel, Kellogg, General Mills, and PepsiCo, entered the organic market and created venture capital funds to invest in startups of organic products. The increase in competition led to lower prices and reduced profitability for early entrants like Hain Celestial, which had been outperforming the stock market for several years.

⁸More precisely, related work examines nominal income inequality (Autor, Katz and Kruger (1998), Autor, Katz and Kearney

The remainder of the paper is organized as follows: Section 2 describes the data; Section 3 describes patterns of increasing product variety and inflation in retail across income groups; Section 4 establishes that increases in demand cause an increase in product variety and lower inflation for continuing products — in such a way that lower inflation in retail for higher-income households is explained by the supply response to changes in the income distribution; Section 5 presents the model, the evidence from the CPI and CEX data on the secular trend of lower inflation for higher-income households, and the implications for inequality. A number of theoretical results, estimation details and robustness checks are reported in appendices.

2 Data Sources and Summary Statistics

In this section, I discuss how product-level scanner data in the retail sector is uniquely suited to accomplish the two main goals of the paper. On the measurement front, this data is ideal to compute income-group-specific inflation rates as well as changes in product variety across the income distribution because I observe the spending of a large panel of consumers at the product level. On the mechanism front, the strength of this data is that the notion of product (barcode) is well defined and I can thus measure whether firms respond to changes in relative market size by skewing product introductions toward market segments that are growing faster.

2.1 Data Sources

2.1.1 Scanner Data

The analysis is primarily based on the Nielsen Homescan Consumer Panel and Nielsen Retail Scanner datasets, which have been widely used in the literature (Einav, Leibtag and Nevo, 2008). This data tracks consumption from 2004 to 2013 at the product level in department stores, grocery stores, drug stores, convenience stores and other similar retail outlets across the US. The data are representative of about 40% of household expenditures on goods and 15% of total household expenditures. Appendix A presents a detailed description of the data sources.

Three features of the data are particularly useful for my analysis. First, product-level data is available on both prices and quantities. Quantity data is rare at the product level (for instance, the BLS does not collect such data), but it is crucial for quality adjustment in price indices.⁹ Second, the Homescan Consumer panel has information on household characteristics such as income, age, education, size, occupation, marital status and zip code. It is therefore possible to directly map products to consumer characteristics. Third,

(2008), Piketty (2013), Song, Price, Guvenen, Bloom and Till von Wachter (2015), Atkinson (2015)), homothetic price indices (Sato (1976), Vartia (1976), Feenstra (1992), Pakes (2003), Broda and Weinstein (2006, 2010), Erickson and Pakes (2011)), non-homothetic price indices (McGranahan and Paulson (2005), Broda and Romalis (2009), Moretti (2013), Diamond (2015), Handbury (2015), Faber and Fally (2015), Argente and Lee (2016), Comin, Lashkari and Mestieri (2015)), innovation in labor markets (Acemoglu (1996, 2002, 2007), Acemoglu and Autor (2011), Autor (2013), Autor and Dorn (2013), Bell, Chetty, Jaravel, Petkova and Van Reenen (2015)), market size effects and endogenous technical change (Acemoglu and Linn (2004)), innovation and inequality in product markets (Schumpeter (1942), Vernon (1966), Matsuyama (2002)), and trade models of monopolistic competition with free entry (Melitz (2003), Melitz and Ottaviano (2008), Zhelobodko, Kokovin, Parenting and Thisse (2012)).

⁹Intuitively, observing shifts in quantities allows me to directly measure substitution patterns (and thus address substitution bias, which is a core concern of the CPI produced by BLS) and to infer the quality of products given their price, their market share, and the demand system. See Section 3 for a complete discussion.

the dataset offers a good measure of product innovations, defined as the introduction of new barcodes, or Universal Product Classification (UPC) codes. It is rare for a meaningful quality change to occur without resulting in a change of UPC¹⁰ and, conversely, for a UPC change to occur without being associated with a quality change that might cause consumers to pay a different price.¹¹ Similarly, discontinued UPCs can be identified.¹²

Nielsen provides a detailed product hierarchy, based on where products are sold in stores. In my sample, about 3 million products (identified by their barcode, or UPC code) are classified into 10 broad departments (dry grocery, general merchandise, health and beauty care, alcoholic beverages, deli, etc.), 125 more detailed product groups (grooming aids, soup, beer, pet care, kitchen gadgets, etc.) and 1,075 very detailed product modules (ricotta cheese, pet litter liners, bathroom scale, tomato puree, women’s hair coloring, etc.). When ranking product modules by mean consumer income¹³, in line with intuition the top five product modules are scotch, natural cheese, gin, fondue sauce and cookware, while the bottom five are tobacco, canned meat, taco filling, insecticide and frozen fruit drinks.

Finally, the data can be disaggregated at the level of 76 local markets, described in Appendix A. According to Nielsen, the dataset is still representative within each of the 76 markets. The data cannot reliably be disaggregated further (e.g. at the county or zip code level).

2.1.2 Additional Datasets: Manufacturer Identifiers, Markups and CPI/CEX Data

A number of datasets in addition to the Nielsen scanner data are used in the analysis. Appendix A describes them in greater detail. First, to measure manufacturer entry and competition, I match manufacturer identifiers from GS1, the company in charge of allocating bar codes in the US, to the UPC code in the Nielsen data, reaching a match rate of 95%. The typical product group is characterized by a few large manufacturers and a competitive fringe of manufacturers with very low market shares. The model developed in Section 5 is consistent with these patterns.

To test additional predictions of the model, I use data on retailer markups in 250 grocery stores, operated by a single retail chain, between January 2004 and June 2007 in 19 U.S. states. To examine whether lower inflation for higher-income households is a secular trend, I match the various expenditure categories of the CEX to 48 item-specific CPI data series going back to 1953 (see Appendix Section A.1.3).

¹⁰First, retailers rely on UPCs for inventory management: if distinct products, offering a different experience to consumers, have the same barcode, then inventory management becomes very difficult. As a result, GS1, the international organization in charge of attributing UPC codes, requires manufacturers to purchase a new barcode following any quality change of a product. Second, I observe a large set of attributes of UPCs in the Nielsen data, e.g. flavor, container description, organic label, scent, etc. I find that this set of characteristics remains constant over time within a UPC.

¹¹Allowing for changing UPC codes for the same product, offering the same experience to consumers, would hamper inventory management by making it difficult for retailers to re-order out-of-stock items. Accordingly, GS1 forbids manufacturers to change UPC codes when there is no quality change. Second, I show in Section 3 that over 50% of new UPCs (weighted by spending) in any given year correspond to entry of a manufacturer that was previously not active in this part of the product space.

¹²Note that these measures of product turnover include any change in products, including those driven by changes in the size of products, their flavor, or other characteristics that can be secondary for the consumer. Nielsen provides identifiers that allows tracking barcodes that are new because of marginal changes, e.g. a change in packaging: all of the results presented in the paper are similar when excluding these products from the definition of “new” products.

¹³For product module l , mean consumer income is defined as $I_l \equiv \sum_n s_{nl} I_n$, where I_n is the income of household n and s_{nl} their share of total spending in l .

2.2 Summary Statistics

Table 1 shows the distribution of spending across the main expenditure categories available in the Nielsen scanner data. Although most of aggregate spending is devoted to food products, a wide variety of product groups are included in the dataset. By examining heterogeneous patterns across these detailed product categories, I can distinguish between various theories that could explain why high-income households experienced a lower inflation rate than low-income households.

Table 1: Distribution of Spending across Nielsen Expenditure Categories

Department	Product Groups	Expenditure Share	Barcode Share
Alcoholic Beverages	beer, liquor, wine	4.4%	3.1%
Dairy	butter and margarine, cheese, sour cream, toppings, dough products, eggs, milk, pudding snacks, spreads, yeast, yogurt	8.8%	3.3%
Dry Grocery	baby food, baking mixes, baking supplies, bread and baked goods, breakfast food, candy, syrup, flour, carbonated beverages, cereal, coffee, condiments, gravies, sauces, cookies, crackers, desserts, gelatins, etc. syrup, flour, canned fruit, dried fruit, gum, jams, jellies, non-carbonated soft drinks, soup, spices, seasoning, sugar, sweeteners, molasses, tea, canned vegetables, etc.	39.9%	29.6%
Fresh Produce	fresh produce	2.6%	1.2%
Frozen Food	frozen baked goods, frozen breakfast foods, frozen desserts, fruits and topping, ice, ice cream, frozen drinks, frozen pizza and snacks, frozen prepared food, frozen seafood and poultry, frozen vegetables	8.5%	4.7%
General Merchandise	automotive, batteries and flashlight, books and magazines, canning, freezing supplies, cookware, electronics, records, tapes, gardening, glassware, tableware, party needs, tools, hosiery, socks, household supplies, appliances, insecticides, pesticides, kitchen gadgets, light bulbs, etc.	8.4%	27.5%
Health and Beauty Aids	deodorant, diet aids, ethnic haba, feminine hygiene, first aid, fragrances, grooming aids, hair care medications, men's toiletries, oral hygiene, sanitary protection, shaving needs, skin care, vitamins	10.8%	16.9%
Non-food Grocery	charcoal, logs, accessories, detergents, disposable diapers, fresheners and deodorizers, household cleaners, laundry supplies, paper products, personal soap and bath additives, pet care, tobacco, wrapping materials and bags	13.4%	12.3%
Packaged Meat	fresh meat, deli packaged meat	3.2%	1.4%

Notes: This table reports summary statistics for the various departments and products groups in the Nielsen Homescan Consumer Panel dataset, from 2004 to 2013. A detailed description of the data source is provided in Sections 2.1.1 and A.1.4. Appendix Table A1 reports additional summary statistics.

The product groups listed in Table 1 may not strike the reader as particularly innovative. Indeed, although some consumer electronics are included, most of the spending is devoted to product categories that are not known for groundbreaking technology innovations in recent decades. However, these product categories are characterized by a relatively high rate of increase in product variety, as further documented in Section 3. The data is therefore ideal to study one particular manifestation of innovation, increasing product variety, and how it benefits households across the income distribution. Appendix A presents more details on the data, with a comparison of aggregate spending share in the Nielsen scanner data, the Consumer Price Index for all urban consumers and the Consumer Expenditure Survey.

3 Measuring Quality-Adjusted Inflation Across Income Groups

I compute quality-adjusted inflation rates across the income distribution, taking into account the welfare gains from increasing product variety. I find that inequality is magnified: annual quality-adjusted inflation is on average 65 basis points lower for households who make above \$100,000 a year, relative to households earning below \$30,000. I show that a large share of the inflation difference between income groups occurs within detailed product categories, which cannot be captured by the existing literature using more aggregated data from the Bureau of Labor Statistics.

3.1 Nonhomothetic Preferences, Product Variety and Inequality

The nonhomothetic nature of preferences means that the baskets of goods and services consumed by households across the income distribution systematically differ. Given that households have a taste for variety, the mapping between nominal income and utility depends on both the quality-adjusted price of products and the number of available varieties. This paper studies how the mapping between nominal income and inequality changes over time.

This section characterizes shifts in the mapping from nominal income to utility at various points of the income distribution using a money metric, the compensating variation. The compensating variation gives the amount of nominal income that one would need to take away from the consumer at the “new” equilibrium to make them indifferent between this new equilibrium (with the new mapping) and the “old” one (with the initial mapping). This approach provides a characterization of changes in purchasing-power inequality. Given the demand system, it is possible to infer the quality of products based on their price and equilibrium market share, and to measure the gains from increasing product variety based on the share of spending on new products. The rest of this section discusses the procedure in detail and shows that the results are robust across price indices, indicating that structural assumptions about the demand system do not drive the results.

I use the term “inflation” to describe my findings throughout the paper because it is an intuitive notion, but my results are invariant to the unit of account. I document changes in the relative prices of goods that cater to high- and low-income households. These relative price changes would be unaffected by shifts in the overall level of inflation; therefore nominal indeterminacy plays no role in my findings.

3.2 Overview of Methodology and Review of Basic Price Indices

The goal is to compute the cost of achieving a certain level of utility in one year relative to the previous year. Such price indices are known as “exact price indices.” The analysis must take into account price changes for continuing products, changes in product variety, as well as the optimizing behavior of consumers who may substitute from one good to another. By definition, this exercise requires taking a stance on a utility function. The role of the utility function is twofold: quantifying the impact on utility of price changes for the goods that exist across periods, but also translating the patterns of product creation and destruction into a welfare metric. To understand what parts of the result are driven by structural assumptions on the utility function, it is useful to split this analysis into two parts, first considering price changes on products that exist across periods and second considering changes in product variety.

First, I consider inflation for the set of products available in two consecutive years, accounting for about 90% of overall spending. The quality of a given product is assumed to be constant over time¹⁴ and data is available on market shares of each product; therefore it is straightforward to compute a price index reflecting product quality and consumers’ substitution behavior. Intuitively, I observe the price change for each product and I only need to decide how to weigh the various products. The exact price index offers a principled way of doing so. The structural assumption on the utility function plays a minor role for the final result, as can be seen by computing standard price indices that do not have an interpretation in terms of utility but can serve as bounds by allowing for an extreme form of substitution (like the Paasche price index, which offers a lower bound on inflation) or making any substitution impossible (like the Laspeyres price index, which offers an upper bound on inflation). To show that the quantitative results on continuing products do not depend on the way substitution effects are handled, I present results for the following price indices¹⁵:

$$\begin{aligned}
 \text{Laspeyres Index : } P_L &\equiv \frac{\sum_{i=1}^n p_i^t q_i^0}{\sum_{i=1}^n p_i^0 q_i^0} = \sum_{i=1}^n \frac{p_i^t}{p_i^0} s_i^0 \\
 \text{Paasche Index : } P_P &\equiv \frac{\sum_{i=1}^n p_i^t q_i^t}{\sum_{i=1}^n p_i^0 q_i^t} = \left(\sum_{i=1}^n \left(\frac{p_i^t}{p_i^0} \right)^{-1} s_i^t \right)^{-1} \\
 \text{Marshall – Edgeworth Index : } P_{ME} &\equiv \frac{\sum_{i=1}^n p_i^t (q_i^t + q_i^0)}{\sum_{i=1}^n p_i^0 (q_i^t + q_i^0)} \\
 \text{Walsh Index : } P_W &\equiv \frac{\sum_{i=1}^n p_i^t \sqrt{q_i^t q_i^0}}{\sum_{i=1}^n p_i^0 \sqrt{q_i^t q_i^0}} \\
 \text{Fisher Index : } P_F &\equiv \sqrt{P_L P_P} \\
 \text{Geometric Laspeyres Index : } P_L^G &\equiv \Pi_{i=1}^n \left(\frac{p_i^t}{p_i^0} \right)^{s_i^0} = \exp \left(\sum_{i=1}^n s_i^0 \cdot \log \left(\frac{p_i^t}{p_i^0} \right) \right) \\
 \text{Geometric Paasche Index : } P_P^G &\equiv \Pi_{i=1}^n \left(\frac{p_i^t}{p_i^0} \right)^{s_i^t} = \exp \left(\sum_{i=1}^n s_i^t \cdot \log \left(\frac{p_i^t}{p_i^0} \right) \right) \\
 \text{Tornqvist Index : } P_T &\equiv \Pi_{i=1}^n \left(\frac{p_i^t}{p_i^0} \right)^{\frac{s_i^t + s_i^0}{2}} = \exp \left(\sum_{i=1}^n \frac{s_i^0 + s_i^t}{2} \cdot \log \left(\frac{p_i^t}{p_i^0} \right) \right)
 \end{aligned}$$

¹⁴The assumption that quality is constant at the UPC level was justified in Section 2, with a description of institutional details (rules to grant new barcodes set by GSI and inventory management system used by retailers) and empirical exercises (showing that the set of available characteristics such as flavor, label and scent, are stable within UPC codes).

¹⁵ i indexes barcodes, t time, q quantities, p prices, and s spending shares. See Appendix B for a discussion of chaining.

Second, I follow standard techniques in the literature to provide an adjustment to the price index depending on the rate of increase in product variety. By definition, for new and discontinued products price changes across years are not available. Intuitively, given that consumers have a taste for variety, an increase in the range of available products should lead to a decrease in the price index. Translating the increase in product variety into welfare gains requires structural assumptions. I use two standard frameworks: nested CES utility (presented in the next subsection) and nested translog utility (presented in Appendices B.3 and C.3). In both of them, higher-income groups benefit more from the dynamics of product creation and destruction. Because the estimated elasticities of substitution of products within modules are large, the gains from increasing product variety turn out to be largely reflected in price changes for existing products (this is shown formally in the next subsection). The patterns of product creation and destruction matter through competition effects in general equilibrium, such that their welfare effect is almost entirely taken into account in price changes for products existing across periods.

3.3 An Exact Price Index for Nested CES with New Products

My preferred estimation approach follows a well-established literature in trade and macroeconomics and computes quality-adjusted inflation using a nested-CES utility function. The key insight is that this utility function yields a simple expression for the price index, which can be written only in terms of prices and market shares even when goods are constantly being replaced.

The estimation framework builds on Feenstra (1994) and Broda and Weinstein (2006, 2010). I split the analysis using three representative agents, one for households making less than \$30,000 a year, one for households making between \$30,000 and \$100,000 a year, and one for households making above \$100,000. Preference parameters in my estimation framework are a flexible function of the income level, which allows for nonhomotheticities.

The remainder of this subsection shows how to derive and estimate the price index for any representative agent. I assume a nested CES utility function. Product groups are indexed by g and G is the set of all product groups. $\sigma = \rho/(\rho - 1)$ is the elasticity of substitution between product groups. The upper level utility function is:

$$U = \left(\sum_{g \in G} (C_{gt})^\rho \right)^{\frac{1}{\rho}}$$

Composite consumption within a product group is given by:

$$C_{gt} = \left(\sum_{m \in M_g} (c_{mgt})^{\rho_g} \right)^{\frac{1}{\rho_g}}$$

where $\sigma_g = \rho_g/(\rho_g - 1)$ is the elasticity of substitution between product modules within product group g .

$$c_{mgt} = \left(\sum_{u \in U_{mgt}} (d_{umgt} c_{umgt})^{\rho_m} \right)^{\frac{1}{\rho_m}}$$

where c_{umgt} is the quantity of UPC u consumed in product module m and product group g in period t . $\sigma_m = \rho_m/(\rho_m - 1)$ is the elasticity of substitution between UPCs within product module m . d_{umgt} is unobserved and reflects the quality of the UPC.

The minimum unit cost function of the subutility function at the product module level is:

$$P_{mgt} = \left(\sum_{u \in U_{mgt}} \left(\frac{p_{umgt}}{d_{umgt}} \right)^{\sigma_m} \right)^{\frac{1}{\sigma_m}}$$

The minimum cost function at the product group level is:

$$P_{gt} = \left(\sum_{m \in M_g} (P_{mgt})^{\sigma_g} \right)^{\frac{1}{\sigma_g}}$$

And the overall price index is given by:

$$P_t = \left(\sum_{g \in G} (P_{gt})^\sigma \right)^{\frac{1}{\sigma}}$$

Consumer optimization also yields:

$$s_{umgt} = \left(\frac{p_{umgt}/d_{umgt}}{P_{mgt}} \right)^{1-\sigma_m}$$

i.e. the quality adjusted price can be backed out as follows:

$$\ln \left(\frac{p_{umgt}}{d_{umgt}} \right) = \frac{\ln(s_{umgt})}{1 - \sigma_m} + \ln(P_{mgt})$$

The key insight for estimation is that the share of consumption of UPC u depends directly on the quality-adjusted price. We can write the price index only in terms of prices and market shares even when goods are constantly being replaced.

Under the assumption that product quality is constant over time ($d_{umgt} = d_{umgt-1}$) and ignoring the introduction of new products, the exact price index of the CES utility function for product module m within product group g is as in Sato (1976) and Vartia (1976):

$$P_{mg}(p_{mgt}, p_{mgt-1}, x_{mgt}, x_{mgt-1}, I_{mg}) = \prod_{u \in I_{mg}} \left(\frac{p_{umgt}}{p_{umgt-1}} \right)^{w_{umgt}} \quad (1)$$

$$w_{umgt} = \frac{(s_{umgt} - s_{umgt-1})/(\ln(s_{umgt}) - \ln(s_{umgt-1}))}{\sum_{c \in I_{mg}} (s_{umgt} - s_{umgt-1})/(\ln(s_{umgt}) - \ln(s_{umgt-1}))} ; \quad s_{umgt} = \frac{p_{umgt}x_{umgt}}{\sum_{u \in I_{mg}} p_{umgt}x_{umgt}}$$

where $I_{mg} = I_{mgt} \cap I_{mgt-1}$ is the set of varieties consumed in both periods t and $t - 1$. x_{mgt} and x_{mgt-1} are the cost-minimizing quantity vectors of products within module m in each of the two periods. A remarkable feature is that the price index does not depend on the unknown quality parameters d_{umgt} . We only need to compute the geometric mean of the individual variety price changes, where the weights are ideal log-change weights w_{umgt} . These weights are computed using spending shares in the two periods and are always bounded between the shares of spending in the t and $t - 1$.

With introduction of new varieties and exit of some old varieties, as shown in Feenstra (1994) the exact price index for product module m within product group g is given by:

$$\pi_{mg}(p_{mgt}, p_{mgt-1}, x_{mgt}, x_{mgt-1}, I_{mg}) = P_{mg}(p_{mgt}, p_{mgt-1}, x_{mgt}, x_{mgt-1}, I_{mg}) \cdot \left(\frac{\lambda_{mgt}}{\lambda_{mgt-1}} \right)^{\frac{1}{\sigma_m - 1}} \quad (2)$$

$$\lambda_{mgt} = \frac{\sum_{u \in I_{mg}} p_{umgt} x_{umgt}}{\sum_{u \in I_{mgt}} p_{umgt} x_{umgt}} ; \quad \lambda_{mgt-1} = \frac{\sum_{u \in I_{mg}} p_{umgt-1} x_{umgt-1}}{\sum_{u \in I_{mgt-1}} p_{umgt-1} x_{umgt-1}}$$

This result states that the exact price index with variety change is equal to the “conventional” price index multiplied by an additional term, which captures the role of new and disappearing varieties.¹⁶ The higher the expenditure share of new varieties, the lower is λ_{mgt} and the smaller is the exact price index relative to the conventional price index. An intuitive way to rewrite this ratio is as follows:

$$\frac{\lambda_{mgt}}{\lambda_{mgt-1}} = \frac{1 + \text{Growth Rate of Spending on Overlapping Products}_{mgt}}{1 + \text{Growth Rate of Total Spending}_{mgt}}$$

which clearly shows that a net increase in product variety (weighted by spending) drives the price index down. The price index also depends on the module-specific elasticity of substitution between varieties σ_m . As σ_m grows, the additional term converges to one and the bias goes to zero. Intuitively, when existing varieties are close substitutes to new or disappearing varieties, a law of one price applies and price changes in the set of existing products perfectly reflect price changes for exiting and new varieties.

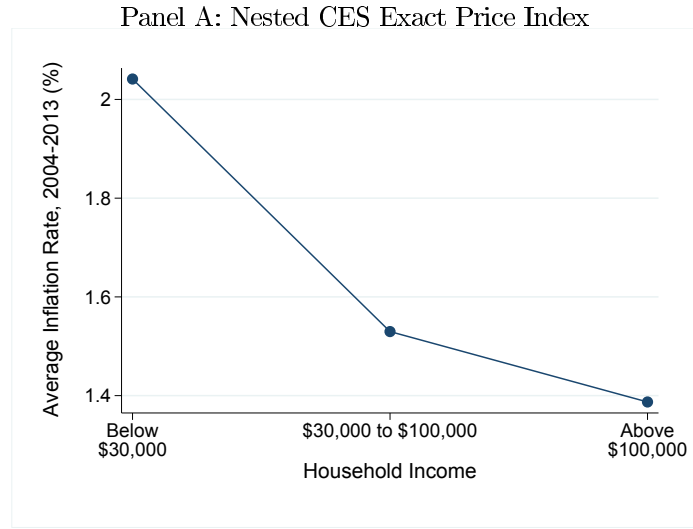
To compute the price index shown by equation (2), a high-dimensional set of elasticities of substitution $\{\sigma_m\}$ must be estimated. In practice, I conduct estimation separately for each income group to allow for non-homotheticities. The main challenge for estimation is that demand and supply parameters must be obtained using only information on prices and quantities. The insight of Feenstra (1994) is that although one cannot identify supply and demand, the data conveys information about the joint distribution of supply and demand parameters: the constant elasticity assumption is essentially sufficient for identification. Due to space constraints, the derivation of the estimation equations and identification conditions is presented in Appendix B.2.

3.4 Inflation across Income Groups for Products Available in Consecutive Years


Panel A of Figure 1 shows the average inflation between 2004 and 2013 on the set of continued products (defined as products that are available in consecutive years) for households across the income distribution. Inflation is computed using the exact price index for the nested CES utility function described in the previous subsection (without the adjustment for new and disappearing products, which is examined later in this section and does not affect the results much). The inflation rate is 0.65 pp lower for households making more than \$100,000 a year, relative to households making less than \$30,000.

¹⁶In principle, the result presented in equation (2) can be used to compute price indices adjusted for increasing product variety over any time horizon. However, two factors make some time horizons more sensible than others in practice. First, it is useful to define periods in years to prevent seasonal factors from driving product turnover. UPCs will be considered destroyed only if they were not purchased at any time during a year-long period. Second, one needs to decide how many years should separate the two periods. While this choice is inherently arbitrary, I decided to present results based on one-year intervals, considering other intervals in robustness checks.


Figure 1: Inflation for Continued Products across Income Groups



Panel B: Stability of Inflation Difference across Price Indices



Panel C: Nested-CES Exact Price Index across Age-Income Groups



Notes: Panels A to C report the average inflation rate for various household groups from 2004 to 2013. In any given year, the sample includes all barcodes observed in the current and previous year. The price indices are described in Sections 3.2 and 3.3. Appendix Tables C5, C6, C7, C8 and C9 show the robustness of these results.

As shown in Panel B of Figure 1, similar results are obtained when considering any of the price indices introduced in Subsection 3.2. In addition, this panel reports the inflation difference when re-defining products as UPCs available in the same store, or as UPCs available in the same local market (see Appendix A.1.4 for a map of local markets). The results with this new definition of products are very similar. Overall, across all price indices and product definitions, the inflation rate is always between 0.56pp and 0.72pp lower for households making more than \$100,000 a year, relative to households making less than \$30,000. Panel C of Figure 1 shows that these results are robust when considering other income groups and when repeating the analysis within age groups. For each age group, inflation is systematically lower for higher-income households.¹⁷

3.5 Changes in Product Variety across Income Groups

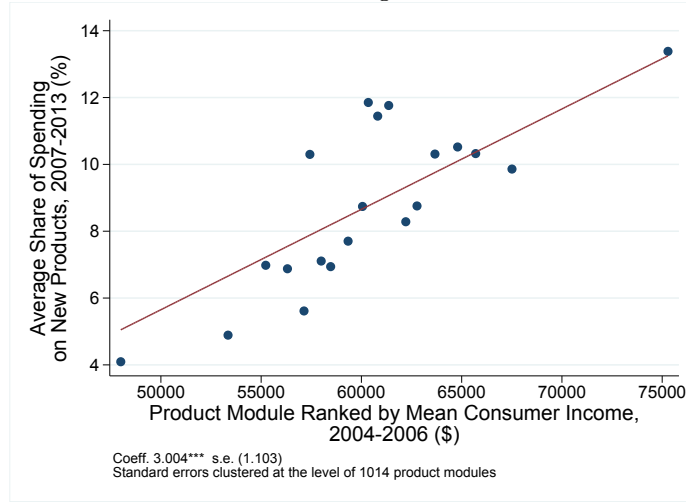
Do welfare effects from increasing or decreasing product varieties also differ across income groups? I find that the rate of increase in product variety is faster in product modules catering to higher-income households, implying that higher-income households benefit more from increasing product variety. Panel A of Figure 2 shows this effect in an intuitive way by using the share of spending on new products (defined as barcodes which did not exist in the previous year) as a measure of the flow of successful product innovations. For every \$10,000 increase in the mean income of the consumers buying from a product module, the share of spending on new products in this product module goes up by 3 percentage points, a large change equal to approximately a third of the average share of spending on new products. Plotting the data in this way, through the lens of the product space rather than by directly looking at the consumption baskets of consumers of different income levels, has the key advantage that the “product cycle” will not mechanically generate differences across income groups. In other words, the fact that new products may first be purchased by higher-income consumers will not generate an increasing relationship between income and share of spending on new products, given that we are looking at patterns *across* product modules while the product cycle operates *within* product modules.

The patterns of product destruction are relatively homogeneous across product modules, regardless of consumer income. In other words, the share of spending on new products is a good proxy for the increase in product variety. Panel A of Appendix Figure C1 shows this directly by plotting the total increase in barcodes across product modules: the rate of increase in the total number of varieties goes up by one percentage point with a \$10,000 dollar increase in the income of the representative consumer. Moreover, Panel B of Appendix Figure C1 plots the welfare-relevant metric that captures the benefits of increasing product variety in the nested CES demand system introduced earlier. Across product modules, the ratio $\frac{\lambda_{mgt}}{\lambda_{mgt-1}} = \frac{1 + \text{Growth Rate of Spending on Overlapping Products}_{mgt}}{1 + \text{Growth Rate of Total Spending}_{mgt}}$ decreases with consumer income, which confirms that higher-income consumers benefit more from product innovations. Similar results hold for other measures of “new products” - new UPCs relative to two, three or four years ago, as well as new brands and new manufacturers, as shown in Appendix Figure C2.

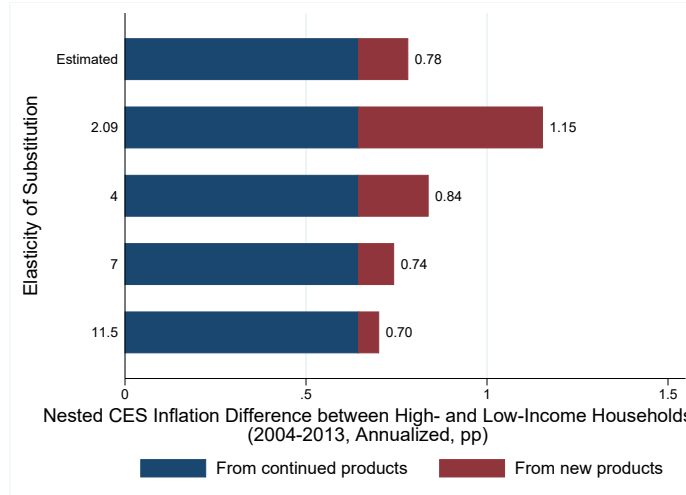
¹⁷In a companion paper, Jaravel (2016) investigates patterns of inflation and product innovations across the age distribution.

Figure 2: Welfare Gains from Changes in Product Variety across Income Groups

Panel A: New Products Benefit Higher-Income Households More



Panel B: Difference in Welfare Gains under Nested CES Utility



Notes: Panel A shows the relationship across product modules between the share of spending on new products (defined in any given year as barcodes that did not exist in the previous year) and the mean (spending-weighted) consumer income. Panel B shows the difference in inflation rates experienced by households earning above \$100,000 (high income) and below \$30,000 (low income), using the exact price index with new products derived in Section 3.3.

Panel B of Figure 2 brings together the patterns of increasing product variety on inflation for continued products, using the formula in equation (2). The results are shown for various values of the within-module elasticities of substitution. With my set of estimated (income-group-specific) within-module elasticities of substitution, I find that between 2004 and 2013, on average, annual quality-adjusted inflation was 78 basis point lower for households earning above \$100,000 a year, relative to households earning below \$30,000 a year. Changes in product variety benefited higher-income households more and contributed another 13 basis points to the inflation difference for continued products of 65 basis points. The distribution of within-module elasticities of substitution and the ratios $\frac{\lambda_{mgt}}{\lambda_{mgt-1}}$ are reported in Appendix Table C2. The elasticities are

relatively high, with a median of 5.5, and are slightly smaller for high-income households.¹⁸ The high values of the estimated elasticities imply that the “product variety” adjustment is relatively small: most of the welfare effects are captured by the inflation difference on goods that exist across consecutive years.¹⁹

Panel B of Figure 2 shows the sensitivity of the quantitative results to various values of the elasticity, based on the estimates obtained by other papers in the literature. A small elasticity of 2.09, as in Handbury (2013), implies a large differential welfare gains to the benefit of high-income households. But a high elasticity of 11.5, as in Broda and Weinstein (2010), implies a small effect. In the marketing literature (Gordon et al., 2013), elasticities in the retail sector are found to be between 4 and 7, which implies modest differential welfare gains from new products. Although the exact values differ, the qualitative finding is the same: new products benefit higher-income households more and the inflation difference for continued products can serve as a lower bound for the full welfare difference between high- and low-income households.²⁰

In sum, due to the high elasticities of substitution within product modules, the patterns of increase in product variety do not matter for the *measurement* of quality-adjusted inflation: they have a small direct effect on the price index. However, increasing product variety may be a fundamental *mechanism* explaining why the price index rises more slowly for higher-income households, because new products compete with existing products and can thus have an indirect effect on the price index. In Section 4, I find strong support for this hypothesis.

3.6 Decompositions

3.6.1 Results

It is possible to decompose the inflation difference between households at different points of the income distribution. For the purpose of this exercise, I focus on comparing households making more than \$100,000 a year to households making less than \$30,000 a year. The inflation difference reflects the combined effects of both price and quantity changes, as well as baseline differences in spending patterns across income groups. For instance, it could be that high-income households spend more on fresh produce and that inflation tends to be lower in this broad item category. Alternatively, it could be the case that high-income households experience different inflation rates compared with low-income households on the same barcodes, for instance because they shop at different stores or have different propensities to use coupons. Accordingly, the inflation difference between high income and low-income households can be decomposed into a “between” component and a “within” component. The “between” component corresponds to the inflation difference that would prevail if households differed only in terms of their expenditure shares across items categories and experienced the same inflation rate within each item category. The “within” component corresponds to the inflation difference

¹⁸The magnitude of these elasticities is consistent with markups in the retail sector, which are between 30% and 40% in the Census of Retail Trade.

¹⁹Indeed, from the derivation in subsection 3.2, quality adjusted inflation is given by $\pi_{mg} \equiv P_{mg} \cdot \left(\frac{\lambda_{mgt}}{\lambda_{mgt-1}} \right)^{\frac{1}{\sigma_m - 1}} \rightarrow P_{mg}$ as $\sigma_m \rightarrow \infty$. Intuitively, a law of one price applies as the elasticity of substitution increases.

²⁰In general, the estimated elasticities of substitution tend to be smaller in papers using the Hausman-type IV approach (e.g. Hausman & Leibtag (2007), Handbury (2013)) and larger in empirical work using the Feenstra (1994) approach for estimation (e.g. Broda & Weinstein (2010), Hottman et al. (2014), and this paper).

that would prevail if households differed only in terms of the inflation rate they experience within an item category and had the same expenditure shares across categories. Formally, for any grouping of products G , the inflation difference between high- and low-income households can be decomposed as follows²¹:

$$\pi^R - \pi^P \equiv \sum_G s_G^R \pi_G^R - \sum_G s_G^P \pi_G^P = \underbrace{\left(\sum_G s_G^R \pi_G^R - \sum_G s_G^P \pi_G^P \right)}_{\text{Between}} + \underbrace{\sum_G \bar{s}_G (\pi_G^R - \pi_G^P)}_{\text{Within}} \quad (3)$$

with s_G^i denoting the share of spending of income group i on product grouping G and π_G^i the inflation experienced by income group i in product grouping G . π_G and \bar{s}_G denote the average inflation rate and the average spending shares for product grouping G , respectively.

Table 2: Decompositions of Differences between High- and Low-Income Households

Panel A: Inflation for Continued Products		
Aggregation Level (Broad to Narrow)	Inflation Difference pp	% of Actual
Department	0.06	8.6
Product Group	0.14	21.4
Product Module	0.28	42.8
UPC	0.476	72.2
UPC-Local Market	0.520	78.8
UPC-Store	0.607	92.1

Panel B: Share of Spending on New Products	
Aggregation Level (Broad to Narrow)	Difference in Share of Spending on New Products (% of Actual)
Department	1.8
Product Group	29.0
Product Module	39.2

Notes: Panel A shows the decomposition of the inflation difference for continued products between households making above \$100,000 (high income) and below \$30,000 (low income) from 2004 to 2013. In any given year, the sample includes all barcodes observed in the current and previous year. Each row reports the “between” component in equation (3) for a given level of aggregation (described in Section 2.1.1). Panel B shows the decomposition of the difference in spending shares on new products between high- and low-income households, from 2004 to 2013. The sample includes all barcodes and each row reports the “between” component in equation (4) for the relevant level of aggregation.

Panel A of Table 2 reports the results of the decomposition at the following levels of aggregation: department, product group, product module, UPC, UPC in a given local market, and UPC in a given store. Inflation is directly observed at the product level for the last three categories, and the definitions of inflation for categories at levels of aggregation above the UPC are given in subsection 3.2. Perhaps not surprisingly, less than 10% of the difference in the inflation rates experienced by high- and low-income households is

²¹Diewert (1975) shows the validity of this decomposition for a large number of price indices.

due to differences in spending across broad departments. More surprisingly, less than 25% of the inflation difference results from different spending patterns across the 125 detailed product groups, and less than 45% of the difference stems from spending patterns across the 1,075 very disaggregated product modules. More than 70% of the inflation difference occurs between UPCs. This is a large share of the overall difference in inflation rates, but a substantial fraction of the difference still occurs *within* UPCs. To assess the mechanism at play, I repeat the decomposition at the level of UPCs in a given local market, which brings the share of the “between” component close to 80%, as well as at the level of UPCs in a given store, which brings the share of the “between” component to 92%.²² Taken together, these results show that most of the difference in inflation rates between high- and low-income households occurs across UPCs, and that some of the effect results from differential price dynamics for the same UPC across stores. The quality ladder plays a key role in the decomposition: close to the entirety of the inflation difference between income groups across UPCs occurs at the level of “product modules by price decile” cells, as shown in Appendix Table C1.

In a way analogous to the exercise conducted for inflation, the difference in the share of spending on new products between high- and low-income consumers can be decomposed at various levels of aggregation. Formally, for any grouping of products G , the decomposition is as follows:

$$N^R - N^P \equiv \sum_G s_G^R N_G^R - \sum_G s_G^P N_G^P = \underbrace{\left(\sum_G s_G^R N_G - \sum_G s_G^P N_G \right)}_{\text{Between}} + \underbrace{\sum_G \bar{s}_G (N_G^R - N_G^P)}_{\text{Within}} \quad (4)$$

with s_G^i the share of spending of income group i on product grouping G and N_G^i the share of spending on new products for income group i in product grouping G . N_G and \bar{s}_G denote the average share of spending on new products and the average spending shares for product grouping G , respectively. Panel B of Table 2 shows that the difference between the shares of spending on new products between high- and low-income consumers largely occurs within product modules. This pattern is very similar to the inflation decomposition shown in Panel A and provides preliminary evidence that there is a tight connection between the patterns of inflation and product innovations.

3.6.2 Relevance for the Methodology of Statistical Agencies

Table 2 indicates that product-level data is needed to capture the magnitude of the difference in inflation rates between households at different points of the income distribution. It is not sufficient to simply reweigh aggregate price series based on income-specific spending shares, even when the level of aggregation is as detailed as product modules. Yet this is precisely the approach followed by the BLS and other statistical agencies. More specifically, the BLS collects prices on 305 different item categories, known as “entry-level

²²Note that the “within UPC” component of the inflation difference between high- and low-income households is difficult to interpret from a welfare perspective, because households can exert search effort - thus incurring a utility cost - to get a better price for a given UPC. Moreover, the Nielsen data is less reliable to document variation in prices paid by different income groups for the same UPC. Indeed, Nielsen often automatically enters the price of the UPC based on the store the panelist reported for their shopping trip. Because most of the inflation difference exists across UPCs, and because the within-UPC patterns have ambiguous welfare implications and are less precisely measured, I focus on the between-UPC patterns in the remainder of the paper.

items” (ELI). Most of these item categories are very coarse. 230 of them are actually in the retail sector, where the level of disaggregation is much higher than in other sectors. Still, this level of aggregation is too high to capture the bulk of the difference between high and low income consumers. This explains why the result presented here may appear inconsistent with the existing literature, which has found small differences between high and low income consumers.

For instance, McGranahan and Paulson (2005) compute income-specific inflation rates based on between-ELI inflation differences and income-specific CEX spending patterns. Using their data, I computed that between 2004 and 2013 the annualized inflation difference for households in the bottom vs. top income quartiles was 0.18 percentage points, which is similar to what I obtained in the Nielsen data with the “between product group” methodology (see Appendix C for details).

Therefore, the conventional wisdom that inflation is not very different across income groups is likely to be misplaced. Statistical agencies like BLS collect data at a broad level of aggregation, which biases the estimate of the difference in inflation across income groups towards zero. Using the Nielsen data, I directly show that the magnitude of this bias is large in the retail sector. Appendix Table C1 shows that a large share of the inflation difference across income groups could be captured by segmenting each of the product modules by price deciles: the confidential micro data collected by statistical agencies like the BLS could be used to replicate this approach, in the retail sector as well as in other sectors.²³

3.6.3 Related Literature

My results are consistent with Argente and Lee (2016). In parallel work, they study inflation across income groups during the Great Recession, find that it is lower for higher-income households, and argue that this effect is driven by substitution patterns. The inflation dynamics I describe in this paper are more general and of a different nature: I show that the difference in inflation rates across income groups is in fact a secular trend, extending well beyond the crisis and continuing to hold even when substitution effects are ignored (indeed, Figure 1 shows that the magnitude of the inflation difference is similar across a variety of price indices that do not allow for substitution, like the Laspeyres index).²⁴

Two other recent papers are closely related to my findings. Pisano and Stella (2015) document that lower-income households pay lower prices than higher-income households for the same products, primarily because they shop more at discount stores. In contrast, I focus on changes in income-specific price indices over time and use the demand system to provide a measure of quality-adjusted inflation. Faber and Fally (2015) explore the implications of firm heterogeneity for household price indices across the income distribution. They find that larger, more productive firms endogenously sort into catering to the taste of wealthier households, and that this gives rise to asymmetric effects on household price indices in their structural model. I provide direct

²³One would then need to infer the spending shares of various income groups along price deciles, which could be done for instance by estimating “quality Engel curves” as in Bils and Klenow (2001).

²⁴Note that both my results and the results of Argente and Lee (2016) appear inconsistent with the findings of Broda and Romalis (2009), who also use Nielsen data and report in an unpublished manuscript that they find that inflation is lower for lower-income households. In a recent working paper, Kaplan and Schulhofer-Wohl (2016) examine inflation at the household level using Nielsen data and confirm that inflation is lower for *higher*-income households.

evidence of differences in inflation rates across income groups and, in Section 4, I show that they are driven by a distinct explanatory mechanism, the supply response to market size effects.

To the best of my knowledge, my paper is the first to measure the difference in inflation rate between high- and low-income households using Nielsen data for a long period of time, to propose decompositions of this difference, and finally to relate these patterns to the dynamics of supply.

3.7 Robustness Checks

Table 3 shows the robustness of the difference between the inflation rates of high- and low-income households: it exists before, during and after the Great Recession,²⁵ and it is not driven by any single department. Appendix C presents various additional robustness tables and figures. First, Tables C5 and C6 describe the level of inflation for various cuts of the income distribution, various price indices and various periods. Figure C5 summarizes this information and shows that the difference in inflation rates is very robust: higher-income households consistently experienced a lower inflation rate. Second, I re-define products to be UPCs available in a given local market (Table C7) or UPCs available in a given store (Table C8) and show that the results continue to hold.

Table 3: Robustness of the Difference in Inflation for Continued Products between Income Groups

Period	Excluded Department	Average Annual Inflation Difference between High- and Low-Income Households
2004-2013	None	0.654
2004-2006	None	0.472
2011-2013	None	0.529
2004-2013	Health and beauty care	0.689
2004-2013	Dry grocery	0.738
2004-2013	Frozen food	0.690
2004-2013	Dairy	0.649
2004-2013	Deli	0.657
2004-2013	Packaged meat	0.654
2004-2013	Fresh produce	0.655
2004-2013	Non-food grocery	0.534
2004-2013	Alcohol	0.638
2004-2013	General merchandise	0.631

Notes: This table reports the inflation difference for continued products between households making above \$100,000 (high income) and below \$30,000 (low income) across subsamples. In any given year, barcodes that are not observed in both the current and previous year are excluded. Appendix Tables C5 to C13 report additional robustness checks.

Appendix C presents many additional robustness checks. First, I document the robustness of the result that new products benefit higher-income consumers more in Appendix Section C.3. I show that the results are similar when defining a “new” product as a UPC code from a new manufacturer, which did not exist at all in previous years or was active in a different part of the product space. I also show that the results are robust to valuing new products using the approximation proposed by Hausman (2003) and by using a

²⁵The difference in inflation rates appears to be larger during the Great Recession. Argente and Lee (2016) argue that the way in which consumers adjusted their shopping behavior to mitigate the crisis can explain the difference in the inflation rates across income groups during this period.

translog expenditure system, which addresses the concern that consumers may perceive varieties as less and less differentiated as the product space gets more crowded. Second, Appendix Section C.4 shows that the results on inflation across income groups are similar when using price information from the Nielsen Retail Scanner dataset, a point-of-sale dataset which addresses the concern that prices may be mismeasured in the Homescan Consumer Panel. Third, the results are not driven by the “product cycle”: the findings do not stem from products that are introduced in the market at a very high price, are initially purchased by high-income households, and start being purchased by lower-income households only after they have converged to their long-run level.²⁶ Fourth, the observed inflation difference between high- and low-income households does not result from price convergence: it is not the case that high-income households initially pay a higher price for the same UPCs as low-income households, and that price then converges to the same level for all households in future periods. Finally, similar results are obtained under a variety of changes to the sample and variables of interest: using alternative measures of household income, adjusting for sample variable across income groups, extending the sample back to 1999 for food products, using unchained price indices, and conducting the analysis at the quarterly level.

4 The Equilibrium Response of Supply to Changes in Demand

In this section, I show that the supply response to market size effects induced by shifts in the income distribution explains why higher-income households experienced a faster increase in product variety and lower inflation from 2004 to 2013. Intuitively, over the sample period demand for premium products increased relative to demand for entry-level products, because of both growth and rising inequality. In response, suppliers directed their product innovation efforts towards premium market segments, which in turn led to increased competitive pressure and lower inflation for products in these market segments. I first present a series of correlations in line with this hypothesis. Second, I estimate the causal impact of a demand shock on product innovations and price dynamics, finding large effects. I then apply these point estimates to changes in demand induced by shifts in the income distribution: the implied inflation patterns are very close to the inflation patterns actually observed across the income distribution. Finally, I present additional evidence on the nature of the supply response, which helps discipline the model developed in the final part of the paper.

4.1 Descriptive Evidence

A variety of patterns in the data are in line with the notion that supply responds to changes in demand and that this process primarily benefits higher-income households. First, Panel A of Figure 3 documents that product modules that grow faster are characterized by a faster-increase in product variety and by lower inflation for continued products. Moreover, product modules catering to higher-income households have grown faster during the sample period.

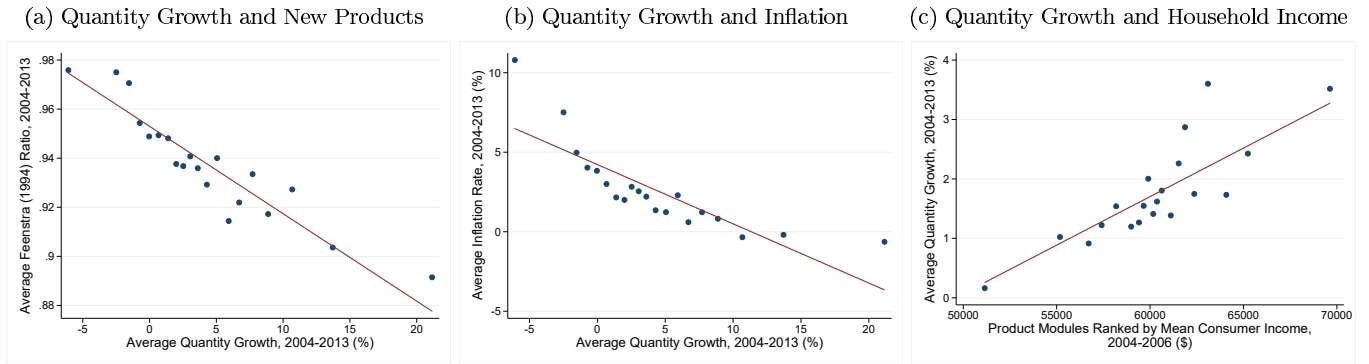
Second, Panel B of Figure 3 shows that the patterns of product introductions and low inflation for

²⁶Similarly, the findings are not driven by the “fashion cycle”, as discussed in Appendix C.1.

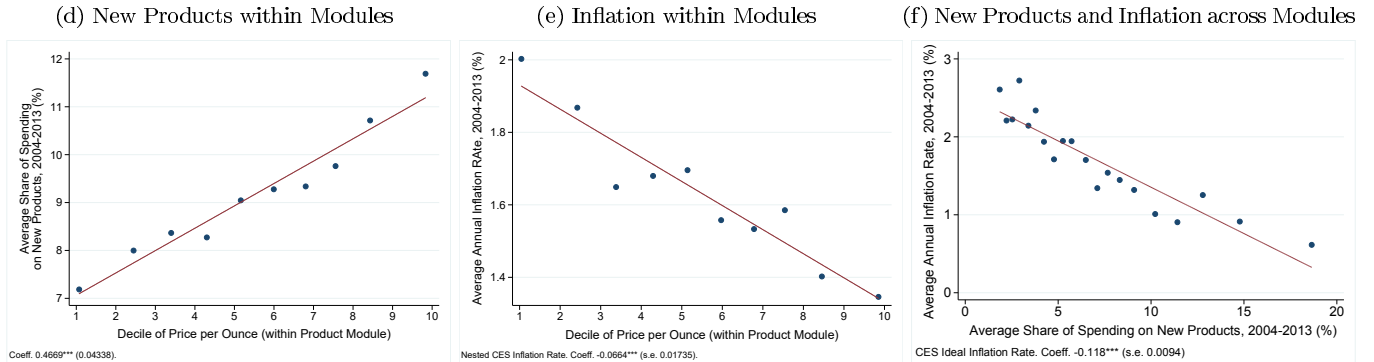
continued products go hand in hand, both across and within product modules. Within product modules, the quality ladder plays an important role. As reported in this panel, there is more product entry and lower inflation for products that belong to higher price deciles. The price deciles are computed within each module based on the average (spending-weighted) unit price of the products that are available in consecutive years. This approach provides a way to segment the product space even within product modules, the highest level of disaggregation provided by Nielsen. Prices in both start and end period are used to classify the UPC across price deciles, so that the classification is not subject to mean reversion. Prices are adjusted for the weight of the item in order to provide a more accurate measure of the unit price. Appendix Figure D1 provides a robustness check using information on the brand of each UPC. In that figure, the deciles are not based on the price of the UPC itself, but rather on pricing behavior at the brand level over the entire dataset. The results are identical to Panel B of Figure 3, which confirms that mean reversion is not driving the results.

Figure 3: Descriptive Evidence

Panel A: Quantity Growth, New Products, Inflation and Household Income across Modules



Panel B: New Products and Inflation Within and Across Modules



Notes: Figures (a), (b), (c) and (f) report the best-fit lines of OLS regressions across 1,075 product modules, as well as binned scatter plots (each dot represents 5% of the data). Figure (a) shows the relationship between changes in product variety (measured by the Feenstra (1994) ratio, derived in Section 3.3) and the growth rate of quantities (using the exact quantity index for the nested CES demand system of Section 3.3). Figure (b) shows the relationship between the inflation rate for continued products (defined in each year as all barcodes observed in the current and previous year) and the growth rate of quantities. Figure (c) shows the relationship between the growth rate of quantities and mean (spending-weighted) consumer income. Figure (f) shows the relationship between the inflation rate for continued products and the share of spending on new products. Figures (d) and (e) present the best-fit lines from OLS regressions at the level of 10,750 product modules by price deciles, with product module fixed effects (the price deciles are built as described in Section 4.1). These figures also report the mean for each (within-module) decile. All regressions are weighted by spending.

Panel B of Figure 3 documents that the negative correlation between inflation and the share of spending on new products *across* product modules is a key feature of the data. A decomposition exercise shows that this correlation is strong enough to explain a large fraction of the inflation patterns across income groups documented in Section 3.²⁷ For any product grouping G , the inflation difference between income groups can be decomposed according to (3), with s_m^i denoting the share of spending of income group i on G and π_G the average inflation rate in G . The “between” component can be decomposed further to examine how much of the inflation difference across categories is explained (predicted) by differences in shares of spending on new products across categories.

$$\underbrace{\left(\sum_G s_G^R \pi_G - \sum_G s_G^P \pi_G \right)}_{\text{“between” component}} = \underbrace{\left(\hat{\pi}_G^R - \hat{\pi}_G^P \right)}_{\text{explained}} + \underbrace{R}_{\text{unexplained}} \quad (5)$$

with

$$\begin{aligned} \hat{\pi}_G^R - \hat{\pi}_G^P &= \sum_G \left(\hat{\beta} X_G \cdot (s_G^R - s_G^P) \right) \\ R &= \sum_G \left(\epsilon_G \cdot (s_G^R - s_G^P) \right) \\ \pi_G &= \beta X_G + \epsilon_G \end{aligned}$$

where X_G is share of spending on new products in G . $\hat{\beta}$ is the OLS estimate of β and ϵ_G the estimated residual. This procedure estimates the extent to which the difference in inflation rates between high- and low-income households results from the fact that high-income consumers tend to devote a higher share of their spending to product categories where the rate of product innovations is higher (i.e. moving to the right along the x-axis in Figure 3 (f)), or from the fact that high-income households tend to spend more on product categories with a lower share of inflation, holding the rate of product innovations constant (i.e. moving down the y-axis in Figure 3 (f)). Table 4 shows that for the various levels of aggregation, around half of the inflation difference between high- and low-income households is explained by differences in patterns of product innovations.²⁸

Table 4: Explaining the Inflation Difference Between High- and Low-Income Households

Aggregation Level (Broad to Narrow)	Share of Inflation Difference Explained by Spending on New Products (% of Actual)
Department	40.9
Product Group	58.3
Product Module	51.3

Notes: This table reports the share of the inflation difference for continued products between households making above \$100,000 (high income) and below \$30,000 (low income), between 2004 and 2013, that can be explained by spending on new products according to the methodology in equation (5). Each row reports the results from the analysis at various levels of aggregation (10 departments, 125 product groups, 1,075 product modules). In any given year, the inflation rate for continued products is computed using all barcodes observed in the current and previous year.

²⁷This exercise is similar in spirit to the reweighting technique introduced in DiNardo, Fortin and Lemieux (1996).

²⁸Note that any measurement error (e.g. UPC relabeling that does not reflect a true product innovation) will bias this estimate downward, therefore these estimates can be viewed as a lower bound.

Additional patterns in support of the theory are reported in Appendix D. First, Appendix Table D1 isolates the contribution of supply factors by showing that the strong correlation between shares of spending on new products and mean consumer income across product modules is unaffected by the inclusion of household fixed effects. This result rejects the hypothesis that the share of spending on new products is higher in product modules catering to higher-income households simply because new products diffuse faster due to a composition effect in demand.²⁹

Second, Appendix Table D2 conducts a decomposition establishing that the inflation difference between high- and low-income households is driven by differences across manufacturers, rather than across retailers or stores. Indeed, most of the inflation difference occurs because high- and low-income households purchase different barcodes within the same store, rather than because they purchase from different stores or retailers in different areas.

Furthermore, Appendix Figure D2 indicates that over time competition, as measured by Herfindahl indices, increases in higher-quality tiers segments of the market, relative to lower-quality tiers. This evidence supports the prediction that increases in market size in higher-quality tiers spur entry and increasing competition.

Finally, Appendix Figure D3 investigates inflation patterns across states. In all states, inflation was lower for households earning above \$100,000 a year, relative to low-income households making below \$30,000 a year. The inflation difference between high- and low-income households was larger in states with a faster increase in inequality, which is consistent with the notion of an endogenous supply response to changes in relative market size.³⁰

In sum, the descriptive evidence provides strong support for the idea that increasing market size, high levels of product introductions and low levels of inflation are closely connected. The relationships described so far are only correlations and should not be interpreted as causal. But they provide transparent evidence on the pervasive nature of the relationship between product innovations and inflation, on its link to changes in market size, and on its relevance for understanding purchasing-power inequality.

4.2 Causal Evidence

The equilibrium relationships between product innovations, price changes and quantities across product modules documented in Section 4.1 do not identify the causal effect of demand, because of reverse causality (better products will have larger markets: causality might run from supply to demand) and omitted variable bias (there might be unobserved heterogeneity in the difficulty of innovating across modules, which could happen to coincide with spending patterns from nonhomothetic preferences). To address this issue, I consider

²⁹Higher-income consumers might have a higher taste for novelty and purchase new products wherever they are introduced in the product space, while the rate of product introduction may be similar across product modules. The inclusion of household fixed effects directly addresses such composition effects.

³⁰Changes in inequality at the state level induce changes in relative market size only for retailers with strong local brand capital (Bronnenberg et al., 2012). See Appendix D for a complete discussion.

changes in market size at the national level driven by changes in the US age and income distributions.³¹

4.2.1 Research Design

This section describes the research design, possible threats to identification and how to address them.

Consider two periods, N socio-demographic household groups denoted by n and L product categories denoted by l . Between the two periods, each household group is characterized by the growth rate of the number of households in this group, denoted g_n , and by growth rates in per-capita spending across the product space, denoted \widetilde{g}_{nl} . The share of total spending in product category l accounted for by household group n in the initial period is denoted s_{nl} .

For an outcome variable Y_l (e.g. spending on new products and inflation for continued products), the relationship of interest is:

$$Y_l = \alpha + \beta X_l + \eta_l$$

where $X_l = \sum_n s_{nl}(\widetilde{g}_{nl} + g_n)$ is the overall growth of demand (total spending) in product category l .

The first empirical challenge is reverse causality. Changes in supply that improve products and lower prices in a given part of the product space will endogenously lead to an increase in spending per capita \widetilde{g}_{nl} in that part of the product space. To address this concern, I use the component of growth of total spending resulting only from changes in the number of households g_n :

$$Y_l = \alpha + \beta Z_l + \epsilon_l \tag{6}$$

where $Z_l = \sum_n s_{nl}g_n$ is the growth of demand in product category l implied by changes in the number of households in each group.

Thus, spending profiles across the product space are kept constant and the variation in predicted demand comes entirely from changes in age-income group size over time. Variations in market size driven by US socio-demographic changes are likely to be exogenous to other, for example scientific, determinants of product innovations. The strategy of using time-invariant spending profiles (s_{nl}) and changes in the size of household groups (g_n) to address reverse causality follows Acemoglu and Linn (2004).³² It can also be viewed as a Bartik research design (Bartik, 1991, Blanchard and Katz, 1992, Goldsmith-Pinkham et al., 2016).

Specification (6) is implemented by building 108 age-income groups and segmenting the product space by product modules by price deciles. Appendix D.2.1 describes how the size of each of these groups varies over the course of the sample, using data from the Annual Social and Economic Supplement of the Current Population Survey.³³ In general, older groups grow faster, as the baby boomers enter retirement, and more

³¹In robustness checks, I use variation in market size both over time and across local markets within the US.

³²Using age spending profiles, Acemoglu and Linn (2004) showed that large R&D efforts in the pharmaceutical industry endogenously respond to market size. By focusing on product innovations in retail, I study innovation dynamics of a different nature — increasing product variety. More importantly, this paper is the first to examine the causal effect of changes in market size on the price of continued products, as well as on the aggregate price index taking into account the welfare gains from increasing product variety.

³³Nine household income groups (<10k, 10k-20k, 20k-30k, 30k-40k, 40k-50k, 50k-60k, 60k-70k, 70k-100k, >100k, in 2004

affluent groups grow faster, in the context of growth and increasing inequality. Moreover, there is substantial variation in growth rates within age groups and within income groups.³⁴ Appendix Figures D5 and D6 and Appendix Table D4 document these trends in detail.

The product space is segmented into 10,750 product-module-by-price-decile cells, because new products and inflation dynamics vary widely even within product modules, as documented in Section 3 and Appendix C. Panel A of Figure 4 shows that income groups are segmented across the quality ladder within product modules, while Panel B shows that some products have strong age profiles, like baby diapers. The age-income spending profiles for each product module by price decile are built using data at the beginning of the sample, from 2004 to 2006.

Given these time-invariant age-income spending profiles, I compute changes in market size implied by shifts in the age and income distributions over time. The resulting patterns of predicted spending growth across the product space are shown on Panel C of Figure 4. The implied changes in market size are smoothly distributed around an annualized growth rate of around 1%, with substantial variation in growth rates. The standard deviation of annualized growth rates is 0.5%, as indicated in Appendix Table D5.

In sum, changes in the US age and income distributions over the course of the sample generate substantial variation in demand across the product space. Intuitively, some products have strong age profiles (e.g. diapers), some have strong income profiles (e.g. organic products) and some have distinct age-by-income profiles (e.g. craft beer and high-end wine). Identification requires that changes in the age and income distributions should have a direct effect on the equilibrium through demand only, and should not directly affect supply. Studying heterogeneity in the effect by age provides a test of this assumption, given that older household groups are likely to be marginally attached to the labor force and innovation activities.³⁵

The second empirical challenge is statistical power. The identification strategy described above requires using spending profiles estimated in the base period. If the spending patterns of the various age-income groups are not stable across the product space, then the research design will have no power. This can be checked directly by testing whether per capita spending at the beginning of the sample, from 2004 to 2006, is a good predictor of per capita spending at the end of the sample, from 2011 to 2013.

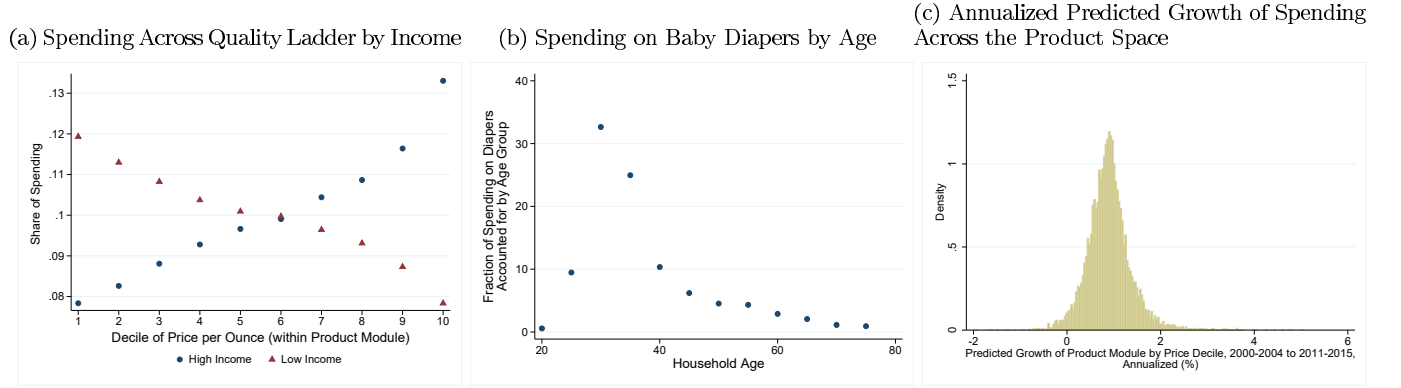
This regression is implemented using the spending patterns across 10,750 product modules by price deciles for the 108 age-income. The results are reported in Panel A of Figure 5 and in column 1 of Table 5. Initial spending patterns are strong predictors of future spending patterns: the point estimate is close to one, is precisely estimated, and the R^2 is above 0.6.

dollars) and twelve household age groups (each spanning five years: 20-25, 25-30, ... 70-75, >75) are considered. Household age is defined as the average age of the household head and their spouse. The patterns are similar when considering 30 coarser groups and when using the Census public-use micro data (IPUMS).

³⁴For instance, considering households in their thirties, there are relatively more high-income and low-income households over time, relative to a shrinking middle class. In contrast, for households in their sixties, growth rates are monotonically increasing in household income.

³⁵Moreover, I introduce another research design in Appendix D.5, based on food stamp policy changes across states, which by construction is immune to direct supply effects.

Figure 4: Changes in Market Size Implied by Changes in the Age-Income Distribution, 2004 to 2013



Notes: Panel A shows the fraction of spending of households earning above \$100,000 a year (“high income”) and below \$30,000 (“low income”) across product modules by price deciles, built as in Section 4.1. Panel B shows the fraction of total spending on diapers accounted for by households across the age distribution. Panel C reports a histogram of the demand growth predictor built in Section 4.2.1 across 10,750 product modules by price deciles. The sample extends from 2004 to 2013.

The third empirical challenge is omitted variable bias due to the endogeneity of initial expenditure shares. The OLS estimator for β in specification (6) can be written as:

$$\hat{\beta} = \beta + \frac{\frac{1}{N} \sum_n (g_n \cdot (\sum_l w_{nl} \epsilon_l))}{\frac{1}{L} \sum_l (\sum_n (s_{nl} - \bar{s}_n) g_n)^2} \quad (7)$$

where $\bar{s}_n = \frac{1}{L} \sum_l s_{nl}$ measures the importance of household group n in an average product category, and $w_{nl} = \frac{N}{L} (s_{nl} - \bar{s}_n)$ is increasing in the share of spending of n in l relative to the full sample. A formal derivation is provided in Appendix D.2.2.

Consistency requires the numerator $\frac{1}{N} \sum_n (g_n \cdot (\sum_l w_{nl} \epsilon_l))$ to go to zero.³⁶ Intuitively, this means that there should be no systematic relationship across socio-demographic groups between the growth rate (g_n) and the weighted-average of the error term ϵ_l across the product space ($\sum_l w_{nl} \epsilon_l$) characterizing this group. For instance, consistency fails if groups that grow faster disproportionately source their consumption from parts of the product space where innovation is intrinsically easier.

Given the observed patterns of growth across age and income groups, without the introduction of more controls the condition for consistency shown in (7) seems unlikely to be met. For instance, higher-income groups tend to grow faster (Appendix Figure D5) and they source their consumption from higher-quality segments of the market (Panel A of Figure 4). It may be intrinsically easier to innovate and push prices down in high-quality segments of the market, where products are more differentiated and consumers are willing to pay for novelty. Similarly, older household groups grow faster and they may spend disproportionately more on product categories where innovation is low and inflation is high because, having defined their tastes earlier in life, these households are less likely to adopt new products.

To address these possible sources of omitted variable bias, I introduce age group fixed effects and income group fixed effects, so that the coefficient β is identified from residual variation of household group growth

³⁶In addition, the denominator should not go to zero. This technical condition is satisfied as long as there is sufficient concentration of spending across household groups in a typical product category, which holds in the data. See Appendix D.2.2 for a complete discussion.

across the joint age-by-income distribution. The specification is:

$$Y_l = \alpha + \beta Z_l + \sum_a \lambda_a s_{al} + \sum_i \lambda_i s_{il} + \tilde{\epsilon}_l \quad (8)$$

where s_{al} (s_{il}) is the share of total spending in l accounted for by households in age group a (in income group i) in the initial period. The OLS estimator for β in specification (8) becomes:

$$\hat{\beta} = \beta + \frac{\frac{1}{N} \sum_n (\tilde{g}_n \cdot (\sum_l w_{nl} \tilde{\epsilon}_l))}{\frac{1}{L} \sum_l (\sum_n (s_{nl} - \bar{s}_n) \tilde{g}_n)^2} \quad (9)$$

where \tilde{g}_n is the growth rate of the number of households in group n after residualizing growth rates by age fixed effects and income fixed effects (Appendix D.2.2 provides the proof). In words, the research design identifies demand shocks from changes in the joint age-by-income distribution, rather than by exploiting broader changes affecting the entire age and income distributions (because the latter are more likely to be correlated with determinants of supply). Much of the overall variation in growth rates across groups occurs within age and income groups, as shown in Appendix Figure D6, which ensures that statistical power is retained even after the inclusion of controls.³⁷

A simple example illustrates the nature of the variation used in the research design: baby diapers. Consistent with the “polarization” of the US labor market and the shrinking US middle class (Autor et al., 2006), over the course of the sample the numbers of 30-year-old households in both high- and low-income brackets grew faster than the number of 30-year-old middle-income households. The various panels of Appendix Figure D5 show that 30-year-old low-income households grew faster than both the average 30-year-old household and the average low-income household. Consequently, specification (8) attributes a positive demand shock to product categories primarily consumed by 30-year-old low-income households, such as low-quality baby diapers (conditional on age and income controls). Likewise, a positive demand shock is given to product categories mostly consumed by 30-year-old high-income households, such as high-quality baby diapers, while a negative demand shock is attributed to product categories primarily consumed by 30-year-old middle-income households, such as mid-range baby diapers. This identifying variation for demand shocks within the baby diaper category appears unlikely to be correlated with determinants of supply: supply shocks would need to vary non-monotonically along the quality ladder for baby diapers.

Finally, to address the concern that some omitted variable bias may survive even conditional on age and income controls, I implement a falsification test using the lagged growth rates of the various age-income groups, computed between 1988 and 1999. Appendix Figure D7 shows that lagged and contemporaneous growth rates are very weakly correlated. Lagged growth rates can be used to implement a falsification test as follows: if the demand predictor in specification (8) spuriously captures properties of expenditure shares, using lagged growth rates should not affect the point estimates much; but if the effect comes from changes in market size, specifications using lagged growth rates should lose significance.

³⁷Residualizing by age group fixed effects and income group fixed effects, the standard deviation of annualized growth rates across household groups becomes 1.25%, which is 56% of the raw standard deviation.

4.2.2 Results

The strength of the relationship between market size growth, product innovations and inflation are best seen graphically using binned scatter plots. Panels B and C of Figure 5 show that the predicted increase in market size (based on the changes in the age and income distributions) is positively correlated with the introduction of new products and negatively correlated with inflation. The relationships are strong and lend clear support to the hypothesis that supply endogenously responds to changes in market size.

Table 5 shows that the relationships between predicted market size growth, product innovations and inflation are large and significant at the 1% level, clustering standard errors by product module. The interpretation of the magnitudes is as follows: a one percentage point increase in the annualized growth of predicted demand causes a 2.73 percentage point increase in the annual share of spending on new products and a 0.43 percentage point decline in the annual inflation rate on goods that are available across years.³⁸

Panels B and C of Table 5 document additional results. Panel B runs a falsification test using lagged growth rates of age-income groups. There is no significant relationship between the placebo predicted change in market size and entry of new products or inflation. Appendix Figure D7 reports the corresponding binned scatter plots. Panel C documents heterogeneity in the treatment effect.³⁹ There is no significant heterogeneity in the effect of changes in demand on supply depending on the age of consumers: the results are not driven by young consumers, for whom direct supply effects are more likely to exist.⁴⁰ This finding addresses the concern that the relationship between predicted demand, innovation and inflation could be spuriously driven by a differential increase in supply across the product space. The second row of Panel C documents that there is no heterogeneity in the effect of demand across the income profile of product categories. This result justifies using the average treatment effect for all income groups in the prediction exercise carried out in Section 4.3.

A variety of robustness checks are reported in Appendix D.2.3. Panel A of Appendix Table D6 shows that increasing demand leads to more exit, but that the entry effects are stronger, such that on net product variety increases when demand increases. This panel also confirms that the relationship between predicted and actual growth of total spending is significant at the 5% level.⁴¹ Panel B shows that the effect of demand on supply does not vary much with the degree of competition across the product space, measured with the Herfindahl index. Panel C shows that the results are similar without spending weights.

³⁸The magnitude of these effects is consistent with the competition channel: Appendix E.3.8 calibrates the strength of these effects using the model derived in Section 5 and shows that the various point estimates are in line with the notion that increasing demand leads to an increase in product variety and lower prices on continued products through the response of markups.

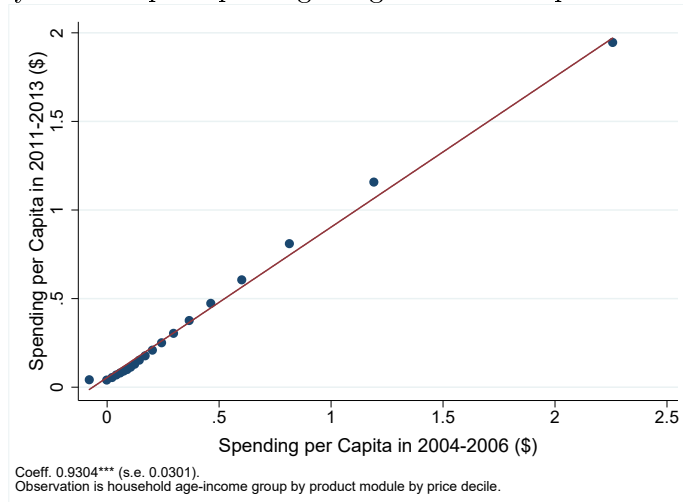
³⁹The interacted regressors of interest are (spending-weighted) average consumer age and average consumer income across product modules by price deciles indexed by l . Average consumer age is defined as $A_l \equiv \sum_n s_{nl} A_n$ and average consumer income as $I_l \equiv \sum_n s_{nl} I_n$, where A_n is the age of household group n and I_n their income. The heterogeneity specifications are similar to equation (8), but including an interaction between the predictor of demand and the interacted regressor of interest. The interacted regressors are standardized by their standard deviation and are demeaned.

⁴⁰Likewise, the point estimates are similar when repeating the analysis in a subsample of product categories whose average consumer age is above 55.


⁴¹The point estimate is close to 1, i.e. the predictor is unbiased. Unbiased prediction wasn't necessarily expected, because the measure of actual total spending growth takes into account both price and quantity effects, while the predicted increase in spending is based on the assumption that spending per capita is fixed.

Figure 5: Causal Effects of Changes in Market Size

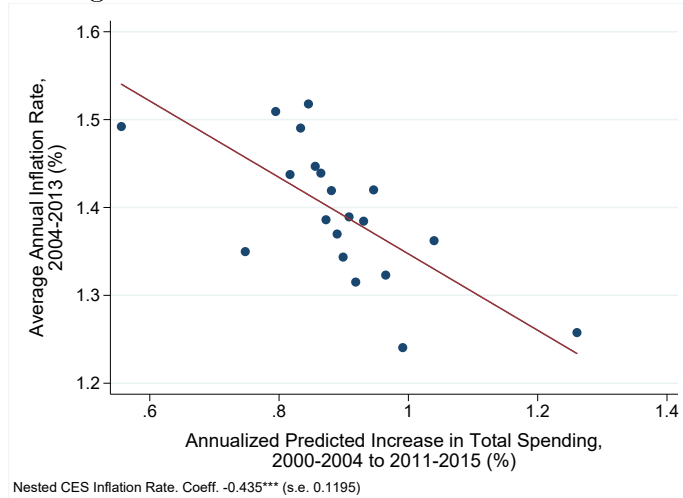
Panel A: Stability of Per-Capita Spending of Age-Income Groups across the Product Space



Panel B: Increasing Market Size Leads to Product Entry



Panel C: Increasing Market Size Leads to Lower Inflation for Continued Products



Notes: Panel A reports the relationship between per capita spending across product modules by price deciles for the 108 age-income groups described in Section 4.2.1. Panel B and C report OLS best-fit lines and binned scatter plots across 10,750 product modules by price deciles. Each dot represents 5% of the data. Panel B shows the relationship between the demand growth predictor built in Section 4.2.1 and the share of spending on new products, while Panel C reports the relationship between the demand growth predictor and the nested CES inflation rate for continued products. Regression results are reported in Panel A of Table 5. The sample extends from 2004 to 2013.

Table 5: Causal Effects of Changes in Market Size

Panel A: Main Results

	Per-Capita Spending in 2011-2013 (\$)	Share of Spending on New Products (pp)	Continued Products Inflation Rate (pp)
Per-Capita Spending in 2004-2006 (\$)	0.93044*** (0.03012)		
Predicted Increase in Spending, Annualized (%)		2.7358*** (0.4887)	-0.4349*** (0.1195)
Age and Income Controls	Yes	Yes	Yes
Product Module Fixed Effects	No	Yes	Yes
R^2	0.61	0.54	0.52
Number of Observations	831,179	10,750	10,750
Number of Clusters	1,075	1,075	1,075

Panel B: Falsification Test Using Lagged Predictor of Spending

	Share of Spending on New Products (pp)	Continued Products Inflation Rate (pp)
Lagged Predictor of Increase in Spending, Annualized (%)	0.8254 (0.9949)	-0.1802 (0.2284)
Age and Income Controls	Yes	Yes
Product Module Fixed Effects	Yes	Yes
R^2	0.55	0.52
Number of Observations	10,750	10,750
Number of Clusters	1,075	1,075

Panel C: Heterogeneity in Effect of Increasing Demand

	Share of Spending on New Products (pp)	Continued Products Inflation Rate (pp)
Annualized Predicted Increase in Spending × Average Consumer Age	-0.06834 (0.19047)	-0.004813 (0.03571)
Annualized Predicted Increase in Spending × Average Consumer Income	0.02806 (0.2043)	-0.002725 (0.02794)
Age and Income Controls	Yes	Yes
Product Module Fixed Effects	Yes	Yes
R^2	0.54	0.52
Number of Observations	10,750	10,750
Number of Clusters	1,075	1,075

Notes: Panels A to C report the results of regressions at the level of product modules by price deciles. The cells and the independent variable are built as described in Section 4.2.1. The inflation rate for continued products is the nested CES price index, computed based on barcodes that are available across consecutive years. The specification in the first column of Panel A is described in Section 4.2.1. The specifications in columns 2 and 3 of Panel A and in Panel B are given by equation (8), with product module fixed effects. In Panel C, the interacted regressors are standardized by their standard deviation and are demeaned, so that the regression coefficient on the treatment effect is similar to Panel A. The specifications for Panel C are described in Section 4.2.2. The sample extends from 2004 to 2013. In all panels, standard errors are clustered by product modules and regressions are weighted by log spending. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

4.3 Implications of Shifts in the Income Distribution for Inflation Inequality

Do historical shifts in the income distribution imply substantial inflation inequality through the equilibrium response of supply? Over the course of the sample, and more broadly over recent decades, demand from high-income consumers has been increasing faster than demand from low-income consumers for two reasons. First, because of economic growth, more and more consumers have become high-income earners over time. Second, because of rising income inequality, the purchasing power of consumers at the top of the income distribution has been increasing faster than that of consumers at the bottom. These trends in the US income distribution have been widely documented in the macro and labor literatures (e.g. Song et al., 2016). Furthermore, Table 5 indicates that the long-term supply curve is downward-sloping. Using this estimated slope, I compute the implied effect on inflation of historical shifts in the income distribution, measured in the Annual Social and Economic Supplement of the Current Population Survey. I find that changes in the income distribution explain close to all of the observed inflation inequality.

To conduct this analysis, it is important to non-parametrically allow for a continuum of tastes across income groups, so that local density changes at a certain point of the income distribution may imply market size effects affecting consumers located in other parts of the income distribution (through common tastes). I do so in three steps. First, I use historical changes in the US income distribution to get changes in demand across detailed cells of the product space. I consider 18 household income groups, denoted by i , which are available in the Nielsen data from 2006 to 2009.⁴² Considering product-module-by-price-decile cells indexed by l , I build changes in demand induced by changes in the income distribution as follows:

$$d_l = \sum_n s_{nl} \cdot (1 + g_n)$$

where s_{nl} is the share of total spending in l accounted for by income group n in 2006-2009 and g_n is the growth rate of the number of households in group n between 1986 and 2006.⁴³ d_l capture changes in demand induced by historical changes in the income distribution.

Second, I use the relevant point estimates from Table 5 to infer the patterns of product innovations and inflation induced by changes in the income distribution through market size effects. The predicted values are computed as follows:

$$\text{Predicted Share of Spending on New Products}_l = 2.73 \cdot d_l \tag{10}$$

$$\text{Predicted Continued Product Inflation Rate}_l = -0.43 \cdot d_l \tag{11}$$

Finally, I compare the actual and predicted relationships between mean consumer income, spending on


⁴²This period is the only period in which Nielsen provides distinct bins of the income distribution above an income of \$100,000. Detailed income bins below \$100,000 are available throughout the sample, from 2004 to 2013. The results are similar for these income bins from 2004 to 2013 — these results are available from the author upon request.

⁴³I report results using secular changes in the income distribution, assuming that suppliers respond to long-term trends. Panels A and B of Appendix Figure D8 document the patterns of growth across income groups: the growth of the number of households in a group is monotonically increasing in the income of that group. The results are similar when using changes in the size of the various income groups over the course of the Nielsen sample (2004-2013) instead of the long-term trend, as shown in Panel C of Appendix Figure D8.

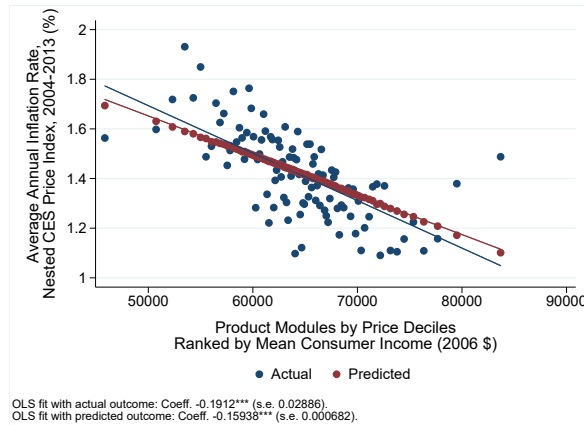
new products and inflation for continued products across the product space.⁴⁴ Mean consumer income in product module by price decile l is defined as $I_l \equiv \sum_n s_{nl} I_n$, with I_n the income level of household group n . I regress the actual and predicted values for spending on new products and inflation for continued products on I_l and I compare the OLS best-fit lines to gauge whether a substantial fraction of actual inflation inequality is induced by changes in the income distribution through market size effects and the endogenous supply response.⁴⁵

Figure 6: Inflation Inequality Implied by Changes in the Income Distribution

Panel A: Consumer Income and Spending on New Products, Actual vs. Predicted



Panel B: Consumer Income and Inflation for Continued Products, Actual vs. Predicted



Notes: Panel A shows the relationship between actual and predicted spending on new products and mean consumer income across product modules by price deciles. Panel B shows the relationship between actual and predicted nested CES inflation rates for continued products and mean consumer income across product modules by price deciles. The predictors are built using equations (10) and (11). The actual data is shown in blue, with the OLS best-fit line and one hundred data points each capturing 1% of the data. All specifications include log spending weights and product module fixed effects. Regression results and robustness checks are reported in Appendix Table D7. The sample extends from 2004 to 2013.

Figure 6 shows that the actual and predicted relationships between mean consumer income, spending

⁴⁴This methodology relies on two assumptions. First, income-group-specific expenditure shares across the product space should be stable over time, which Panel A of Figure 5 shows. Second, the response of supply to changes in demand should be similar regardless of the income profile of the product category, which Panel C of Table 5 lends support to.

⁴⁵The specifications are of the form $Y_l = \beta I_l + \lambda_m + \epsilon_l$, where Y_l is in turn the actual share of spending on new products, the predicted share of spending on new product from equation (10), the actual nested CES inflation rate for continued products, and the predicted nested CES inflation rate for continued products from equation (11). λ_m denotes product module fixed effects.

on new products and inflation for continued products are closely aligned.⁴⁶ The predicted slopes are about 83% of the actual slopes, as reported in Appendix Table D7. This table shows that the results are similar regardless of the inclusion of product module fixed effects in the regression, indicating that the predictor is potent both across and within product modules. Overall, the results show that changes in the household income distribution explain a large fraction of inflation inequality across household income groups, through the endogenous response of supply.⁴⁷

4.4 Additional Evidence

To discipline the model developed in the final part of the paper, I present a series of additional results. First, I confirm that the supply response is induced by *changes* in market size, rather than by the *level* of market size. Second, I present evidence of a disproportionate fall in markups in parts of the product space catering to higher-income households. Third, I show that a similar supply response exists for changes in per capita spending at the bottom of the income distribution. Fourth, I find that other channels that may affect inflation inequality, such as international trade and the product cycle, are quantitatively less important than the supply dynamics previously documented.

Changes in market size vs. level of market size. To test whether the supply response is driven by changes in market size, rather than the level of market size, I use a research design similar to the national age-income group research design, but exploiting cross-state variation. Specifically, I predict the level of spending in a state based on the initial age and income distribution in that state and the age-income spending per capita profiles estimated using data in *other* states (thus addressing the identification concern that cheaper products typically attract more spending). I then predict the change in spending using the observed change in the size of the various age-income groups in each state.

The results, reported in Panel A Table 6, show that the fall in inflation is entirely predicted by the *increase* in spending, rather than by the initial level of spending. Appendix Table D8 presents similar results with OLS regressions at the national level. These results are consistent with models in which an increase in market size only has a *temporary* effect on the level of innovation. In other words, *changes* in market size are the relevant predictors of innovation, not the level of market size.

The role of markups. As mentioned in Section 2, I observe retailer price p_{it} and wholesale cost c_{it} from 2004 to 2007 for a subset of the products. The retailer’s gross margin m_{it} is defined by: $p_{it} = m_{it} + c_{it}$. Do prices rise more slowly for high-income consumers because retailer margins decline more quickly or because wholesale costs rise more slowly? To answer this question, first note that a first-order Taylor expansion

⁴⁶The predicted patterns for new products and inflation for continued products are almost perfectly linear. This is not mechanical but results from the pattern of growth of the number of households across the income distribution, which happens to be linearly increasing in the relevant range, as documented on Panel A of Appendix Table D8.

⁴⁷Figure 1 depicts a convex pattern of inflation across income groups, which contrasts with the linear relationship between mean consumer income and inflation across the product shown in in Figure 6. The contrast is due to the nature of non-homotheticities in the retail sector: above some income level, households tend to converge to similar consumption baskets because they all purchase high-quality items, and therefore have the same inflation rate. Appendix Figure D9 confirms this intuition by plotting a dissimilarity index characterizing consumption patterns across the income distribution.

yields a convenient additive expression for the log price change: $\Delta^t \log(p_{it}) \approx \Delta^t \log(c_{it}) + \Delta^t \frac{m_{it}}{c_{it}}$. Next, with I_i denoting mean consumer income in the module of product i and with λ_{st} denoting module-store-year fixed effects, I run the following regressions across product modules: $\Delta^t \log(p_{it}) = \beta I_i + \lambda_{st} + \epsilon_{it}$; $\Delta^t \log(c_{it}) = \tilde{\beta} I_i + \tilde{\lambda}_{st} + \tilde{\epsilon}_{it}$; $\Delta^t \frac{m_{it}}{c_{it}} = \bar{\beta} I_i + \bar{\lambda}_{st} + \bar{\epsilon}_{it}$. Note that $\beta \approx \tilde{\beta} + \bar{\beta}$, which provides a convenient decomposition of inflation inequality (β) into wholesale cost effects ($\bar{\beta}$) and retailer margin effects ($\tilde{\beta}$).⁴⁸ The fixed effects absorb rent and labor costs, addressing the concern that the retailer gross margin m_{it} may fall due to changes in rent and labor cost. Empirically, inclusion of these fixed effects does not affect the results, which suggests that the relationship is driven by changes in markups over wholesale cost rather than differential changes in rents and labor costs across the product space.

Panel B of Table 6 shows that differential changes in retailer margins account for about half of the differential inflation between high- and low-income households, while differential changes in wholesale costs account for the other half. The result that about half of inflation inequality is explained by changes in retailer margins can be thought of as a lower bound on the total share of changes in markups in inflation inequality, because wholesalers, and in turn manufacturers, themselves have a markup.⁴⁹ Appendix Figure D10 depicts changes in retailer margins and wholesale cost graphically. These relationships are robust across years and when using other specifications, changing the sets of fixed effect and weights. These results provide motivation for a model featuring variable markups.

Food stamp research design. Appendix Section D.5 introduces a research design based on food-stamp policy changes across US states. Between 2001 and 2007, the take-up rate for food stamps substantially increased in some states due to a series of policy changes that made it easier for eligible individuals to enroll in the program. This policy variation generates variation in purchasing power for food products at the bottom of the income distribution.

This identification strategy is a useful complement to the previous analysis based on changes in the number of consumers across the product space at the national level over time. First, it is useful to examine whether variation in demand coming from changes in per capita spending generates similar effects to variation in demand coming from changes in the number of consumers. Second, the SNAP-based research design has a number of advantages from the point of view of identification: there is clearly no direct supply effect, the market size change occurs at the bottom of the distribution (thus breaking the usual collinearity between level of income and rate of growth in income), and the time frame and the location of the market size change are known very precisely. Third, these findings are of direct policy relevance. I find a large effect, which can be summarized as follows: a 1 percentage point increase in spending per capita lowers the inflation rate by about 10 basis points.

Alternative mechanisms. In Appendix Section D.6, I study in depth two mechanisms that may

⁴⁸As can be checked from the regression table, the margins are sufficiently small for the Taylor expansion to be almost exact, which in turn implies that the relationship between the regression coefficients is almost exact.

⁴⁹See Appendix E.4 for a formal double marginalization model making this point.

disproportionately benefit the poor: the product cycle and international trade. I show that although these mechanisms appear to indeed play a role and benefit the poor relatively more, they are quantitatively less important than the endogenous supply dynamics documented previously. I then investigate a series of other possible mechanisms — aggregate shocks, online retail, innovation dynamics independent of changes in market size, and household search behavior — and find that they can’t be the primary drivers of the patterns found in the data.

Table 6: Additional Evidence on Supply Response to Changes in Demand

Panel A: Lower Inflation is Caused by Increases in Market Size

	Inflation Rate for Continued Products (pp)
Predicted Increase in Spending, 2004-2006 to 2011-2013 (%)	-0.04796*** (0.0111)
Predicted Level of Spending in 2004-2006 (Log)	-0.00437 (0.0071)
Department Fixed Effects	Yes
Spending Weights	Yes

Panel B: Changes in Wholesale Costs vs. Changes in Retailer Margins

	Log Price Change	Log Wholesale Cost Change	Retailer Margin Change (pp)
Mean Income of Consumer in Product Module (\$10,000)	-0.777*** (0.188)	-0.341*** (0.103)	-0.448*** (0.212)
Store-Year Fixed Effects	Yes	Yes	Yes
Spending Weights	Yes	Yes	Yes
Number of Observations	6,002,235	6,002,235	6,002,235
Number of Clusters	628	628	628

Notes: Panel A reports the result of regressions at the level of product modules by US states, following the methodology described in Section 4.4, in the full sample extending from 2004 to 2013. Panel B reports the results of barcode-level regressions, using information on wholesale cost and retail margin variables available in a subsample of the data, from 2004 to 2007, described in Appendix Section A.1.2. For both panel, the sample is restricted to continued products, defined as barcodes that are available across consecutive years. In both panels, standard errors are clustered by product modules. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

5 Model and Implications

In this section, I develop a model motivated by the empirical results from the previous sections. On the demand side, the model flexibly accounts for non-homotheticities and generates tractable exact price indices across the income distribution. On the supply side, it features a downward-sloping long-term supply curve through both the endogenous introduction of new products and endogenous markups. Using CEX and CPI data, I find support for the key prediction of the model that lower-inflation for higher-income households is a long-term trend in retail and beyond. I then use the model for welfare calculations and find that the overall welfare effect of inflation inequality in the retail sector across income groups is large. Taking into account inflation inequality in retail, the rate of increase in purchasing-power inequality between the bottom and top

income quintiles is 25% faster than when only considering changes in nominal income. I conclude this section by discussing the implications of the findings for policy and for our understanding of innovation dynamics.

5.1 Model

I develop a model explaining the patterns of changes in product variety and inflation across the income distribution through the endogenous supply response to changes in demand implied by shifts in the income distribution. I describe the setting, solve for the equilibrium and the comparative statics, and relate the results to the empirical evidence and other models in the literature.⁵⁰

Non-homothetic CES aggregator across sectors. Consider an arbitrary number of household types differing in their productivity level l^i , $i \in I$. The numbers of households of each type in each period t is denoted L_t^i . These agents consume and produce in K different sectors of the economy. The number of products available in each sector is endogenous and denoted N_k . In period t , households maximize aggregate consumption C_i^t using a non-homothetic CES aggregator of sectoral consumptions $\{C_{ik}^t\}_{k=1}^K$:

$$C_i^t = \left(\sum_{k=1}^K \tilde{\Omega}_{ik}^{\frac{1}{\sigma_i}} (C_{ik}^t)^{\frac{\sigma_i-1}{\sigma_i}} \right)^{\frac{\sigma_i}{\sigma_i-1}}$$

where the income-group-specific elasticity of substitution σ_i and income-group-specific sectoral weights $\tilde{\Omega}_{ik}$ capture non-homotheticities. Appendix E.3 provides a micro-foundation for the $\tilde{\Omega}_{ik}$ weights in terms of total consumption C_i^t .⁵¹

Translog preferences within sectors. Each sectoral consumption k is itself a consumption aggregator. Within a sector k , all households have the same translog preferences. Let \tilde{N}_k be the total number of products (or, alternatively, firms)⁵² conceivably available in sector k and treat this number as fixed. Dropping the k and t subscripts for convenience and denoting by p_n the price of product n , within each sector the translog expenditure function is defined as:⁵³

$$\ln(E) = \ln(U) + \alpha_0 + \sum_{n=1}^{\tilde{N}} \alpha_n \ln(p_n) + \frac{1}{2} \sum_{n=1}^{\tilde{N}} \sum_{m=1}^{\tilde{N}} \gamma_{nm} \ln(p_n) \ln(p_m) \quad \text{with } \gamma_{nm} = \gamma_{mn} \quad \forall m, n$$

Labor supply and budget constraint. Labor is supplied inelastically. Households of type i are endowed with l^i effective units of labor. The wage for one effective unit of labor is the numeraire. Each

⁵⁰A number of simplifying assumptions are made for tractability but are all relaxed in Appendix E.3: 1. firms in each sector are homogeneous, i.e. have the same marginal and fixed costs of production; 2. agents maximize current-period consumption (hand-to-mouth); 3. consumer preferences are non-homothetic *across* sectors (i.e. different agents place different weights on the various sectors, depending on their income levels) but are homothetic *within* sectors (i.e. at the lowest level of aggregation, all agents have the same spending patterns); 4. all agents enter the production function in a similar way across sectors, which implies that there is no feedback effect of shifting demand on wages across agent types; 5. all agents pay the same price for each barcode (within sectors).

⁵¹Intuitively, each sector k is characterized by an income elasticity parameter $\tilde{\epsilon}_k$. As aggregate consumption C_i increases, the weight given to the consumption of good k varies at a rate controlled by the parameter $\tilde{\epsilon}_k$. This generalization of the standard (homothetic) CES aggregator follows Comin, Lashkari and Mestieri (2016).

⁵²In my model presentation below and in the discussion of results, there is a one-to-one identification between a producer, a product, and a firm. The results are robust to the introduction of multi-product firms, as shown in Appendix E.3.9.

⁵³See Diewert (1974) and Feenstra (2003). The restrictions $\sum_{n=1}^{\tilde{N}} \alpha_n = 1$ and $\sum_{n=1}^{\tilde{N}} \gamma_{nm} = 0$ ensure that the expenditure function is homogeneous of degree one. Following the literature, I impose that all goods enter the expenditure function symmetrically: $\alpha_n = \frac{1}{\tilde{N}}$, $\gamma_{nn} = \frac{-\gamma(\tilde{N}-1)}{\tilde{N}}$ and $\gamma_{nm} = \frac{\gamma}{\tilde{N}}$ for $m \neq n$, with $n, m = 1, \dots, \tilde{N}$.

household type is subject to the period budget constraint, with c_{inkt} denoting consumption of variety n in sector k by household i at time t :

$$\sum_{k=1}^K \sum_{n=1}^{N_k} c_{inkt} \cdot p_{nkt} = l^i$$

Monopolistic competition between homogeneous firms within sectors. Within each sector k , symmetric firms compete monopolistically. Given symmetry, firm subscripts n are dropped in what follows. The quantity produced by a single firm is given by $q_k = Z_k l_k$, where l_k is labor demand for production and Z_k is a productivity factor specific to sector k . Moreover, all firms pay a sunk entry cost of f_k effective units of labor, i.e. the required amount of labor for entry per firm is $\frac{f_k}{Z_k}$. Given this cost structure (see Appendix E.2.2 for a flexible cost structure) and the residual demand curve $p_k(q_k)$ implied by households' preferences, firms maximize profits:

$$\max_{q_k} \pi_k(q_k) = p_k(q_k) \cdot q_k - \frac{q_k}{Z_k} - \frac{f_k}{Z_k}$$

Labor demand. The total amount of labor required by all firms in sector k is $L_k = N_k \cdot \left(\frac{q_k + f_k}{Z_k}\right)$.

Free entry. Firms enter sector k until profits are brought to zero: $\pi_k = 0$.

Proposition 1. *In period t , across sectors k and income groups i , the equilibrium is characterized by:*

$$\text{Number of varieties } \forall k : N_{kt}^* = \frac{(\gamma_k - 1) + \sqrt{(1 - \gamma_k)^2 + 4\gamma_k \frac{(\sum_i L_t^i \cdot l^i \cdot s_{ikt}) \cdot Z_k}{f_k}}}{2\gamma_k} \quad (12)$$

$$\text{Demand elasticity } \forall k : \epsilon_{kt}^* = 1 + (N_{kt}^* - 1) \cdot \gamma_k \quad (13)$$

$$\text{Markup } \forall k : M_{kt}^* = \frac{1}{(N_{kt}^* - 1) \cdot \gamma_k} \quad (14)$$

$$\text{Price of varieties } \forall k : p_{kt}^* = (1 + M_{kt}^*) \cdot \frac{1}{Z_k} \quad (15)$$

$$\text{Sectoral price index } \forall k : \ln(P_{kt}^*) = \alpha_{0k} + \frac{1}{2} \frac{1 - \frac{N_{kt}^*}{N_k}}{\gamma_k N_{kt}^*} + \frac{1 + (N_{kt}^* - 1) \cdot \gamma_k}{(N_{kt}^* - 1) \cdot \gamma_k} \frac{1}{Z_k} \quad (16)$$

$$\text{Sectoral spending share } \forall (i, k) : s_{ikt}^* = \frac{\tilde{\Omega}_{ik} (P_{kt}^*)^{1-\sigma_i}}{\sum_{k=1}^K \tilde{\Omega}_{ik} (P_{kt}^*)^{1-\sigma_i}} \quad (17)$$

$$\text{Aggregate price index } \forall i : \mathbb{P}_t^i = \left[\sum_{k=1}^K \tilde{\Omega}_{ik} (P_{kt}^*)^{1-\sigma_i} \right]^{\frac{1}{1-\sigma_i}} \quad (18)$$

The full proof is given in Appendix E.3.

The expressions for the various equilibrium objects and the implied comparative statics with respect to market size in Proposition 1 are intuitive. Equation (12) shows that the equilibrium number of varieties in sector k is increasing in market size (total spending) in that sector, which is given by the term $(\sum_i L_t^i \cdot l^i \cdot s_{ikt})$. Market size is higher if there are more consumers (increased L_t^i), if their purchasing power is higher (increased l^i) and if their spending share in sector k is higher (increased s_{ikt}). The elasticity of demand in a sector is increasing in the number of varieties in this sector, as indicated by (13): as the number of varieties increases,

consumers perceive them as less and less differentiated.⁵⁴ This elasticity is the key feature of preferences which determines firms' optimal markups in equilibrium. The markup on any given variety is decreasing in the price elasticity of demand, hence in the total number of varieties available in that sector, as shown in (14). The equilibrium price of varieties is given by the optimal markup over marginal cost in (15).

The comparative statics of the number of varieties and the price of varieties with respect to changes in market size are in line with the results in Table 5. Increasing market size causes the introduction of more products and, through increased competitive pressure and decreasing markups, lower inflation.⁵⁵ Moreover, Appendix Section E.3.8 shows that the elasticity of the price of continuing products to changes in product variety implied by the model is closely aligned with the data.

The model delivers a tractable non-homothetic price index. Equation (16) shows that the sectoral price index goes down, i.e. welfare goes up, as the equilibrium number of varieties increases because of two forces. First, consumers love variety, which is captured by the term $\frac{1 - \frac{N_k^*}{N_k}}{2 \gamma_k N_k^*}$; note that this term decreases at a decreasing rate as N_k^* increases, because the product space gets filled and there are decreasing returns to increasing product variety. Second, an increasing number of varieties leads to lower markups, which is reflected by the term $\frac{1 + (N_k^* - 1) \cdot \gamma_k}{(N_k^* - 1) \cdot \gamma_k} \frac{1}{Z_k}$. Sectoral spending shares are determined by (17) and, in turn, (18) gives the price index for each income group.⁵⁶

Proposition 2. *Represent changes in the income distribution between periods $t - 1$ and t using a set $\{g_{it}\}_{i=1}^I$ of growth rates in the number of households with income (productivity) l^i , such that $L_t^i = (1 + g_{it})L_{t-1}^i$. For each income group i , with k denoting sectors, define the welfare-relevant market size effect implied by changes in the income distribution as:*

$$\widetilde{g}_{it} = \sum_{k=1}^K s_{ik(t-1)} \cdot \underbrace{\left(\sum_{j=1}^I \widetilde{s}_{kj(t-1)} \cdot g_{jt} \right)}_{\text{Spending growth for } k}$$

with $s_{ik(t-1)}$ as defined in (17) and $\widetilde{s}_{kj(t-1)}$ denoting the share of income group j in total spending in k at $t - 1$. Then, changes in prices indices $\pi_{it} \equiv \log(\mathbb{P}_t^i) - \log(\mathbb{P}_{t-1}^i)$ across the income distribution satisfy:

$$\widetilde{g}_{it} > \widetilde{g}_{mt} \iff \pi_{it} < \pi_{mt} \quad (19)$$

The full proof is given in Appendix E.3.7.

⁵⁴This property of the equilibrium results from the use of translog preferences at the lowest nest of the utility function. In contrast, with CES as the lowest nest, the elasticity and in turn the markup and the price remain constant regardless of how many varieties are available.

⁵⁵Variable markups are often studied in the macro literature in the context of short-run business cycle fluctuations. The fact that declining markups explains inflation inequality for continued products between high- and low-income households does not mean that these dynamics are bound to be short-lived. Indeed, the set of available products changes over time. Adjusted for quality, the marginal cost of the new products is lower than that of existing products, which are forced to reduce their markups. In other words, the price effects show up largely through changes in markups, but these changes reflect the productivity gains brought about by new products. See Appendix Section E.3.9 for an extension of the model with heterogeneous firms and heterogeneous markups.

⁵⁶The exact price index for the translog expenditure function on continuing products is the Tornqvist price index and is shown in Panel B of Figure 1. The estimation equations for this demand system are derived in Appendix Section B.3. The results for inflation across income groups, taking into account both existing and new varieties, are reported in Appendix Figure C3 and are similar to the results from Section 3.

Intuitively, Proposition 2 indicates that if a household’s preferences are skewed toward parts of the product space that grow faster, then their price index *falls* relative to that of other households. As the market size becomes relatively larger, the price index decreases because of both increasing product variety (equation (12)) and decreasing markups (equation (15)). Households who source their consumption from parts of the product space that grow faster, due in particular to changes in the income distribution, benefit disproportionately from this process and face a lower inflation rate, in line with the patterns displayed in Figure 6.

Three important lessons follow from Proposition 2. First, changes in the price index at a given point of the income distribution are determined by market size effects that take into account changes in the entire income distribution. For instance, to the extent that they share similar preferences, households in the middle class can benefit from increased spending from the upper middle class. Second, changes in the price index over time are determined by *changes* in market size rather than by the *level* of market size. Intuitively, when market size in a given part of the product space grows, it is profitable for firms to enter this part of the product space because the returns to paying the fixed cost of entry are higher in a larger market. However, markups endogenously decrease as more firms enter and competition intensifies, such that the process of entry eventually stops. In other words, product innovations and low inflation are features of growing markets, not of large markets, consistent with the evidence in Table 6. Finally, given that households taste diverge gradually across the income distribution (Appendix Figure D9), the market size effects implied by growth and rising inequality disproportionately benefit higher-income households. Given that the US economy is characterized by long-term trends of growth and rising top income shares (Song et al., 2016), the model predicts a long-term trend of lower inflation for higher-income households.

Robustness. Although upward-sloping supply curves are a typical feature of standard price theory, a variety of models can generate the key prediction that in general equilibrium the quality-adjusted price goes down when demand increases. There are three broad classes of such models: endogenous growth macro models with scale effects (e.g. Romer, 1990, Aghion and Howitt, 1992, and Acemoglu and Linn, 2004), trade models with free entry and endogenous markups through variable-elasticity-of-substitution preferences (e.g. Melitz, 2003, and Zhelobodko et al., 2012), and industrial organization models with free entry and endogenous markups through strategic interactions between firms (e.g. Sutton, 1991, and Berry and Reiss, 2006). The model I developed has the advantage of being particularly tractable, consistent with all features of the data presented in Sections 3 and 4 (in contrast, existing models make counterfactual predictions, discussed in Appendix E.1) and robust to several extensions presented in Appendix E.3.9 (multi-product firms, heterogeneous productivity across firms, non-homotheticities within sectors, feedback effects of shifting demand on the relative income of the various household types, endogenous savings, and additional nests in the utility function).


5.2 Long-Term Inflation Inequality across Income Groups

To test the key prediction from Proposition 2 that long-run changes in the income distribution should induce a long-term trend of lower-inflation for higher-income households, I use BLS and CEX data. I proceed in two steps. First, I collect CPI price series on 48 CEX expenditure categories going back to 1953, which cover the full basket of consumer goods and services and are described in Appendix A.1.3. These categories are matched by hand across the CPI and CEX surveys. Second, I build price indices for the consumption baskets of households in the top and bottom quintiles of the income distribution, using expenditure shares fixed at 1980-1985 levels (which are observed in the CEX data).

The data I thus obtain covers the full basket of consumption - in particular, housing, auto purchases and medical care are included. The advantage of this dataset is its broad coverage, as well as the fact that it goes much further back in time than the Nielsen data. This of course comes at a price: the data series are relatively aggregated, therefore it is more difficult to capture the segmentation of consumption across income groups, and quality adjustments are difficult to carry in many of the product categories.⁵⁷

To probe the external validity of the core findings of the paper, I use the CPI and CEX data to ask two questions. First, is inflation lower for the consumption basket of college graduates, relative to high-school dropouts, over a long horizon? Second, does the difference increase after the 1970s, in the broader context of increasing inequality? The answer to both questions is yes.

Figure 7: Full-Basket Inflation Inequality across Income Groups in the Long Run



Notes: This figure reports the relative price index of households in the bottom income quintile (low income) relative to households in the top income quintile (high income) from 1953 to 2015. The relative price index is normalized to one in 1953. Income-group-specific price indices are built using CPI and CEX data as described in Section 5.2.

Figure 7 plots the price index of households in the bottom income quintile relative to the price index of households in the top income quintile from 1953 to 2015. Relative to high-school dropouts, average annual inflation for college graduates was about 10 basis points lower during 1953-1970 and about 25 basis points

⁵⁷I do not make any further adjustment to the price series provided by BLS, which are meant to adjust for quality changes over time.

lower during 1970-2015. The magnitude of the inflation difference is lower than in the Nielsen data, but this was expected: the relatively broad level of aggregation of product categories biases the inflation difference towards 0, as Panel A of Table 2 showed. These results are not driven by any single broad product category and are robust to considering other base years for the spending shares, as well as other education and income groups (Appendix Figure F1).

Finally, in Appendix Figure F2, I present complementary evidence showing that technical change disproportionately benefited high-income households over a long time horizon. Despite aggregation bias, a clear pattern emerges: the number of granted patents and TFP growth have been substantially higher in sectors of manufacturing targeting high-income households.

5.3 Implications

Inequality. Using the exact price index from the model derived in equation 18, I compute the implication of inflation inequality in the retail sector for purchasing-power inequality between the bottom and top quintiles of the income distribution. I conduct a simple exercise assuming that households' preferences outside the retail sector are Cobb-Douglas and using expenditure shares from the CEX. The match of the Nielsen to the CEX data reported in Appendix 2 indicates that the Nielsen data captures 18% of spending for households in the bottom quintile of the income distribution, and 12% of spending for households in the top quintile. I compute the overall increase in nominal income inequality between 2004 and 2013, the sample period, from the Census public-use microdata (IPUMS). I find that nominal income increased 0.93 percentage points faster per year in the top income quintile, relative to the bottom income quintile. I then compute the impact of inflation inequality in retail on purchasing-power inequality between these two groups:

$$\begin{aligned} \Delta \text{Purchasing Power Inequality} = & \underbrace{(\Delta \log(Y^{Q1}) - \Delta \log(Y^{Q5}))}_{\text{Income: } -0.93pp} - \underbrace{(\alpha^{Q1} \Delta \log(\mathbb{P}^{Q1}) - \alpha^{Q5} \Delta \log(\mathbb{P}^{Q5}))}_{\text{Retail Inflation: } 0.22pp} \\ & - \underbrace{\left((1 - \alpha^{Q1}) \Delta \log(\tilde{\mathbb{P}}^{Q1}) - (1 - \alpha^{Q5}) \Delta \log(\tilde{\mathbb{P}}^{Q5}) \right)}_{\text{Inflation Beyond Retail}} \end{aligned}$$

where $\tilde{\mathbb{P}}^Q$ denotes the price index outside retail, Y^Q income and α^Q the retail expenditure share for income quintile Q . From 2004 to 2013, inflation inequality in the retail sector had a large impact on purchasing-power inequality between the top and bottom income quintiles, equal to about one fourth of the impact of increasing income inequality. Figure 7 suggests that inflation patterns beyond retail may increase inflation inequality even further.⁵⁸

⁵⁸Assuming that there was no inflation inequality beyond retail, purchasing-power inequality increased 23.7% faster than nominal-income inequality. This number is a lower bound according to Figure 7, which suggests that inflation is likely to be lower for higher-income households outside retail as well. Using local house price indices, the analysis in Moretti (2013) suggests that inflation in the housing sector is large for college workers relative to high-school dropouts. However, Diamond (2015) shows that once changes in amenities are taken into account, housing inflation is higher for lower-income groups. My CEX analysis relies on a common price index for the housing sector for all households and ignores these within-housing differences, which is an example of aggregation bias. According to Diamond (2015), my estimates in Figure 7 are conservative. In Appendix F.1, using the demand system and changes in expenditure patterns in the CEX over time, I derive a lower bound of 0.65 percentage point for the overall level of inflation inequality between the top and bottom income quintiles. Although striking, this lower-bound result must be approached with caution because it remains subject to aggregation bias, one of the major findings of Section 3, which pleads for additional work with suitable micro data in other sectors of the economy.

In sum, the two main lessons of this paper regarding inequality are that purchasing-power inequality is increasing faster than commonly thought, at least in the retail sector but probably also beyond, and that changes in nominal income inequality have an amplification effect, because of the response of supply to changes in relative market size. But two caveats should be kept in mind. First, a more unequal income distribution may have other effects on the equilibrium dynamics of innovation that are not captured in my analysis. For instance, because the early adopters are typically high-income households, it could be the case that a more unequal income distribution allows for the introduction of more new technologies that eventually “trickle down” to the rest of the income distribution and benefit everyone (e.g. as in Matsuyama, 2002). My analysis does not speak to this general equilibrium effect. Second, much of the debate about inequality in the US has been revolving around the income share of the top 1%, and my results do not speak to that part of the income distribution, where quality-adjusted consumption is very difficult to measure.⁵⁹

Policy. The various findings in this paper have two broad implications for public policy. First, accurately measuring quality-adjusted inflation across income groups is of the utmost importance. Indeed, I have shown that the inflation difference is large across income groups in the retail sector (cf. Section 3) and is likely to persist beyond retail (cf. Figure 7). Several government transfers are indexed to food-at-home CPI (e.g. food stamps); many others are indexed on the full-basket CPI (e.g. Social Security), as are income poverty thresholds⁶⁰ and tax brackets. The magnitude of the effects I documented in the retail sector is large: from 2004 to 2013, according to the food price index for households who were eligible for food stamps, the nominal increase in food stamp benefits required to preserve purchasing power was 35.5%, instead of the 24.8% actual increase implied by indexation on the overall food CPI. To appropriately account for income-group-specific inflation rates, it appears essential for BLS to improve on its ability to measure income-group-specific spending patterns, so that rigorous measurement of quality-adjusted inflation across income groups becomes possible in all sectors of the economy (as opposed to only in those sectors for which barcode scanner data happens to be available). A first step could be to record information on income in the Telephone Point of Purchase Survey (TPOPS) administered by BLS. Note that with the existing micro-data available to researchers and staff at BLS, it is already possible to measure price changes at different points of the quality (price) distribution within detailed product categories. Combining this information with simple estimates of quality Engel curves within categories (as in Bils and Klenow, 2001) may be sufficient to capture the bulk of the inflation difference across income groups.⁶¹

The second major lesson for public policy is that taking into account the supply response to market size changes induced by policy is key for cost-benefit analysis. Food stamps, the EITC, UI and DI insurance, the

⁵⁹Indeed, the consumption of very high-income households is not well covered in scanner data and, in general, tends to be much more idiosyncratic (e.g. luxury products that are extremely customized and make quality adjustments very difficult, such as luxury cruises).

⁶⁰Following Orshansky (1962), poverty is measured according to an “absolute” scale in the US, which makes the adjustments for non-homothetic price indices even more important than in countries using relative measures of poverty, like most European countries.

⁶¹Indeed, Table C1 showed that inflation inequality largely results from inflation difference across the quality distribution, at least in retail.

minimum wage, Social Security transfers, the possible introduction of a universal basic income, and so on — these policies will all affect the relative market size of different groups of agents, which will induce a targeted response of supply, with price effects which will determine the equilibrium *real* effects of the policy change. In Section 4 and Appendix D.5, I have shown that such effects are large in retail and make food stamp policy more potent than previously understood, because it induces a supply response that lowers the equilibrium price for the recipients. Estimating the equilibrium incidence of other policies in the broader economy is a key task for future research, which could rely on the model developed in Section 5.

Innovation. The various results of the paper show the importance of increasing product variety, and how it differs across income groups. In retail, product innovations are typically simple “customizations”, such as a new flavor or a new size, as opposed to radically new products that usher in a new technological era — like smart phones, electric cars, etc. But these simple product innovations do change people’s lives by providing more variety and lower prices for everyday purchases, which account for an important share of total spending. I have shown that the dynamics of product variety are largely governed by changes in market size, and for that reason they disproportionately benefit high-income households. This stands in contrast with the “product cycle” view, according to which to a first-order approximation innovation benefits everyone equally. The product cycle does characterize some parts of the product space relatively well, e.g. consumer electronics, but in many large sectors of the economy the logic of increasing product variety may be the dominant force at play — I have shown in this paper that it is the case for retail, and a similar logic might apply in other sectors, as suggested by the evidence presented in Section 5.2 using CPI and CEX data.

6 Conclusion

This paper has shown that quality-adjusted inflation was substantially lower for higher-income households in the retail sector between 2004 and 2013, which amplified inequality. The current methodology of statistical agencies like the Bureau of Labor Statistics cannot capture this variation, which arises primarily because income groups differ in their spending patterns along the quality ladder *within* detailed item categories. The Bureau of Labor Statistics currently collects data measuring income-group-specific spending patterns *across* broad item categories, leading to aggregation bias.

Furthermore, the paper has established that product introductions and prices endogenously respond to changes in market size implied by changes in the income distribution, in a way that magnifies purchasing-power inequality. The supply response to market size effects over the past decade explains most of the observed difference in inflation rates across income groups during this period. Finally, analysis of CPI and CEX data on the full consumption basket of American households going back to 1953 supports the prediction from the model that lower-inflation for higher-income households is a secular trend.

This paper raises a number of questions for future research. First, a similar analysis could be carried out with suitable micro data in sectors of the economy beyond retail and in countries other than the United

States. Second, more work is needed to characterize the impact of a range of policies on inflation and product variety across the income distribution, which could fundamentally alter their cost-benefit analysis (e.g. minimum wage laws, monetary policy, sales tax changes, or transfer programs like the EITC). Finally, from a theoretical perspective, it would be fruitful to build on Mirrlees (1971) and adjust optimal redistributive taxation formulas to take into account the endogenous response of supply to changes in market size induced by redistribution.

References

- Acemoglu, Daron**, “Why Do New Technologies Complement Skills? Directed Technical Change and Wage Inequality,” *Quarterly Journal of Economics*, 1998, pp. 1055–1089.
- , “Directed Technical Change,” *The Review of Economic Studies*, 2002, 69 (4), 781–809.
- , “Equilibrium Bias of Technology,” *Econometrica*, 2007, 75 (5), 1371–1409.
- **and David Autor**, “Skills, tasks and technologies: Implications for employment and earnings,” *Handbook of labor economics*, 2011, 4, 1043–1171.
- **and Joshua Linn**, “Market Size in Innovation: Theory and Evidence from the Pharmaceutical Industry,” *Quarterly journal of economics*, 2004, (3), 1049–1090.
- Aghion, Philippe and Peter Howitt**, “A Model of Growth through Creative Destruction,” *Econometrica*, 1992, 60 (2).
- , **Nicholas Bloom, Richard Blundell, Rachel Griffith, and Peter Howitt**, “Competition and Innovation: An Inverted U Relationship?,” *Quarterly Journal of Economics*, 120(2), 701-728, 2005.
- Aguiar, Mark and Erik Hurst**, “Consumption versus Expenditure,” *Journal of Political Economy*, 2005, 113 (5).
- **and –** , “Life-cycle prices and production,” *The American Economic Review*, 2007, 97 (5), 1533–1559.
- Argente, David and Munseob Lee**, “Cost of Living Inequality during the Great Recession,” *Working Paper*, 2016.
- Autor, David and David Dorn**, “The growth of low-skill service jobs and the polarization of the US labor market,” *The American Economic Review*, 2013, 103 (5), 1553–1597.
- , – , **and Gordon Hanson**, “The China syndrome: Local labor market effects of import competition in the United States,” *The American Economic Review*, 2013, 103 (6), 2121–2168.
- , **Lawrence Katz, and Alan Krueger**, “Computing inequality: have computers changed the labor market?,” *Quarterly Journal of Economics*, 1998, 113 (4).

- Bartik, Timothy**, “Who Benefits from State and Local Economic Development Policies?,” *W.E. Upjohn Institute for Employment Research*, 1991.
- Bell, Alex, Raj Chetty, Xavier Jaravel, Neviana Petkova, and John Van Reenen**, “The lifecycle of inventors,” *Working Paper*, 2015.
- Bell, David, Jose Alvarez, James Weber, and Mary Shelman**, “The Hain Celestial Group,” *Harvard Business School Case 516-007*, 2015.
- Beraja, Martin, Erik Hurst, and Juan Ospina**, “The Aggregate Implications of Regional Business Cycles,” *Working Paper*, 2016.
- Berry, Steve and Peter Reiss**, “Empirical models of entry and exit,” *Handbook of Industrial Organization*, 2006, 3.
- Blanchard, Olivier and Lawrence Katz**, “Regional evolutions,” *Brookings papers on economic activity*, 1992, 1992 (1), 1–75.
- Boppart, Timo and Franziska Weiss**, “Non-homothetic preferences and industry directed technical change,” *Working Paper*, 2016.
- Borusyak, Kirill and Xavier Jaravel**, “A Note on Bartik Research Designs,” *Working Paper*, 2016.
- Broda, Christian and David Weinstein**, “Globalization and the Gains From Variety,” *The Quarterly Journal of Economics*, 2006, 121 (2), 541–585.
- and – , “Product Creation and Destruction: Evidence and Price Implications,” *The American economic review*, 2010, 100 (3), 691–723.
- and **John Romalis**, “The welfare implications of rising price dispersion,” *University of Chicago mimeo*, 2009, 3.
- Chevalier, Judith and Anil Kashyap**, “Best Prices: Price Discrimination and Consumer Substitution,” *Working Paper*, 2014.
- Coibion, Olivier, Yuriy Gorodnichenko, and Gee Hee Hong**, “The cyclicity of sales, regular and effective prices: Business cycle and policy implications,” *The American Economic Review*, 2015, 105 (3), 993–1029.
- Comin, Diego A, Danial Lashkari, and Martí Mestieri**, “Structural Change with Long-Run Income and Price Effects,” *Working Paper*, 2015.
- Diamond, Rebecca**, “The determinants and welfare implications of US workers’ diverging location choices by skill: 1980–2000,” *The American Economic Review*, 2016, 106 (3), 479–524.

- Diewert, Erwin**, “Exact and superlative index numbers,” *Journal of econometrics*, 1976, 4 (2), 115–145.
- Eizenberg, Alon**, “Upstream innovation and product variety in the us home pc market,” *The Review of Economic Studies*, 2014, 81 (3), 1003–1045.
- Erickson, Tim and Ariel Pakes**, “An Experimental Component Index for the CPI: From Annual Computer Data to Monthly Data on Other Goods,” *The American Economic Review*, 2011, 101 (5), 1707–1738.
- Faber, Benjamin and Thibault Fally**, “Firm Heterogeneity in Consumption Baskets: Evidence from Home and Store Scanner Data,” *Working Paper*, 2015.
- Feenstra, Robert**, “New product varieties and the measurement of international prices,” *The American Economic Review*, 1994, pp. 157–177.
- , “A homothetic utility function for monopolistic competition models, without constant price elasticity,” *Economics Letters*, 2003, 78 (1), 79–86.
- Ganong, Peter and Jeffrey Liebman**, “The decline, rebound, and further rise in SNAP enrollment: Disentangling business cycle fluctuations and policy changes,” *Working Paper*, 2013.
- Goldsmith-Pinkham, Paul, Isaac Sorkin, and Henry Swift**, “Econometrics Properties of Bartik and Simulated Instruments,” *Working Paper*, 016.
- Gopinath, Gita, Pierre-Olivier Gourinchas, Chang-Tai Hsieh, and Nicholas Li**, “International prices, costs, and markup differences,” *The American Economic Review*, 2011, 101 (6), 2450–2486.
- Hall, Bronwyn, Adam Jaffe, and Manuel Trajtenberg**, “The NBER patent citation data file: Lessons, insights and methodological tools,” *National Bureau of Economic Research*, 2001.
- Handbury, Jessie**, “Are poor cities cheap for everyone? non-homotheticity and the cost of living across us cities,” *The Wharton School Research Paper*, 2013, (71).
- Hausman, Jerry**, “Sources of bias and solutions to bias in the consumer price index,” *the Journal of Economic Perspectives*, 2003, 17 (1), 23–44.
- Hobijn, Bart and David Lagakos**, “Inflation inequality in the United States,” *Review of income and Wealth*, 2005, 51 (4), 581–606.
- Jaravel, Xavier**, “Consumer Aging and Secular Stagnation,” *Working Paper*, 2016.
- Jones, Charles**, “R & D-based models of economic growth,” *Journal of political Economy*, 1995, pp. 759–784.
- , “Time series tests of endogenous growth models,” *The Quarterly Journal of Economics*, 1995, pp. 495–525.
- Kaplan, Greg and Sam Schulhofer-Wohl**, “Inflation at the Household Level,” *Working Paper*, 2016.

- Kopczuk, Wojciech, Emmanuel Saez, and Jae Song**, “Earnings inequality and mobility in the United States: evidence from social security data since 1937,” *The Quarterly Journal of Economics*, 2010, *125* (1), 91–128.
- Matsuyama, Kiminori**, “The rise of mass consumption societies,” *Journal of political Economy*, 2002, *110* (5), 1035–1070.
- Mayer, Thierry, Marc Melitz, and Gianmarco Ottaviano**, “Market size, competition, and the product mix of exporters,” *The American Economic Review*, 2014, *104* (2), 495–536.
- McGranahan, Leslie and Anna Paulson**, “Constructing the Chicago Fed Income Based Economic Index–Consumer Price Index: Inflation Experiences by Demographic Group: 1983–2005,” *Federal Reserve Bank of Chicago Working Paper*, 2006.
- Melitz, Marc**, “The impact of trade on intra-industry reallocations and aggregate industry productivity,” *Econometrica*, 2003, *71* (6), 1695–1725.
- **and Gianmarco Ottaviano**, “Market size, trade, and productivity,” *The review of economic studies*, 2008, *75* (1), 295–316.
- Mirrlees, James**, “An exploration in the theory of optimum income taxation,” *The review of economic studies*, 1971, *38* (2), 175–208.
- Moretti, Enrico**, “Real Wage Inequality,” *American Economic Journal: Applied Economics*, 2013, *5* (1), 65–103.
- Pakes, Ariel**, “A Reconsideration of Hedonic Price Indexes with an Application to PC’s,” *The American Economic Review*, 2003, *93* (5), 1578–1596.
- Pisano, Luigi and Andrea Stella**, “Price Heterogeneity and Consumption Inequality,” *Working Paper*, 2015.
- Romer, Paul**, “Endogenous Technological Change,” *Journal of Political Economy*, 1990, *98* (5 pt 2).
- Sato, Kazuo**, “The ideal log-change index number,” *The Review of Economics and Statistics*, 1976, pp. 223–228.
- Schumpeter, Joseph**, “Socialism and democracy,” 1942.
- Song, Jae, David Price, Fatih Guvenen, Nicholas Bloom, and Till von Wachter**, “Firming up inequality,” *Working Paper*, 2015.
- Stroebel, Johannes and Joseph Vavra**, “House prices, local demand, and retail prices,” *Working Paper*, 2014.

Sutton, John, *Sunk costs and market structure: Price competition, advertising, and the evolution of concentration*, MIT press, 1991.

Vartia, Yrjö O, “Ideal log-change index numbers,” *Scandinavian Journal of Statistics*, 1976, pp. 121–126.

Vernon, Raymond, “International investment and international trade in the product cycle,” *The Quarterly Journal of Economics*, 1966, pp. 190–207.

Zhelobodko, Evgeny, Sergey Kokovin, Mathieu Parenti, and Jacques-Francois Thisse, “Monopolistic competition: Beyond the constant elasticity of substitution,” *Econometrica*, 2012, 80 (6), 2765–2784.

Appendix

A Data Appendix

A.1 Data Sources

A.1.1 Manufacturer Identifier Data

In order to measure manufacturer entry and competition, I use data from GS1, the company in charge of allocating bar codes in the US, on the universe of barcodes and manufacturers. I match the bar codes observed in the Nielsen data to manufacturers using the first few digits of the bar code - the match rate is close to 95%. Since the cutoff size for a manufacturer to appear in this dataset is to make a sale rather than an arbitrary number of workers, I can observe the full distribution of manufacturers in each product group. There are about 500 manufacturers on average in each product group, with 90 percent of the product groups having more than 200 manufacturers. The median number of products supplied by a manufacturer is 5 and the average is 14.

Consistent with the findings reported by Hottman et al. (2016), while on average half of all output in a product group is produced by just five manufacturers, around 98 percent of manufacturers have market shares below 2 percent. Thus, the typical product group is characterized by a few large manufacturers and a competitive fringe of manufacturers with very low market shares. Another important feature of the data is that even the largest manufacturers are not close to being monopolists: the largest manufacturer in a product group on average has a market share of 22 percent. The model presented in Section 5 is consistent with these patterns.

A.1.2 Retailer Markup Data

To test specific predictions of the model in Section 5, I use data on retailer markups. I have access to weekly product-level data between January 2004 and June 2007 in 19 U.S. states, for 250 grocery stores operated by a single retail chain. This dataset contains information for 125,048 unique products (UPCs), mostly in the food and beverages categories, housekeeping supplies, books and magazines, and personal care products. Most of the stores are located in the western and eastern corridors, in the Chicago area, Colorado and Texas. For every store in every week, data is available on the price, the wholesale cost and the marginal cost of each product. I infer the markups of the retailer based on the price and wholesale cost. Note that I do not measure other costs like labor, rent and utilities. In the analysis carried out in Section 4, store-year fixed effects are used to absorb these costs. The dataset also reports “adjusted gross profits” per unit for each product, defined as the net price minus the sum of wholesale costs and transportation costs plus net rebates from the manufacturer - I use this adjustment in robustness checks.

In addition, I can measure wholesale prices from 2006 to 2011 using data from National Promotion Reports’ PRICE-TRAK database. These data contain wholesale price changes and deal offers by UPC in 48 markets during this period, along with associated product attributes such as item and pack sizes. The data

are sourced from one major wholesaler in each market, which is representative due to the provisions of the Robinson-Patman (Anti-Price Discrimination) Act. I compute retail margins by matching wholesale prices with retail prices by UPC, item size, and year.

A.1.3 BLS Consumer Price Index and Consumer Expenditure Survey Data

In order to provide suggestive evidence about the external validity of the findings obtained with the Nielsen data, I rely on additional data and find that the results are likely to extend to earlier periods and to other product groups. Specifically, I use the Consumer Expenditure Survey (CEX) to compute the full consumption baskets of various income and education groups. In order to price the items in these consumption baskets, I manually match the various CEX product categories to 48 item-specific Consumer Price Index (CPI) data series. These price series extend back to 1953 and I thus obtain estimates of income-group-specific inflation rates for the full consumption basket over a long time horizon. The results are reported in Section 5 and support the idea that the findings obtained in the Nielsen sample apply more broadly.⁶²

The product categories are matched by hand and are as follows: cereals, bakery, beef, pork, other meat, poultry, fish, egg, dairy, fresh fruit, fresh vegetables, sugar, fat and oils, other food, beverages, food away from home, beer at home, whiskey at home, wine at home, spirits at home, alcohol away from home, shelter, rent, fuel, utilities, electricity, oil, water, furniture, men's apparel, boys' apparel, girls' apparel, infants' apparel, footwear, other apparel, new vehicles, used vehicles, motor fuel, vehicle maintenance, vehicle insurance, public transportation, medical care products, medical care services, tobacco, personal care products, personal care services.

A.1.4 More on Nielsen Scanner Data

Description of Homescan Consumer Panel Data: I primarily rely on the Home Scanner Database collected by AC Nielsen and made available through the Kilts Center at The University of Chicago Booth School of Business. AC Nielsen collects these data using hand-held scanner devices that households use at home after their shopping in order to scan each individual transaction they have made. Faber and Fally (2015) report that on average each semester covers \$105 million worth of retail sales across 58,000 individual, across more than 500,000 barcodes belonging to 180,000 brands.

Description of Retail Scanner Data: The Retail Scanner Data consist of weekly price and quantity information for more than one hundred retail chains across all US markets between January 2006 and December 2013. The database includes about 45,000 individual stores. The stores in the database vary in terms of the channel they represent: e.g. food, drug, mass merchandising, liquor, or convenience stores. Faber and Fally (2015) report that on average each semester covers \$110 billion worth of retail sales across 25,000 individual stores, across more than 700,000 barcodes belonging to 170,000 brands.

The strength of the home scanner database is the detailed level of budget share information that it

⁶²These results are based on relatively aggregated data and are therefore much cruder than those obtained with the Nielsen microdata. But the consistency of the results across samples is striking.

provides alongside household characteristics. Its relative weakness in the comparison to the store-level retail scanner data is that the home scanner samples households and, therefore, has higher sampling error at the product level. Relative to the home scanner data, the store-level retail scanner data records more than one thousand times the retail sales in each semester. I primarily rely on the home scanner data in the paper, but I present robustness checks based on the retail scanner data.

Examples. The food industry has undergone a revolution in the past fifteen years, with the rise of organic and natural food products, which illustrates the price and quantity dynamics discussed in the paper particularly well. As shown in Figure A1, organic products constitute an increasing share of the market and their price relative to nonorganic food products has been steadily decreasing. For a detailed study of the sector, the [US Department of Agriculture’s Economic Research Service report](#).

Another particularly good example illustrating the forces discussed in the paper is the market for snacks. In recent years, meat snacks have grown tremendously - for instance premium beef jerky, with sustained double-digit growth for over five years nationwide. Premium beef jerky is a high-protein, low-fat and low-calorie snack - a practical and healthy snack that particularly appeals to young and high-income households. The branding of premium beef jerky is fundamentally different from that of traditional jerky - favorite of truckers and staple of gas-station checkouts - and so is its production process. In particular, many of the varieties of premium beef jerky are fully organic - for instance, beef jerky made from 100% grass-fed cattle from networks of small family farms. The so-called “jerky renaissance” is largely driven by demand. It is answering the demand of high-income consumers concerned with healthy living and eager to support a sustainable, more humane agriculture. And it is taking place in a broader context of increased demand for snacks - a Nielsen survey found that one in ten Americans say they eat snacks instead of meals - and for proteins - according to the NPD group, more than half of Americans say they want more protein in their diet. The competition for the premium beef jerky market has intensified in recent years, with an ever-increasing number of small, local players but also with the entry of established companies through acquisitions. For instance, Krave, one of the early players in premium jerky who led the market in the late 2000s, was acquired in 2015 by Hershey’s, the largest chocolate manufacturer in North America. Accordingly, premium beef jerky prices have fallen and varieties have increased. Similar - although less spectacular - dynamics are visible in other segments of the snack industry, like hummus and protein bars, but not so in segments catering to lower-income consumers, like chips, bars and nuts.

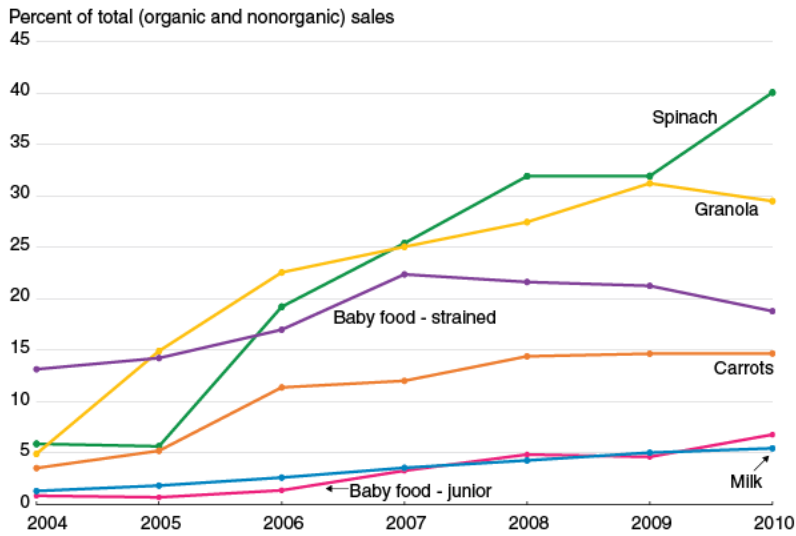
Another case in point is craft beer - the number of microbreweries in the United States went from about 30 in the early 1990s to 300 in the early 2000 to more than 3,000 today, and the relative price of craft beers relative to entry-level beers has plummeted.⁶³

⁶³Source: <https://www.brewersassociation.org>.

Figure A1: The Rise of Organic Food Products

Panel A: Quantities

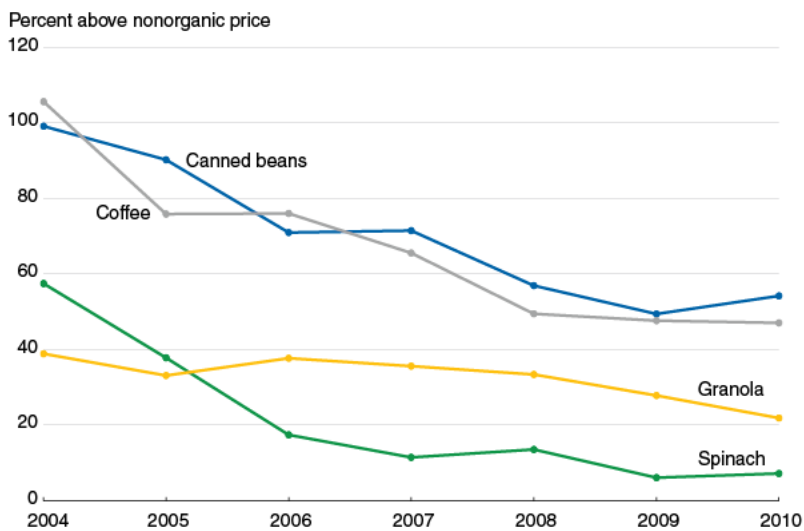
Organics' share of total product sales are rising



Source: USDA, Economic Research Service using 2004-10 Nielsen Homescan data.

Panel B: Relative Prices

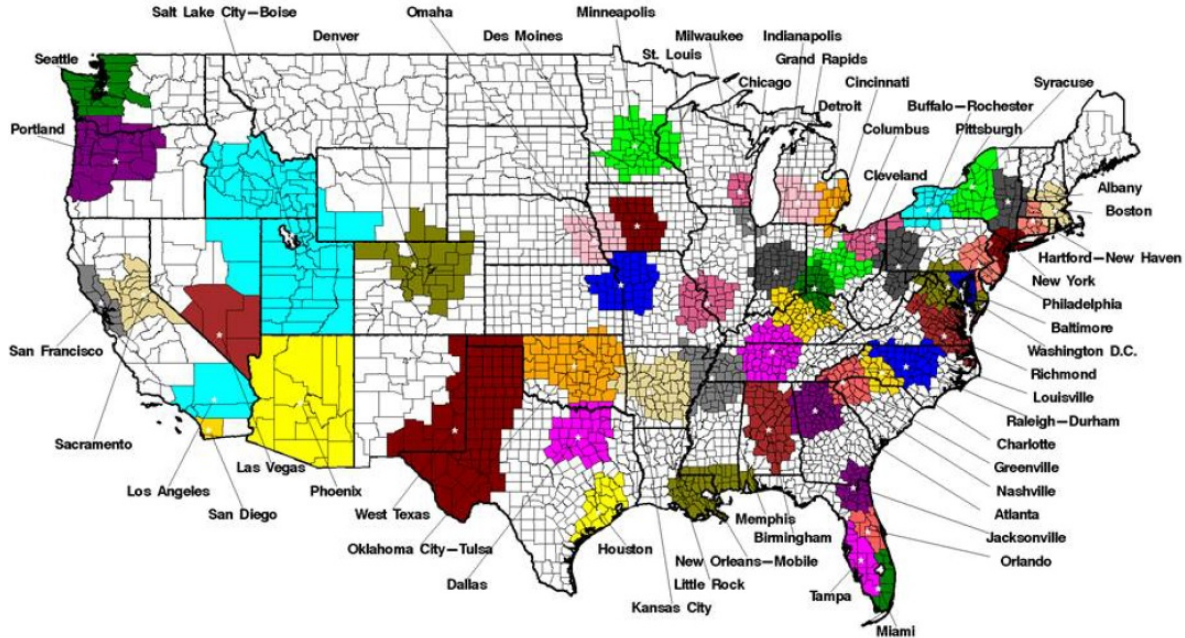
Organic price premiums for only a few of the products studied have fallen over 2004-10



Source: USDA, Economic Research Service using 2004-10 Nielsen Homescan data.

Local Markets: Both the home scanner and retail scanner data can be disaggregated into 76 local markets, which are shown on the map below.

Figure A2: Map of the 76 Local Markets Tracked in the Nielsen Datasets



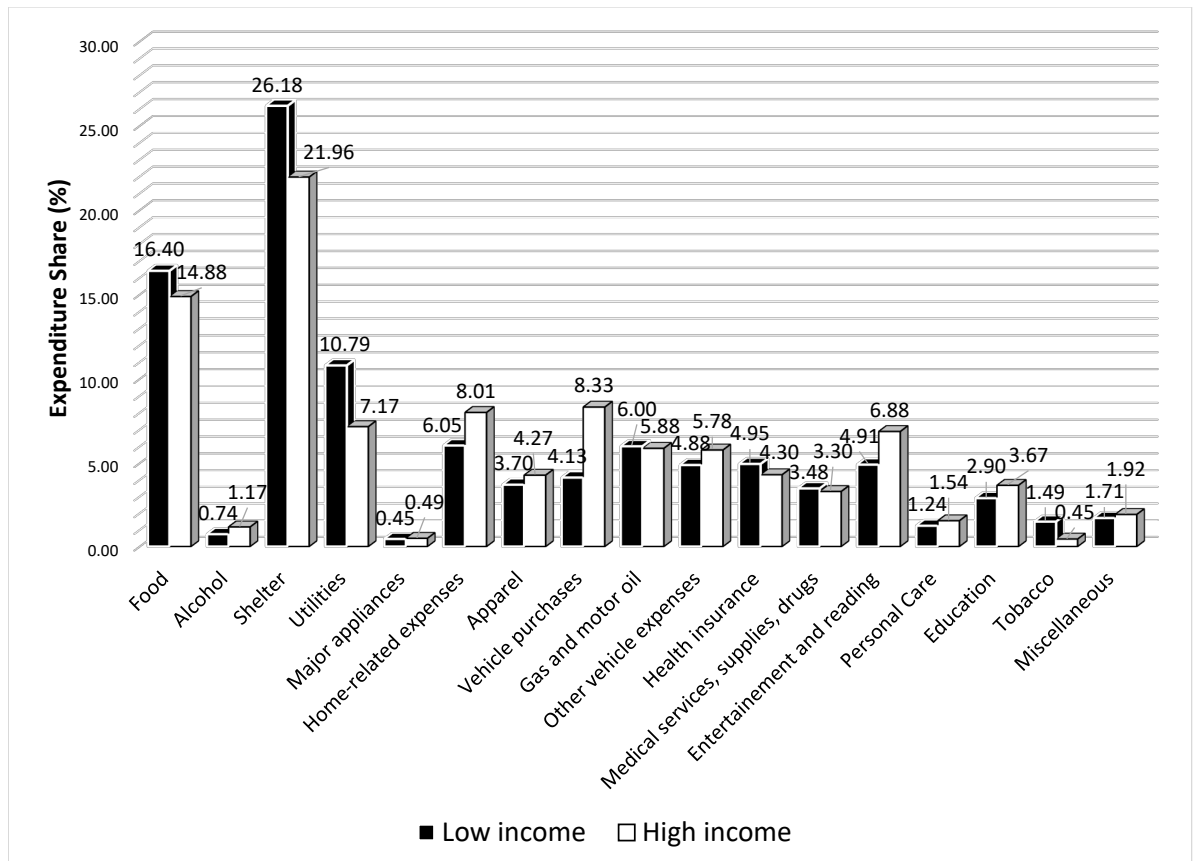
A.2 Summary Statistics

Table A1 compares aggregate spending share in the Nielsen scanner data with the Consumer Price Index for all urban consumers (CPI-U) and the Consumer Expenditure Survey (CEX). As expected, the Nielsen products are not representative of the full consumption basket. Accordingly, in order to probe the external validity of the findings based on the Nielsen data, I extend the analysis using CPI and CEX data in Section 5. Although spending shares differ between Nielsen and the full consumption basket, price series do not: in a given expenditure category, the price indices built from the Nielsen data closely match the patterns from the CPI (Beraja et al., 2016; Kaplan and Schulhofer-Wohl, 2016). Figure A3 reports the spending shares of the top and bottom income quintiles across categories.

Table A1: Comparing Spending in Nielsen Basket and Full Consumption Basket

Spending Category	Expenditure Shares (%)		
	CPI-U	CEX	Nielsen
Food and beverages	14.8	16.2	58.8
Food	13.2	14.9	55.2
Food at home	8.6	8.9	53.1
Cereals and bakery products	1.2	1.2	7.7
Cereal products	0.4	0.4	2.9
Bakery products	0.8	0.8	4.8
Meats, poultry, fish, and eggs	2.0	1.9	7.5
Meats, poultry, and fish	1.8	1.7	6.7
Eggs	0.2	0.2	0.5
Dairy and related products	0.9	0.8	8.1
Fruits and vegetables	1.3	1.6	7.2
Nonalcoholic beverages, beverage materials	0.9	0.7	6.9
Other food at home	2.4	2.2	14.8
Sugar and sweets	0.4	0.3	2.8
Fats and oils	0.4	0.3	1.4
Other foods	1.6	1.6	10.4
Food away from home	5.6	6.1	4.1
Alcoholic Beverages	1.0	1.0	3.1
Housing	41.9	35.7	9.3
Shelter	31.1	22.6	0
Fuels and utilities	5.1	5.4	0.1
Household furnishings and operations	4.0	7.6	9.1
Window and floor coverings and other linens	0.3	0.3	0
Furniture and bedding	0.7	0.9	0
Appliances	0.3	0.7	1.3
Other household equipment and furnishings	0.6	0	1.0
Tools, hardware, outdoor equipment, supplies	0.7	0	1.1
Housekeeping supplies	0.8	1.4	5.8
Household operations	0.8	2.8	0
Apparel	3.7	4.0	8.2
Transportation	16.7	20.4	0.2
Private transportation	15.5	19.1	0.2
Public transportation	1.2	1.1	0
Medical care	7.5	8.1	6.9
Recreation	5.8	6.3	6.3
Video and audio	1.9	2.3	2.2
Pets, pet products and services	1.2	0	4.1
Sporting goods	0.4	0	0
Photography	0.1	0	0.2
Other recreational goods	0.6	0	0.1
Other recreational services	1.8	0	0
Recreational reading materials	0.2	0	0
Education and communication	6.9	5.7	0
Others goods and services	3.2	4.0	7.9
Tobacco and smoking products	0.8	0.9	1.8
Personal care	2.6	1.5	4.4

Figure A3: Expenditures Share of Top and Bottom Income Quintiles (2012 CEX data)



Notes: This figure reports expenditure shares across expenditure categories for households in the top income quintile (high income) and in the bottom income quintile (low income). The data is from the 2012 CEX survey.

B Price Index Appendix

B.1 Chaining

An important consideration is whether or not to chain the price index. In a chain index, each link consists of an index in which each period is compared with the preceding one, the weight and price reference being moved forward each period. A chain index is therefore path dependent: it depends on the prices and quantities in all the intervening periods between the first and last period in the index series. When there is a gradual economic transition from the first to the last period, chaining is advantageous because it smoothes trends in relative prices and quantities and tends to reduce the index number spread between the various price indices listed above.

But if there are fluctuations in the prices and quantities in the intervening periods, chaining may not only increase the index number spread but also distort the measure of the overall change between the first and last periods. For example, suppose all the prices in the last period return to their initial levels in period 0, which implies that they must have fluctuated in between. A chain Laspeyres index will not return to 100: it will tend to be greater than 100. If the cycle is repeated with all the prices periodically returning to their original levels, a chain Laspeyres index will tend to drift further and further above 100 even though there may be no long-term upward trend in the prices. Chaining is therefore not advised when the price fluctuates.

Accordingly, I present robustness checks with and without chaining the indices.

B.2 Estimation Equations for (Nested) CES Exact Price Index

Given the formula reported in the main text, we only need to estimate the module-specific elasticities. We do this by first modeling the supply and demand conditions for each good within a module.

The demand equation comes from the following transformation, which exploits the panel nature of the data:

$$\begin{aligned} \ln(s_{umgt}) - \ln(s_{umg(t-1)}) &= \Delta \ln(s_{umgt}) \\ &= (1 - \sigma_m) [\ln(p_{umgt}) - \ln(p_{umg(t-1)})] + \ln(P_{mgt}) - \ln(P_{mg(t-1)}) \\ &= (1 - \sigma_m) [\ln(p_{umgt}) - \ln(p_{umg(t-1)})] + \lambda_{mt} \end{aligned}$$

where the second line uses (1) and the fact that quality/taste is assumed to be constant over time. The fixed effect corresponds to the change in the price index of the module. In practice, there will be an estimation error, which for instance could come from yearly change in taste (which would affect the d parameters). So we can write the demand curve as:

$$\Delta \ln(s_{umgt}) = (1 - \sigma_m) \Delta \ln(p_{umgt}) + \lambda_{mt} + \epsilon_{umgt}$$

Then, we assume an isoelastic supply curve (with $\alpha > 0$ assumed to be the same for all UPCs within a module):

$$\begin{aligned} \ln(c_{umgt}) &= \alpha \ln(p_{umgt}) + \chi_{mg} \\ \ln(s_{umgt}) &= \alpha \ln(p_{umgt}) - \ln(E_{mgt}) + \chi_{mg} \end{aligned}$$

Differencing over time:

$$\ln(s_{umgt}) - \ln(s_{umg(t-1)}) = \alpha [\ln(p_{umgt}) - \ln(p_{umg(t-1)})] + \ln(E_{mgt}) - \ln(E_{mg(t-1)})$$

so

$$\Delta \ln(p_{umgt}) = \frac{1}{\alpha} \Delta \ln(s_{umgt}) - \frac{1}{\alpha} \Delta \ln(E_{mgt}) = \frac{1}{\alpha} \Delta \ln(s_{umgt}) + \psi_{mgt}$$

The fixed effect corresponds to the change in total expenditures in the module (which is observed). In practice there will be estimation errors, e.g. due to assembly line shocks, so we write:

$$\Delta \ln(p_{umgt}) = \frac{1}{\alpha} \Delta \ln(s_{umgt}) + \psi_{mgt} + \delta_{umgt}$$

We now want to eliminate the fixed effects in the demand and supply equations. We take a difference relative to the UPC k with the largest market share:

$$\Delta^k \ln(s_{umgt}) = (1 - \sigma_m) \Delta^k \ln(p_{umgt}) + \epsilon_{umgt}^k \quad (20)$$

$$\Delta^k \ln(p_{umgt}) = \frac{1}{\alpha} \Delta^k \ln(s_{umgt}) + \delta_{umgt}^k \quad (21)$$

with $\Delta^k X = \Delta X_{umgt} - \Delta X_{kmgt}$, $\epsilon_{umgt}^k = \epsilon_{umgt} - \epsilon_{kmgt}$ and $\delta_{umgt}^k = \delta_{umgt} - \delta_{kmgt}$.

Now we can set up the moment condition, based on the assumption that the upc-specific demand and supply shocks are uncorrelated over time, i.e $E_t[\epsilon_{umgt}^k \delta_{umgt}^k] = 0$.

$$v_{umgt} = \epsilon_{umgt}^k \times \delta_{umgt}^k$$

$$G(\beta_m) = E_t(v_{umgt}(\beta_m)) = 0 \quad \forall u, m \text{ and } g$$

This can be written as:

$$\begin{aligned} v_{umgt}(\beta_m) &= \epsilon_{umgt}^k \times \delta_{umgt}^k \\ &= (\Delta^k \ln(s_{umgt}) - (1 - \sigma_m) \Delta^k \ln(p_{umgt})) \times \left(\Delta^k \ln(p_{umgt}) - \frac{1}{\alpha} \Delta^k \ln(s_{umgt}) \right) \\ &= (\sigma_m - 1) (\Delta^k \ln(p_{umgt}))^2 - \frac{1}{\alpha} (\Delta^k \ln(s_{umgt}))^2 + \frac{\alpha + (1 - \sigma_m)}{\alpha} \Delta^k \ln(s_{umgt}) \Delta^k \ln(p_{umgt}) \end{aligned}$$

The moment condition $E_t[v_{umgt}(\beta_m)] = 0$ means:

$$E_t \left[(\Delta^k \ln(p_{umgt}))^2 \right] = \frac{1}{\alpha(\sigma_m - 1)} E_t \left[(\Delta^k \ln(s_{umgt}))^2 \right] - \frac{\alpha + (1 - \sigma_m)}{\alpha(\sigma_m - 1)} E_t \left[\Delta^k \ln(s_{umgt}) \Delta^k \ln(p_{umgt}) \right] \quad \forall u, m \text{ and } g$$

Rewriting $\alpha \equiv \frac{1+\omega_m}{\omega_m}$ yields:

$$E_t \left[(\Delta^k \ln(p_{umgt}))^2 \right] = \underbrace{\frac{\omega_m}{(1+\omega_m)(\sigma_m-1)}}_{\equiv \theta_1} E_t \left[(\Delta^k \ln(s_{umgt}))^2 \right] - \underbrace{\frac{1-\omega_m(\sigma_m-2)}{(1+\omega_m)(\sigma_m-1)}}_{\equiv \theta_2} E_t \left[\Delta^k \ln(s_{umgt}) \Delta^k \ln(p_{umgt}) \right]$$

The parameters ω_m and σ_m are estimated under the restriction that $\omega_m > 0$ and $\sigma_m > 1$. To do this, θ_1 and θ_2 are first estimated by weighted least squares, as in Feenstra (1994). Then I go back to the primitive parameters. If this produces imaginary estimates or estimates of the wrong sign, a grid search is performed for the objective function for values of $\sigma_m \in [1.05, 131.5]$ at intervals that are 5 percent apart.

Given estimates for σ_m , the price index formula given in equation (2) can be implemented. An example of how equation (2) captures the impact of different types of creation and destruction is given in Broda and Weinstein (2010): “Let’s consider the case of a new type of sunscreen that replaces an earlier type. If the new sunscreen is just a repackaging of last year’s sunscreen without a noticeably different quality or price, then, ceteris paribus, the new sunscreen will have a market share equal to that of the old sunscreen. If this is true, then the share of common goods will be unchanged and our measured quality bias from the replacement of the old model would be zero. If, instead, the new sunscreen is priced identically but is of a higher quality than the old model, then, ceteris paribus, its market share will rise. This result comes directly from the optimizing behavior of the consumer, because the new sunscreen will have a lower price per unit quality than the old sunscreen. If this is the case, the higher share of the new good relative to the old good implies that there is a “quality bias” in the conventional price index that only considers products existing across periods.”

B.3 Estimation Equations for Translog Exact Price Index

This section derives estimation equation for the translog price index. The translog price index is used as a robustness check for the results with the CES price index presented in Section 3. The results are very similar and are presented in Appendix Table C3. Moreover, the translog price index comes out of the model developed Section 5, where translog preferences are nested in CES preferences. The derivations below show how to estimate the translog elasticity parameter of the lower nest of the preference structure from Section 5. Since the logic of the derivation of the estimation equations for translog is similar to what was done for CES in Appendix Section B.2, only the main steps are reported here.

The derivation of the estimation equations borrows from Feenstra and Weinstein (2016), who derive similar equations at the firm level (rather than at the variety level). As shown in Feenstra and Weinstein (2016), with a changing set of varieties, at time t the rate of inflation in the translog price index for product module m is given by:

$$\tilde{\pi}_m = \tilde{P}_m \cdot \exp\left(\frac{\tilde{\lambda}_{mt} - \tilde{\lambda}_{m(t-1)}}{\gamma_m}\right) \cdot \exp\left(\frac{-\left(\sum_i s_{it}^2 - \sum_i s_{i(t-1)}^2\right)}{2\gamma_m}\right)$$

where \tilde{P}_m is the Tornqvist price index introduced in Section 3.2, $\tilde{\lambda}_{mt}$ is the share of spending at time t on products that did not exist at time $t-1$, $\tilde{\lambda}_{m(t-1)}$ is the share of spending at time $t-1$ on products that no

longer exist at time t , s_{it} denotes the spending share on i at time t . Intuitively, the second term means that inflation declines when product variety increases, and the third term means that inflation increases when the Herfindahl index of expenditure shares decreases. These terms reflect the fact that the translog demand system features love-of-variety (term 2), but at a declining rate (term 3). As more products get introduced, the product space becomes more crowded and consumers benefit less from an additional variety (see Section 5 for a discussion of these effects and their implications for markups).

The key object to estimate is γ_m , which govern the degree of substitutability between products within product module m . If γ_m takes a very large value, i.e. if product are very substitutable, then inflation is correctly by the Tornqvist price index on continued products: $\tilde{\pi}_m \rightarrow \tilde{P}_m$ as $\gamma_m \rightarrow \infty$. The intuition is similar to the CES case: when products are highly substitutable, a law of one price applies.

To estimate γ_m , it is important to take into account that prices are endogenous, as in a conventional supply and demand system. For the CES case, Feenstra (1994) showed how this endogeneity could be overcome by specifying the supply equation and assuming that the demand and supply errors are uncorrelated. Identification of the model parameters from this moment condition depended on having heteroskedasticity in second-moments of the data, so this is an example of “identification through heteroskedasticity,” as discussed more generally by Rigobon (2003). A similar logic applies to the translog case.

The key equations for estimation are as follows. On the supply side, we have the pricing equation:

$$(1 + \omega)\ln(p_{it}) = \omega_{i0} + \omega\ln(s_{it}) + \omega\ln(E_t) + \ln\left(1 + \frac{s_{it}N_t}{\gamma(N_t - 1)}\right) + \delta_{it} \quad (22)$$

where δ_{it} is an idiosyncratic supply shock. The demand side is characterized by the equation:

$$s_{it} = \alpha_{it} + \alpha_t - \gamma(\ln(p_{it}) - \overline{\ln(p_t)})$$

Next, model α_{it} as a barcode fixed effect plus an idiosyncratic error term: $\alpha_{it} = \alpha_i + \epsilon_{it}$. The demand curve becomes:

$$s_{it} = \alpha_i + \alpha_t - \gamma(\ln(p_{it}) - \overline{\ln(p_t)}) + \epsilon_{it} \quad (23)$$

Equations (22) and (23) can now be transformed to yield the estimation equation. We difference with respect to product k and with respect to time, thereby eliminating the terms $\alpha_i + \alpha_t$ and the overall average prices $\overline{\ln(p_t)}$. Denoting $X_t - X_{t-1}$ by ΔX , we obtain:

$$\begin{aligned} \frac{\Delta\delta_{it} - \Delta\delta_{kt}}{(1 + \omega)} &= (\Delta\ln(p_{it}) - \Delta\ln(p_{kt})) - \frac{\omega}{(1 + \omega)} (\Delta\ln(s_{it}) - \Delta\ln(s_{kt})) - \frac{1}{1 + \omega} \left(\Delta\ln\left(1 + \frac{s_{it}N_t}{\gamma(N_t - 1)}\right) - \Delta\ln\left(1 + \frac{s_{kt}N_t}{\gamma(N_t - 1)}\right) \right) \\ \frac{\Delta\epsilon_{it} - \Delta\epsilon_{kt}}{\gamma} &= \frac{\Delta s_{it} - \Delta s_{kt}}{\gamma} + (\Delta\ln(p_{it}) - \Delta\ln(p_{kt})) \end{aligned}$$

Multiplying these equations together yields:

$$\bar{Y}_i = \frac{\omega}{1 + \omega} \bar{X}_{1i} + \frac{\omega}{\gamma(1 + \omega)} \bar{X}_{2i} - \frac{1}{\gamma} \bar{X}_{3i} + \frac{1}{1 + \omega} \bar{Z}_{i1}(\gamma) + \frac{1}{\gamma(1 + \omega)} \bar{Z}_{i2}(\gamma) + \bar{u}_i \quad (24)$$

where the over-bar indicates averaging variables over time and:

$$\begin{aligned}
Y_{it} &= (\Delta \ln(p_{it}) - \Delta \ln(p_{kt}))^2 \\
X_{1it} &= (\Delta \ln(s_{it}) - \Delta \ln(s_{kt})) \cdot (\Delta \ln(p_{it}) - \Delta \ln(p_{kt})) \\
X_{2it} &= (\Delta \ln(s_{it}) - \Delta \ln(s_{kt})) \cdot (\Delta s_{it} - \Delta s_{kt}) \\
X_{3it} &= (\Delta \ln(p_{it}) - \Delta \ln(p_{kt})) \cdot (\Delta s_{it} - \Delta s_{kt}) \\
Z_{1it} &= \left(\Delta \ln\left(1 + \frac{s_{it}N_t}{\gamma(N_t - 1)}\right) - \Delta \ln\left(1 + \frac{s_{kt}N_t}{\gamma(N_t - 1)}\right) \right) \cdot (\Delta \ln(p_{it}) - \Delta \ln(p_{kt})) \\
Z_{2it} &= \left(\Delta \ln\left(1 + \frac{s_{it}N_t}{\gamma(N_t - 1)}\right) - \Delta \ln\left(1 + \frac{s_{kt}N_t}{\gamma(N_t - 1)}\right) \right) \cdot (\Delta s_{it} - \Delta s_{kt}) \\
u_{it} &= \frac{(\Delta \epsilon_{it} - \Delta \epsilon_{kt})(\Delta s_{it} - \Delta s_{kt})}{\gamma(1 + \omega)}
\end{aligned}$$

Under the assumption that the contemporaneous (differenced) demand and supply shocks are uncorrelated, $\bar{u}_i \rightarrow 0$ as $T \rightarrow \infty$ and equation (24) can be used to estimate γ by nonlinear least squares. I do so using grid search and imposing $\gamma > 0.5$, given that the results are sensitive to small values of γ . The results are presented in Appendix Table C3 and Appendix Figure C3.

C Robustness Checks on Quality-Adjusted Inflation across Income Groups

C.1 Summary of Results

Increasing product variety and manufacturer competition. Appendix Figures C1 and C2 show that the results are similar when defining a “new” product as a UPC code from a new manufacturer, which did not exist at all in previous years or was active in a different part of the product space. About 50% of the increase in product variety across product modules ranked by consumer income can be explained by the entry of manufacturers in a product module by price decile that is new to them. This directly addresses the concern that the patterns of increasing product variety would be directly due to spurious “relabeling” of UPC codes.

Valuing new products. Appendix Figure C3 shows that the result that higher-income households benefit more from the dynamics of product entry and exit is robust to valuing new products using the approximation of Hausman (2003) and using a translog demand system. In other words, this result is not a feature of the CES demand system.

Price data from point-of-sale dataset. A potential concern is that the price data in the Nielsen Homescan Consumer Panel is mismeasured. To address this issue, I have repeated the analysis using price data from the Nielsen Retail Scanner dataset, where prices are recorded at the point of purchase. Appendix Section C.4 shows that the results on inflation across income groups are similar using these prices.

Selection effects. A potential concern is that the inflation patterns for continued products across income groups could result from selection effects. For instance, it could be the case that low-income households overwhelmingly consume goods whose characteristics are rendered obsolete by the entry of new products. In such a case, a relatively higher share of the goods consumed by the poor would be exiting the market in any given year - the price changes for these goods are not observed, but if they were they would be negative because these products face tougher competition.⁶⁴ Appendix Tables C10 to C12 show that such selection effects are in fact not at play in the data.

The product cycle. One may worry that the patterns about inflation and new products are driven by the “product cycle” - namely, products start in the market with a very high price, and at that point are only purchased by high-income households, and then converge to their long-run, stable price, at which point they start being purchased by lower-income households. I address this concern in several ways. First, my results hold across the product space, as shown in Figures 2 and 3. If the product cycle was driving the results, then the measured differences in inflation and product innovation should only be visible from the point of view of each individual consumer and not across the product space. Second, I have repeated the analysis by considering only products in the middle of their lifecycle. Specifically, in any given year I have restricted the sample to products that had been in the market for at least two years and that would remain in the market for at least two more years. The inflation patterns obtained with this approach are similar to those reported

⁶⁴See Pakes and Erickson (2011) for a discussion of such selection effects.

above. Third, I have shown that the product cycle is not an important force in the data as barcodes do not travel down the income distribution (empirically, barcodes tend to remain in the same price decile during their entire lifecycle, which is intuitive for the retail sector and stands in contrast with other products like computers). Fourth, even if the product cycle were an important force in the data, under the assumption described at the beginning of this section the nested CES demand system would provide an accurate estimate of the quality-adjusted inflation rate for each of the various income groups, given the speed of the product cycle. In particular, in this analysis the “novelty” of a product is determined separately for each income group based on the basket of goods consumed by this income group in the previous year.

The fashion cycle. A distinct concern is that the inflation patterns may be driven by a phenomenon analogous to the “fashion cycle” - the fact that products exhibit seasonality patterns and that the price of older products falls disproportionately. For instance, because of the fashion cycle measured inflation is negative in the apparel industry - yet productivity gains for apparel are small and it would be incorrect to infer large welfare gains from the observed price patterns.⁶⁵ Conceptually, the fashion cycle means that the assumption that the “quality” of a barcode is fixed over time fails - if newness is a key feature of the utility derived from a product, the observed price of this product will fall over time but this may not reflect any change in the quality-adjusted price. I address the concern that high-income households are more affected by the fashion cycle in two ways. First, the fashion cycle is about churn of products and not about a net increase in the number of available varieties. I show that there is a faster increase in varieties in the parts of the product space that cater to higher income households, but there is not more churn. Similarly, the price patterns across product modules are predicted by the net increase in product variety, rather than by churn. Second, the results hold even with product categories where the fashion cycle is unlikely to exist, such as food products.

Price convergence Another potential concern is that the observed inflation difference between high- and low-income households could be driven by the fact that high-income households might initially pay a higher price for the same UPCs than low-income households, and the price would then converge to the same level for all households in future periods. The last three rows of Table 2 reject the hypothesis by showing that the “within-UPC” share of the total inflation difference is modest. A more direct way of showing that this mechanism is not the driving force, without the need for any assumption about the demand system, is to run a regression of the unit price of the UPC on a UPC fixed effect and an indicator for whether the household is high income (restricting attention to products purchased by both income groups). Appendix Table C13 reports the results of such a regression and shows that, in any given year, households making more than \$100,000 a year tend to pay about 2.9% more for the same UPC, compared with households making less than \$30,000 a year. This result is consistent with the findings of Pisano and Stella (2015). The magnitude of this effect is negligible compared with the 0.65pp difference in inflation rates, which over the course of a few years leads to a much bigger welfare difference between high- and low-income households than the difference

⁶⁵The Bureau of Labor Statistics addresses this by making hedonic adjustments and by ignoring sale prices.

in price levels in any given year.⁶⁶ Appendix Figure C6 provides complementary evidence by showing that the distribution of average unit prices paid by high- and low-income households is very similar, restricting attention to the set of products purchased by both income groups.

Alternative measures of household income. I repeated the analysis with three alternative measures of household income: reported income divided by household size; total retail expenditures per capita within a household; and whether the head of household is a college graduate. The results are similar.

Sampling variability. To ensure that the results are not driven by differing degrees of sampling variability across income groups, I built a random subsample of the data with an equal number of households in each of the income bins (following Handbury, 2013). I have also checked that the results across product modules hold in the Retail Scanner Data (which is based on information recorded directly at the store, not obtained from households, and contains many more observations, as described in Appendix A).

Extending the sample back to 1999 for food products. I have obtained Nielsen data on food products going back to 1999 (similar to Broda and Romalis, 2009). In ongoing work, I am repeating the previous analysis on this extended sample. Preliminary results show that inflation was also lower for higher-income households between 1999 and 2004.

Base drift. I have repeated the analysis in this section using unchained price indices instead of chained indices and obtained similar results.

Quarterly data. Appendix Table C9 shows that the results are very similar when repeating the analysis at a quarterly frequency.

C.2 Decomposition of Inflation Difference across UPC codes

Table C1 shows that differences in the spending patterns of high- and low-income households across price deciles within product modules explain more than 85% of the inflation difference between high- and low-income households that exists across UPCs. In other words, the decomposition shows that the inflation difference between high- and low-income households can be accounted for almost entirely by the fact that inflation is lower for higher-quality products (with higher unit prices), which primarily cater to higher-income consumers. Similar patterns exist when decomposing the share of spending on new products.

⁶⁶Note that my focus on inflation allows me to take into account changes in product variety and consumer substitution across products over time, as well as to characterize how these patterns differ across the income distribution. The static analysis of the levels of prices paid for the same barcodes by individuals across the income distribution does not speak to these dynamic considerations, which are first order in the data.

Table C1: Decomposition of the Inflation Difference Between High- and Low-Income Households Relative to Across-UPC Benchmark

Aggregation Level (Broad to Narrow)	Inflation Difference	
	pp	% of benchmark
Department	0.061	12.8
Product Group	0.143	30.0
Product Module	0.282	59.2
Product Module*Price Decile	0.408	85.7
UPC	0.476	100

C.3 Welfare Effects from Increasing Product Variety across Income Groups

Figure C1 shows that the patterns of increasing product variety across product modules are similar when measured by using the growth of the total count of UPC codes or the ‘Feenstra ratio’

$$\frac{\lambda_{mgt}}{\lambda_{mgt-1}} = \frac{1 + \text{Growth Rate of Spending on Overlapping Products}_{mgt}}{1 + \text{Growth Rate of Total Spending}_{mgt}}$$

which is the welfare-relevant metric in equation 2.

Figure C2 shows that manufacturer entry occurs disproportionately in parts of the product space catering to higher-income households. The introduction of new products in product modules by price deciles in which a manufacturer was previously never active explains over 50% of the patterns of increasing product variety across income groups.

Table C2 shows the distribution of ‘Feenstra ratios’ $\frac{\lambda_{mgt}}{\lambda_{mgt-1}}$ and of the elasticities of substitution in 1,075 product modules for households making above \$100,000 a year (‘high income’) and households making less than \$30,000 (‘low income’). A back-of-the-envelope calculation, based on the formula in equation 2 and using the mean Feenstra ratios and the mean elasticity, yields a difference in welfare gain between high- and low-income households in line with the results in Figure 2⁶⁷:

$$\Delta\pi = \frac{1}{6.2 - 1} ((1 - 0.9515) - (1 - 0.9448)) = 12.88bp$$

Table C3 shows the distribution of translog elasticities of substitution across product modules for households making above \$100,000 a year (‘high income’) and households making less than \$30,000 (‘low income’). Figure C3 shows that under the translog demand system, inflation patterns on continued products contribute to a 62 basis point lower rate of (annual) inflation for high-income household, relative to low-income households, and the dynamics of entry and exit contribute another 2 basis points. Intuitively, because elasticities of substitution are high, price changes on continued products capture most of the differential welfare effects.

⁶⁷This formula uses the approximation $\log(1 + x) \approx x$ for small x .

Figure C3 also reports the results when valuing new products using the “lower bound” formula of Hausman (2003), using the nested CES elasticity. The technique introduced by Hausman (2003) is a conservative way to value new products, using the slope of the demand curve at the observed prices and quantities: the use of a linear demand curve to estimate infra-marginal consumer surplus will provide a lower bound for the true infra-marginal consumer surplus (unless the true demand curve is concave to the origin, which is theoretically possible but is not expected for most products). The compensating variation under the linear demand curve is easily calculated as

$$\hat{\pi}_m^{Hausman} = \frac{0.5 \times (\tilde{\lambda}_{mt} - \tilde{\lambda}_{mt-1})}{\sigma_m}$$

where $\tilde{\lambda}_{mt}$ is the share of spending at time t on products that did not exist at time $t - 1$, and $\tilde{\lambda}_{mt-1}$ is the share of spending at time $t - 1$ on products that no longer exist at time t . The results show that high-income households benefit more from new products, but the magnitude of the effect is relatively small (3 basis points per year) due to high elasticities of substitution.

Table C2: Within-module CES Elasticities of Substitution and Feenstra ratios


Variable	Group	Mean	p25	p50	p75
Feenstra Ratios	High-Income	0.9448	0.9109	0.9566	0.9913
	Low-Income	0.9515	0.9229	0.9666	0.9942
CES Elasticities	High-Income	6.2680	3.9873	5.5027	7.5624
	Low-Income	6.3272	4.0974	5.7874	7.5196

Table C3: Within-module Elasticities of Substitution for Translog Demand System

Variable	Group	Mean	p25	p50	p75
Translog Semi-Elasticities	High-Income	54.93	0.9902	3.6522	150.5
	Low-Income	61.43	1.0231	4.6534	150.5

Figure C1: Increasing Product Variety across Income Groups

Panel A: Annual Growth in Total UPC Count across the Product Space



Panel B: Feenstra Ratio across the Product Space

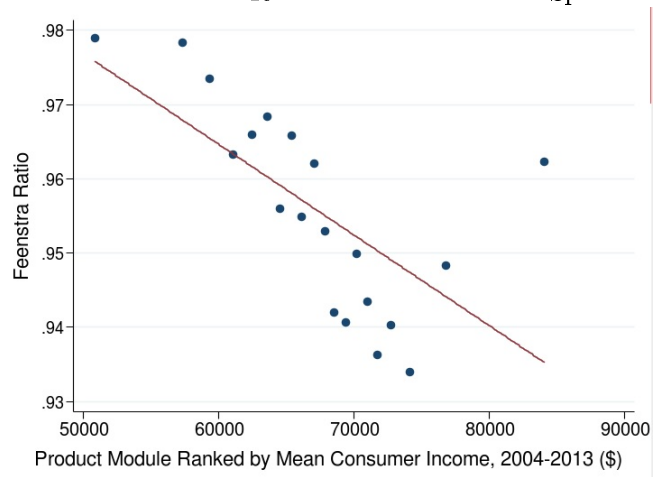
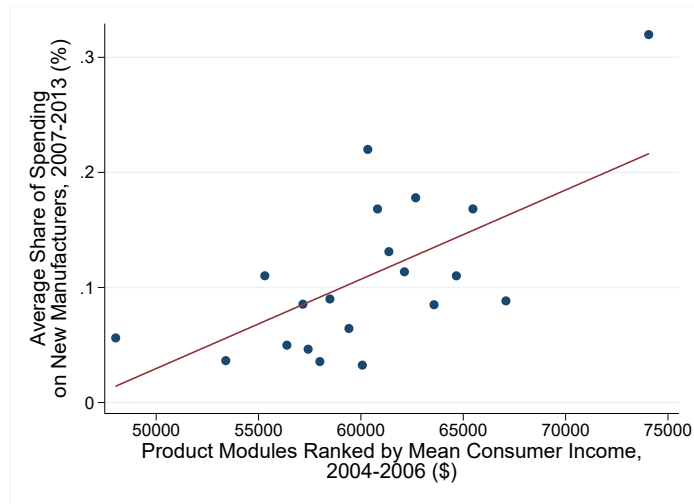


Figure C2: Manufacturer Entry Benefits Higher-Income Consumers More

Panel A: Manufacturer Entry



Panel B: Manufacturer Entry in a New Product Module - Price Decile

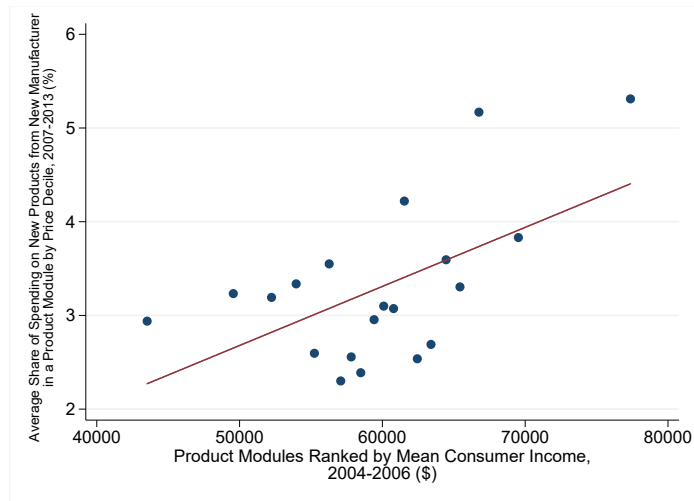
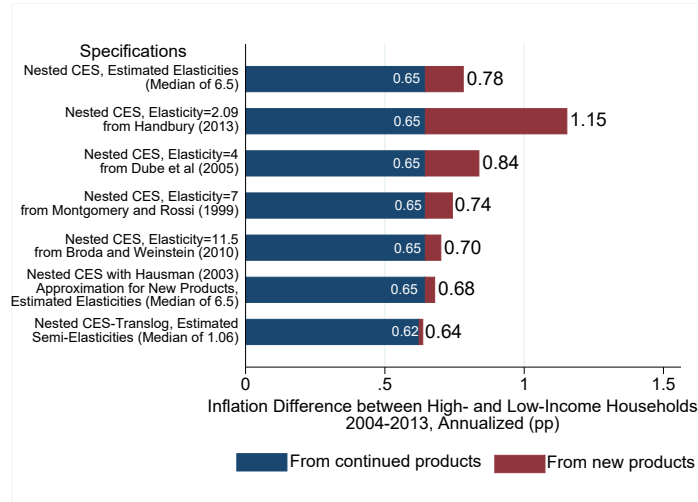


Figure C3: Robustness of Welfare Gains from New Products across Models and Parameters



C.4 Inflation Across Income Groups Using the Retail Scanner Dataset

The patterns of inflation on continued products across income groups are similar when the price information is obtained from the Retail Scanner dataset instead of the Homescan Consumer Panel dataset. Specifically, I merged the barcodes present in both datasets and computed price indices across income groups using the price information from the Retail Scanner dataset and the (income-group-specific) spending information from the Homescan Consumer Panel dataset.

The results are as follows: from 2006 to 2014, on average inflation was 47.57 basis points smaller for households making more than \$100,000 a year, relative to households earning below \$30,000 a year. This effect is consistent with the result obtained based on the Homescan Consumer Panel dataset alone, when income-group-specific inflation dynamics within UPC are ignored (47.6 basis points, reported in the last row of Table C1). Indeed, using the Retail Scanner data restricts the analysis to be carried across-UPCs, since for a given UPC prices are no longer allowed to vary across income groups.

Figure C4 and Table C4 show that (spending-weighted) average unit prices are very closely aligned in the Retail Scanner and Homescan Consumer Panel datasets. The data extends from 2006 to 2014, prices are winsorized at the 1% level and standard errors are clustered by product modules. Overall, this analysis confirms that the patterns of inflation across income groups reported in the main text of the paper do not depend on the choice of the dataset.

Figure C4: Relationship between Average Unit Prices in Retail Scanner and Consumer Panel Datasets, Graphical Analysis (2006-2014)

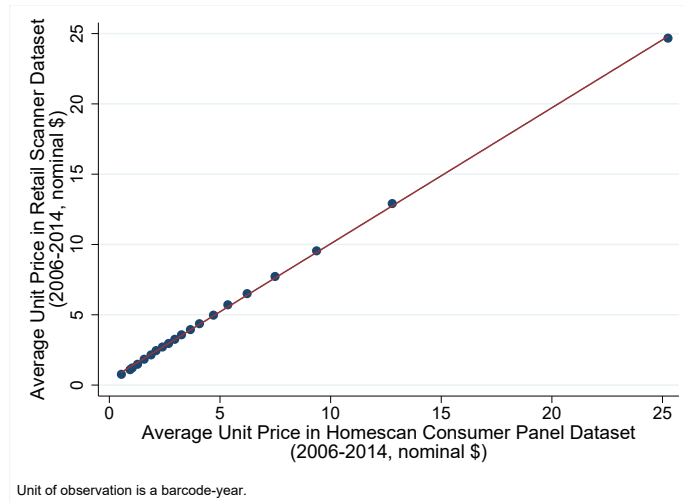


Table C4: Relationship between Average Unit Prices in Retail Scanner and Consumer Panel Datasets, Regression (2006-2014)

	Average Unit Price Retail Scanner Data
Average Unit Price Consumer Panel Data	0.9672*** (0.003626)
R^2	0.9149
N	3,117,983
# Clusters	1,019

C.5 Average Inflation Rates for Various Income Groups According to Various Price Indices

Table C5: Average Annual Inflation Rates Across Three Income Groups

Panel A: Full Sample (Percentage Points)

	Income < \$30k		Income ∈ [\$30k-\$100k]		Income > \$100k	
	Arithmetic Avg.	Geometric Avg.	Arithmetic Avg.	Geometric Avg.	Arithmetic Avg.	Geometric Avg.
Geometric Laspeyres	1.212	1.204	0.912	0.951	0.561	0.639
Truncated Geometric Laspeyres	1.544	1.536	1.137	1.157	0.862	0.909
Paasche	1.580	1.571	0.985	1.010	0.965	0.979
Truncated Paasche	1.719	1.710	1.182	1.194	1.117	1.126
Tornqvist	1.938	1.929	1.426	1.418	1.296	1.290
Fisher	1.983	1.974	1.425	1.418	1.327	1.320
Marshall-Edgeworth	1.992	1.984	1.440	1.433	1.330	1.323
CES Ideal	2.041	2.032	1.529	1.522	1.387	1.380
Truncated CES Ideal	2.063	2.054	1.541	1.534	1.413	1.406
Walsh	2.076	2.067	1.571	1.563	1.423	1.416
Truncated Laspeyres	2.257	2.502	1.724	1.910	1.554	1.721
Laspeyres	2.387	2.379	1.867	1.860	1.689	1.682
Truncated Geometric Paasche	2.433	2.424	1.742	1.734	1.822	1.815
Geometric Paasche	2.669	2.660	1.942	1.934	2.037	2.031

Panel B: All Years but Great Recession (Percentage Points, Arithmetic Average)

	Income < \$30k	Income ∈ [\$30k-\$100k]	Income > \$100k
Geometric Laspeyres	0.870	0.642	0.318
Truncated Geometric Laspeyres	1.179	0.876	0.627
Paasche	1.246	0.732	0.768
Truncated Paasche	1.380	0.928	0.919
Tornqvist	1.586	1.144	1.085
Fisher	1.625	1.161	1.111
Marshall-Edgeworth	1.633	1.176	1.116
CES Ideal	1.674	1.254	1.169
Truncated CES Ideal	1.695	1.268	1.192
Walsh	1.707	1.297	1.204
Truncated Laspeyres	1.891	1.448	1.316
Laspeyres	2.006	1.592	1.455
Truncated Geometric Paasche	2.071	1.467	1.623
Geometric Paasche	2.308	1.648	1.858

Table C5: Average Annual Inflation Rates Across Three Income Groups (*Continued*)

Panel C: Years Prior to Great Recession (Percentage Points, Arithmetic Average)

	Income < \$30k	Income ∈ [\$30k-\$100k]	Income > \$100k
Geometric Laspeyres	1.210	0.808	0.607
Truncated Geometric Laspeyres	1.545	1.060	0.938
Paasche	1.521	0.906	1.161
Truncated Paasche	1.670	1.186	1.276
Tornqvist	1.854	1.303	1.384
Fisher	1.884	1.281	1.428
Marshall-Edgeworth	1.892	1.294	1.414
CES Ideal	1.966	1.452	1.493
Truncated CES Ideal	2.019	1.493	1.532
Walsh	2.001	1.501	1.536
Truncated Laspeyres	2.249	1.658	1.695
Laspeyres	2.249	1.658	1.695
Truncated Geometric Paasche	2.317	1.688	1.971
Geometric Paasche	2.502	1.802	2.167

Panel D: Years After Great Recession (Percentage Points, Arithmetic Average)

	Income < \$30k	Income ∈ [\$30k-\$100k]	Income > \$100k
Geometric Laspeyres	0.615	0.519	0.101
Truncated Geometric Laspeyres	0.905	0.738	0.394
Paasche	1.03	0.601	0.473
Truncated Paasche	1.16	0.735	0.651
Tornqvist	1.386	1.024	0.860
Fisher	1.430	1.070	0.873
Marshall-Edgeworth	1.439	1.088	0.892
CES Ideal	1.456	1.106	0.926
Truncated CES Ideal	1.452	1.099	0.937
Walsh	1.485	1.143	0.954
Truncated Laspeyres	1.675	1.321	1.088
Laspeyres	1.823	1.542	1.275
Truncated Geometric Paasche	1.886	1.301	1.362
Geometric Paasche	2.162	1.532	1.625

Table C5: Average Annual Inflation Rates Across Three Income Groups (*Continued*)

Panel E: During the Great Recession (Percentage Points, Arithmetic Average)

	Income < \$30k	Income ∈ [\$30k-\$100k]	Income > \$100k
Geometric Laspeyres	2.408	1.857	1.411
Truncated Geometric Laspeyres	2.821	2.049	1.685
Paasche	2.751	1.872	1.657
Truncated Paasche	2.903	2.069	1.811
Tornqvist	3.168	2.413	2.036
Fisher	3.235	2.351	2.081
Marshall-Edgeworth	3.249	2.364	2.081
CES Ideal	3.323	2.492	2.147
Truncated CES Ideal	3.352	2.498	2.186
Walsh	3.369	2.531	2.189
Truncated Laspeyres	3.721	2.831	2.506
Laspeyres	3.721	2.831	2.506
Truncated Geometric Paasche	3.700	2.705	2.518
Geometric Paasche	3.933	2.973	2.666

Table C6: Average Annual Inflation Rates Across Four Income Groups

Panel A: Full Sample (Percentage Points)

	Income < \$25k	Income ∈ [\$25k-\$50k]	Income ∈ [\$50k-\$100k]	Income > \$100k
Geometric Laspeyres	1.236	1.029	0.785	0.561
Truncated Geometric Laspeyres	1.561	1.293	1.025	0.862
Paasche	1.647	1.249	0.962	0.965
Truncated Paasche	1.766	1.414	1.132	1.117
Tornqvist	2.000	1.668	1.365	1.296
Fisher	2.045	1.687	1.377	1.327
Marshall-Edgeworth	2.052	1.698	1.396	1.330
CES Ideal	2.086	1.763	1.462	1.387
Truncated CES Ideal	2.106	1.778	1.474	1.413
Walsh	2.116	1.800	1.501	1.423
Truncated Laspeyres	2.293	1.984	1.657	1.554
Laspeyres	2.445	2.126	1.795	1.689
Truncated Geometric Paasche	2.527	2.090	1.738	1.822
Geometric Paasche	2.769	2.311	1.949	2.037

Panel B: All Years but Great Recession (Percentage Points, Arithmetic Average)

	Income < \$25k	Income ∈ [\$25k-\$50k]	Income ∈ [\$50k-\$100k]	Income > \$100k
Geometric Laspeyres	0.843	0.729	0.529	0.318
Truncated Geometric Laspeyres	1.164	1.007	0.769	0.627
Paasche	1.289	0.964	0.723	0.768
Truncated Paasche	1.405	1.148	0.885	0.919
Tornqvist	1.613	1.374	1.097	1.085
Fisher	1.660	1.392	1.127	1.111
Marshall-Edgeworth	1.666	1.403	1.148	1.116
CES Ideal	1.692	1.469	1.201	1.169
Truncated CES Ideal	1.713	1.489	1.211	1.192
Walsh	1.722	1.504	1.241	1.204
Truncated Laspeyres	1.895	1.683	1.394	1.316
Laspeyres	2.033	1.823	1.533	1.455
Truncated Geometric Paasche	2.143	1.805	1.474	1.623
Geometric Paasche	2.388	2.024	1.669	1.858

Table C7: Average Annual Inflation Rates across the Income Groups at UPC*Geography Level (Full Sample, Percentage Points, Arithmetic Average)

	Income < \$30k	Income ∈ [\$30k-\$100k]	Income > \$100k
Paasche	2.065	1.401	1.341
CES Ideal	2.434	1.902	1.722
Laspeyres	2.789	2.365	2.08

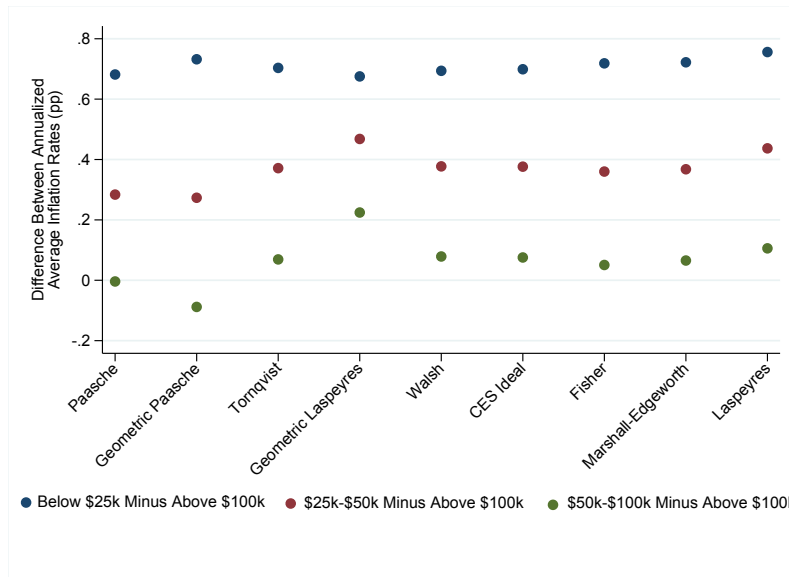
Table C8: Average Annual Inflation Rates across the Income Groups at UPC*Store Level (Full Sample, Percentage Points, Arithmetic Average)

	Income < \$30k	Income ∈ [\$30k-\$100k]	Income > \$100k
Paasche	2.239	2.002	1.692
CES Ideal	2.471	2.248	1.901
Laspeyres	2.710	2.471	2.072

Table C9: Average Annual Inflation Rates across the Income Groups at Quarterly Level (Full Sample, Percentage Points, Arithmetic Average)

	Income < \$30k	Income ∈ [\$30k-\$100k]	Income > \$100k
Paasche	-1.161	-2.268	-2.124
CES Ideal	1.911	1.107	1.066
Laspeyres	5.429	5.042	4.956

Figure C5: Inflation Difference Between Various Income Groups For Various Price Indices (Fixed Basket)



C.6 Is Differential Inflation on Continued Products Across Income Groups Driven by a Selection Effect?

Table C10: Products that are about to exit have a lower inflation rate

Subsample	Laspeyres Inflation Rate	Median Laspeyres Inflation Rate (Across Product Modules)
Continued	2.03%	2.06%
About to Exit	-1.33%	-0.52%
Justed Entered	0.03%	1.3%

Table C11: Products that are about to exit have a higher price level

Subsample	Average Price Level	Median Price Level (Across Product Modules)
Continued	3.67	2.75
About to Exit	3.95	2.68
Justed Entered	4.91	3.05

Table C12: Share of spending on new and discontinued products across income groups

Household Income	Share of Spending on Products...	
	About to Exit	Just Entered
> \$100,000	3.04%	10.94%
\$30,000 – \$100,000	2.71%	10.01%
< \$30,000	2.59%	9.26%

C.7 Differences in Prices Paid for Same Products for High- and Low-Income Households

Figure C6: The Distribution of Average Unit Prices Paid is the Same Across Income Groups (Reweighting by Spending Shares)

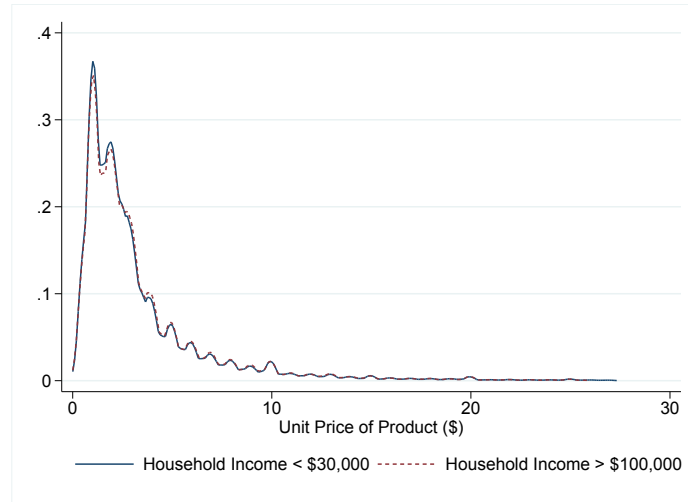


Table C13: Differences in Price Level Paid for Same UPC by High- and Low-Income Households (\$)

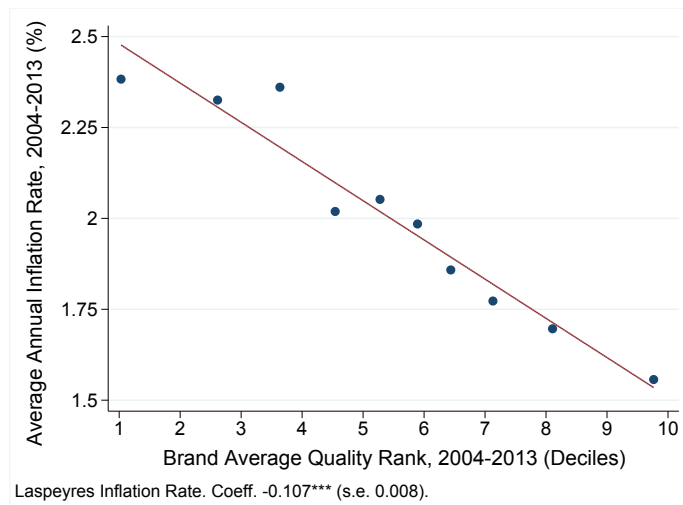
	Average Unit Price
High-Income Household	0.0664*** (0.00118)
Constant	2.2825*** (0.00061)
UPC*Year Fixed Effect	Yes
R^2	0.9954

D Robustness Checks on the Equilibrium Response of Supply to Changes in Demand

D.1 Descriptive Evidence

Inflation across brand price deciles. Figure D1 shows inflation patterns across UPCs ranked by the average (leave-one-out) price per ounce of UPCs of the same brand in the same product module available during the sample period. The deciles are not based on the price of the UPC itself and the results are identical to Panel B of Figure 3, which confirms that mean reversion is not driving these patterns.

Figure D1: Inflation across Brand Price Deciles, within Product Modules



The role of supply effects. Do the product variety patterns across income groups come from supply or demand? As shown on Figure 2, the share of spending on new products increases with mean consumer income across product modules. It could be the case that more new products are introduced in product modules catering to high-income consumers because of supply effects, which may be exogenous (e.g. it may be inherently easier to introduce new products at the high-end of the product space) or endogenous (e.g. if innovators and suppliers decide to specifically target higher-income consumers). Alternatively, it could be the case that higher-income consumers have a higher taste for novelty and purchase new products wherever they are introduced in the product space. In other words, the share of spending on new products may be higher in product modules catering to higher-income households simply because new products diffuse faster due to a basic composition effect in demand (while the rate of product introduction may be similar across modules).

To isolate the contribution of supply, the ideal regression would compare the same household moving across the product space. Such a regression can be directly run in the Nielsen data, at the household $H \times$

product module M level with household fixed effects:

$$\begin{aligned} & \text{ShareSpendingNewProducts}_{HM} \\ &= \alpha + \beta \text{ProductModuleIncomeRank}_M + \alpha_H + \epsilon_{HM} \end{aligned}$$

where α_H is a household fixed effect and $\text{ProductModuleIncomeRank}_M$ is the rank of the product module by income of the representative consumer in the product module (computed using 2004-2006 data). The results are reported in Table D1, with standard errors clustered at the household level. As in the previous graphs, I find a strong positive relationship between the share of spending on new products and the mean income of the consumer in the product module - the point estimate is almost identical to the specification without household fixed effect shown in Figure 2. This analysis confirms that supply plays a role in this process, because household fixed effects ensure that the relationship is not driven by a composition effect across modules (i.e. different propensities of consumers to buy new products wherever they show up in the product space). I also present specifications with interaction terms for whether the household is “high-income” (income above \$100,000) or “low-income” (income below \$30,000). The magnitude of the interaction effects is small, around 10% of the effect for middle-income households.

Table D1: New Products Target Higher-Income Consumers

	<i>ShareSpendingNewProducts_{hm}</i>	
<i>ProductModuleIncomeRank_M</i>	2.79*** (1.024)	2.82*** (1.031)
<i>ProductModuleIncomeRank_M × HighIncome_H</i>		-0.24*** (0.063)
<i>ProductModuleIncomeRank_M × LowIncome_H</i>		0.11* (0.058)
<i>Household Fixed Effects</i>	Yes	Yes

Standard errors clustered by product modules.

Retailers vs. Manufacturers Decomposition. In order to establish whether the supply effects documented above are driven by retailers or manufacturers, I carry out an additional decomposition of the inflation difference between high- and low-income households. For this exercise I use the Laspeyres price index, which can be written as follows:

$$P_L^i \equiv \sum_{i=1}^n \frac{p_i^t}{p_i^0} s_i^0 = \sum_{i=1}^n \frac{p_i^t}{p_i^0} s_{local\ market}^i \cdot s_{store}^i \cdot s_{upc}^i$$

where i indexes the income group, $s_{local\ market}^i$ the share of spending in a given local market (MSA), s_{store}^i the share of spending in a given store within a local market, and s_{upc}^i the share of spending on a given UPC within a store. In other words, the difference in inflation rates between high- and low-income households

across UPCs could come from the fact that these consumers shop in different local markets or different stores or buy different UPCs within stores.

Table D2 presents the results. The third row shows that differences in spending patterns across local markets (MSAs) explain only about 3% of the inflation difference across UPCs between high- and low-income households. The second row gives an upper bound for the contribution of store-specific price dynamics, which account for at most about 40% of the total difference. It is an upper bound because in several stores I only observe spending from either the low- or high- income, therefore I cannot separately identify the contribution of UPC dynamics within stores. Overall, these results show that at least 60% of the inflation difference comes from UPC effects within stores, suggesting that manufacturer-level dynamics are a key channel.

Table D2: Isolating the Contribution of Stores and Local Markets to the Overall Inflation Difference between High- and Low-Income Households

Price Change	Local Market Shares	Store Shares	UPC Shares	Inflation Difference (% of Benchmark)
Counterfactual	Actual	Actual	Actual	100
Counterfactual	Counterfactual	Actual	Counterfactual	43.2
Counterfactual	Actual	Counterfactual	Counterfactual	3.1

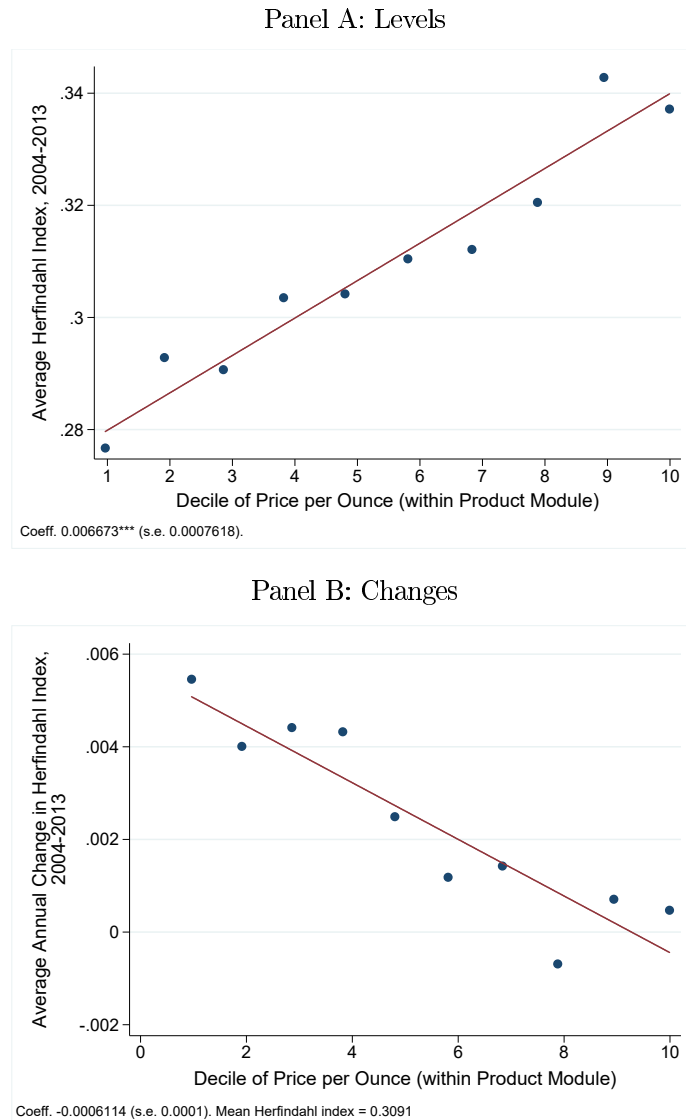
Evidence on Herfindahl Indices Across the Quality Ladder. Herfindahl indices vary substantially across the product space. Table D3 presents the distribution of Herfindahl indices across product modules by price deciles. The level of observation is a year by product module by price decile. Statistics are weighted by log spending and show wide dispersion in Herfindahl indices.

Figure D2 shows that Herfindahl indices across the product space, in levels and changes, are systematically related to the quality ladder. Panel A shows that, from a static perspective, on average competition tends to be lower in higher-quality tiers of the market, where quality is proxied for by prices deciles within modules as in Section 4.1. Panel B indicates that over time competition increases in higher-quality tiers, relative to lower-quality tiers. The magnitude of these differential changes in competition across the quality ladder is large. Over a period of ten years, on average the Herfindahl indices of the top and bottom deciles within a product module converge by 0.061 points, which is equal to 27% of the standard deviation of Herfindahl indices across the product space. This evidence supports the prediction of the model in Section 5 that increases in market size in higher-quality tiers spur entry and increasing competition.

Table D3: Summary Statistics on Herfindahl Indices Across Product Modules by Price Deciles


Variable	Mean	SD	p25	p50	p75
Herfindahl Index	0.3091	0.2282	0.1382	0.2410	0.4112

Figure D2: Herfindahl Indices in Levels and Changes across the Quality Ladder



Evidence from Variation in the Rate of Growth of Inequality Across US States. Using Census public use microdata between 2004-2006 and 2012-2014, I measure the change in the total income accruing to households who earned more than 100k and less than 30k in each state. Inequality has increased in all 50 states but the rate of increase varied across states. The increase in inequality was fastest in California, Texas and New York and slowest in West Virginia, New Mexico and North Dakota. I aggregate the Nielsen data at the state level to examine how variation in the rate of inequality growth relates to patterns of inflation. In all states, inflation was lower for high-income households earning above \$100,000 a year, relative to low-income households making below \$30,000 a year. But this difference in inflation rates was relatively larger in states with a faster increase in inequality. Figure D3 shows this result.

Figure D3: The Inflation Difference Between High- and Low-Income Increases as Inequality Increases Faster

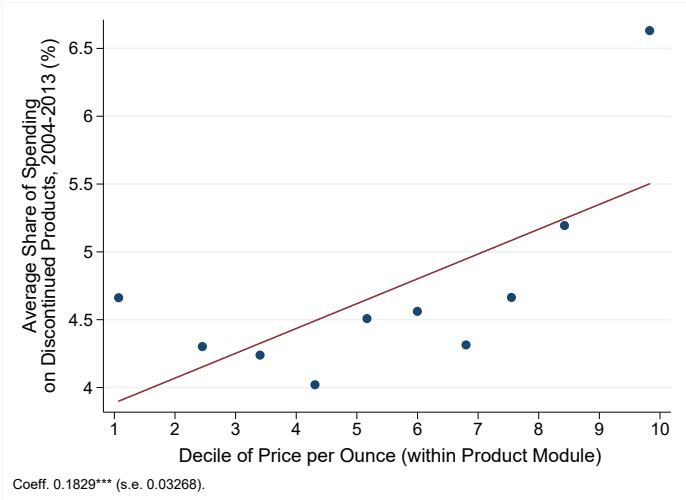


Evidence on Product Destruction and Net Product Creation Across the Quality Ladder.

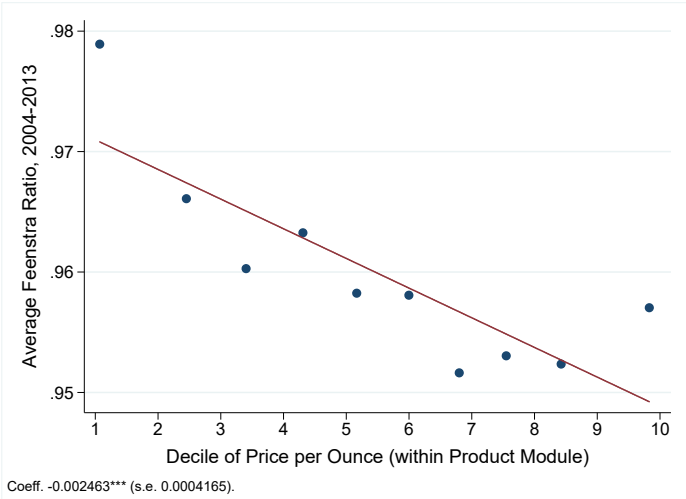
Panel A of Figure D4 shows that there is relatively more exit of products at the top of the quality ladder. This differential exit effect is smaller than the differential entry patterns documented in Figure 3. As a result, product variety increases faster in higher-quality tiers of the market. Panel B of Figure D4 quantifies the differential increase in product variety across the quality by plotting the “Feenstra ratio”, derived using a nested CES demand system in Section 3.3.

Figure D4: Product Destruction and Net Product Creation Across the Quality Ladder

Panel A: Spending on Discontinued Product



Panel B: Feenstra Ratio



D.2 Causal Evidence

D.2.1 Source of variation

Figure D5: Changes in Age-by-Income Distributions, 2011-2015 relative to 2000-2004

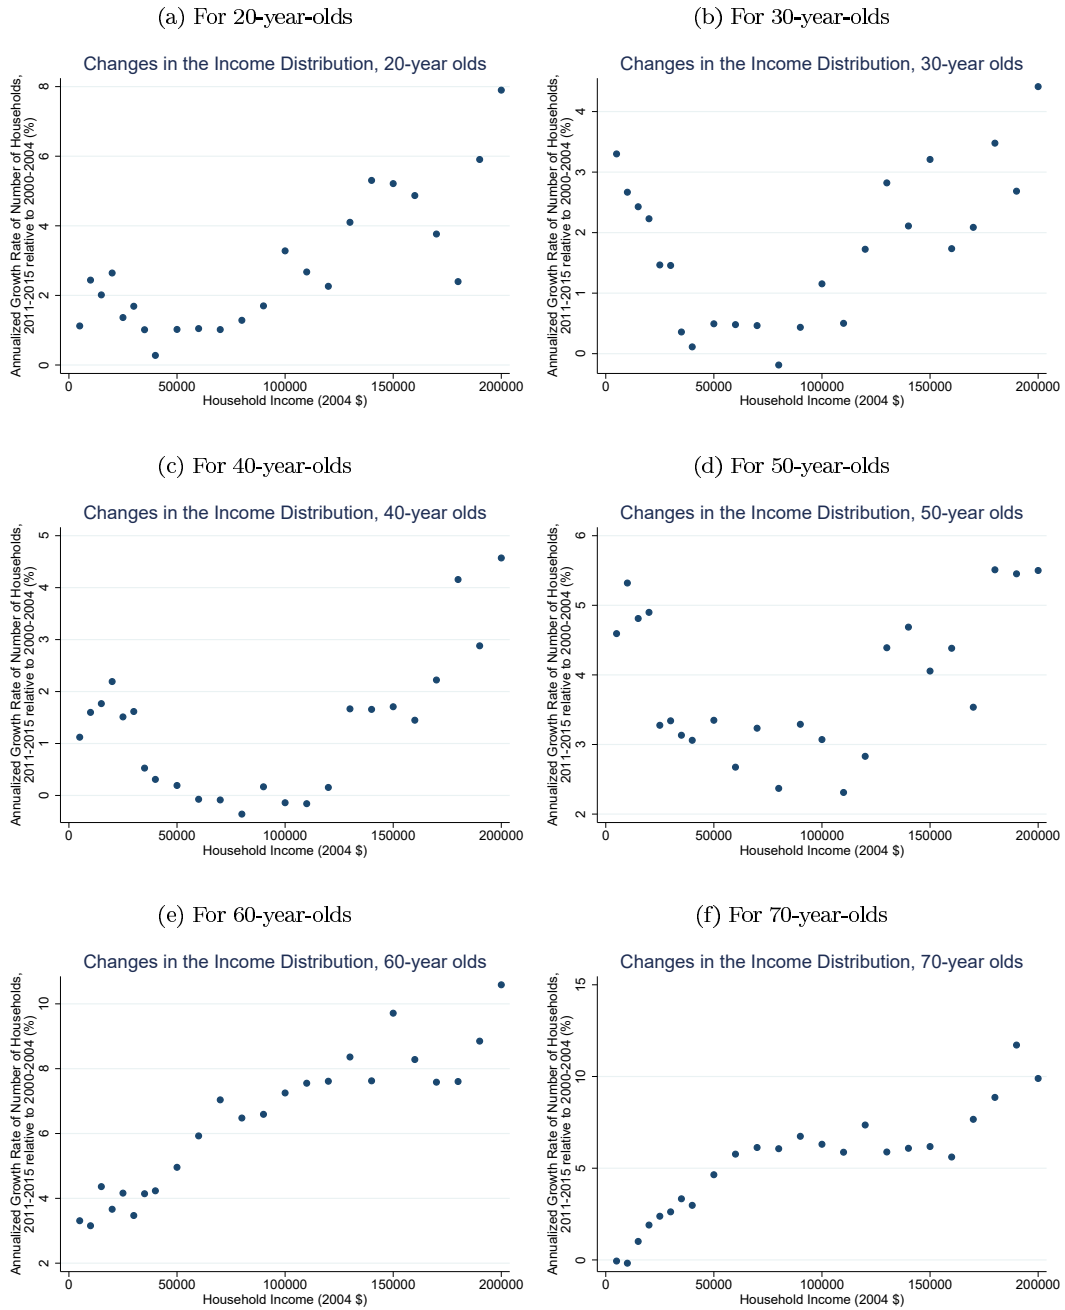


Table D4: Distribution of Changes in Number of Households across Age-Income Groups

	Annualized Change (%), 2011-2015 relative to 2000-2004		
Mean	1.23	p10	-1.62
SD	2.24	p25	0.41
Min	-2.04	p50	0.81
Max	7.56	p75	2.72
N	108	p90	4.50

Figure D6: Changes in Number of Households by Age-Income Groups, 2011-2015 relative to 2000-2004

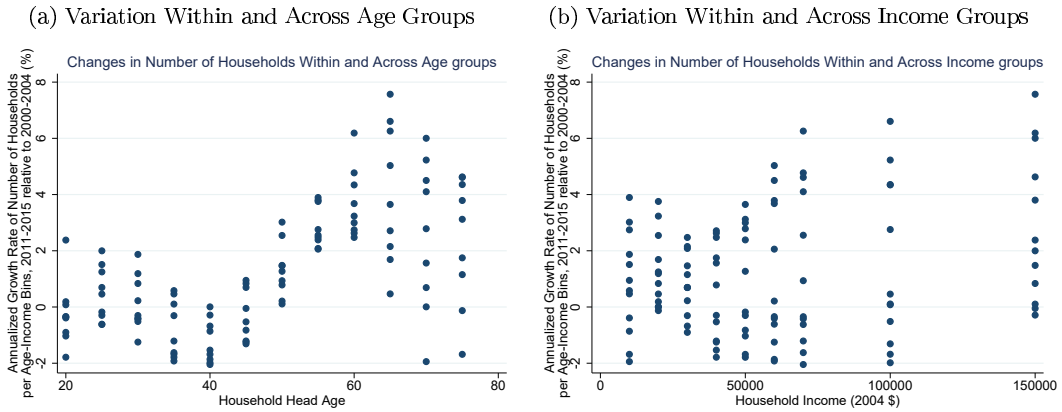


Table D5: Distribution of Predicted Growth Across the Product Space

	Annualized Change (%), 2011-2015 relative to 2000-2004		
Mean	0.90	p10	0.36
SD	0.49	p25	0.62
Min	-1.94	p50	0.88
Max	5.04	p75	1.13
N	10,394	p90	1.45

D.2.2 Identification and Consistency of OLS Estimator⁶⁸

Assumptions on data generating process. We assume that $\epsilon_l = f(s_{l1}, \dots, s_{lN}) + \eta_l$ where $f(\cdot) \in [-B, B]$ for some $0 < B < \infty$.⁶⁹ We consider a sequence of statistical models. In each of them, g_n , s_{ln} and η_l are drawn for all groups indexed by n and parts of the product space indexed by l — these random variables are assumed to be jointly independent.

⁶⁸For more details on identification and consistency in Bartik research designs, see Goldsmith-Pinkham et al. (2016) and Borusyak and Jaravel (2016).

⁶⁹ $f(\cdot)$ is a random function. This notation illustrates that identification can be preserved even when the error term is correlated with expenditure shares.

Consistency of estimator. Denote the average Z_l across the product space as:

$$\bar{Z} = \frac{1}{L} \sum_l Z_l = \frac{1}{L} \sum_l \sum_n s_{ln} g_n = \sum_n g_n \frac{1}{L} \sum_l s_{ln} = \sum_n \bar{s}_n g_n \quad (25)$$

where $\bar{s}_n = \frac{1}{L} \sum_l s_{ln}$ measures the importance of household group n in an average product category. The OLS estimator for β is:

$$\hat{\beta} = \frac{\frac{1}{L} \sum_l \left[\left(\overbrace{\sum_n (s_{ln} - \bar{s}_n) g_n}^{=Z_l - \bar{Z}} \right) \cdot Y_l \right]}{\frac{1}{L} \sum_l \left[\sum_n (s_{ln} - \bar{s}_n) g_n \right]^2} = \beta + \frac{\frac{1}{L} \left[\sum_l \left(\sum_n (s_{ln} - \bar{s}_n) g_n \right) \epsilon_l \right]}{\frac{1}{L} \sum_l \left(\sum_n (s_{ln} - \bar{s}_n) g_n \right)^2} \quad (26)$$

Conditions under which the second term goes to 0 as N and L go to infinity are provided below.

First, the denominator should not go to zero. It could go to zero if there were many household groups and very dispersed household shares in each product category. Sufficient concentration of spending across household groups in a typical product category is required. For instance, a group of high-income household might account for most of the demand in a product category like scotch. This holds, for instance, when $s_{ln} = \frac{X_{ln}}{\sum_m X_{lm}}$ with X_{ln} following a Pareto distribution, or when each product category is perfectly specialized in catering to one household group, $s_{ln} = 1 [n = n_l]$.

Second, we are looking for the conditions under which the numerator goes to 0. Intuitively, we flip the order of the summation for the $f(\cdot)$ component of the error term, so that we can make an asymptotic statement in household group space N . The expression for the numerator can be re-written as follows:

$$Num = \frac{1}{N} \sum_n \left[g_n \cdot \frac{1}{L} \sum_l N(s_{ln} - \bar{s}_n) f(s_{ln}) \right] + \frac{1}{L} \sum_l (Z_l - \bar{Z}) \eta_l \quad (27)$$

Intuitively, we are averaging over household group objects, $g_n \cdot \sum_l \frac{N}{L} (s_{ln} - \bar{s}_n) f(s_{ln})$. A loose intuition is that for the numerator to go to 0, there should be no systematic relationship between the growth rate of a household group g_n and the cross-product category covariance between the spending share of this group (relative to the full sample) and the error term (induced by household group composition). For instance, it should not be the case that household groups that grow faster have higher spending shares in product categories with a larger error term (due to anything related to changing household share composition across product categories).

This intuition is imperfect, however, because the relevant size of the household group in the expression above is $N s_{ln}$, which may explode without further assumptions. Regularity conditions on $\{s_{ln}\}$ ensure this does not happen. For instance, one can consider a special case where each product category l is fully specialized in catering to one household group n_l , i.e. $s_{ln} = 1 [n = n_l]$. The standard \sqrt{N} convergence rate is achieved when all household groups are approximately of the same size—that is, $\exists a \in (1, \infty)$ such that $\bar{s}_n < a/N$ for all n .

Controls. In some cases, it may be the case that the assumption that g_n is randomly assigned holds only conditional on some controls. For instance, each household group n is characterized by a variable x_n taking

values $a \in [1, 2, \dots, A]$. x_n is not independent from the spending shares s_{ln} , which are themselves correlated with the error term ϵ_l . Consider the following model:

$$g_n = \tilde{g}_n + \sum_a \mu_a \cdot 1(x_n = a)$$

where μ_a is a growth fixed effect common to all household groups for which $x_n = a$ and \tilde{g}_n is the residual variation in growth rate (assumed to be independent of x_n). So $\sum_n s_{nl} g_n = \sum_n s_{nl} \tilde{g}_n + \sum_n s_{nl} (\sum_a \mu_a \cdot 1(x_n = a))$. Note that $\sum_n s_{nl} (\sum_a \mu_a \cdot 1(x_n = a)) = \sum_a \mu_a \underbrace{\left(\sum_n s_{nl} \cdot 1(x_n = a) \right)}_{\equiv s_{al}}$, where s_{al} denotes the share of spending accounted for by sectors with $x_n = a$ in product category l .

Consider the following regression:

$$\begin{aligned} Y_l &= \alpha + \beta \left(\sum_n s_{nl} \tilde{g}_n \right) + \beta \left(\sum_a \mu_a s_{al} \right) + \epsilon_l \\ &= \alpha + \beta \left(\sum_n s_{nl} \tilde{g}_n \right) + \left(\sum_a \tilde{\mu}_a s_{al} \right) + \epsilon_l \end{aligned} \quad (28)$$

Equation (11) is an OLS regression with parameters β and $\tilde{\mu}_a$. Compared with the baseline model, we are adding controls for the weighted distribution of variable x_n in location l , where the weights are local spending shares. Intuitively, we are considering a ‘‘fixed effects’’ regression in sector space N , and when we move to the location space L to run the regression we obtain a linear estimator in the spending shares accruing to sectors with characteristics indexed by a .

A consistent estimator of β is given by running the equation above using g_n instead of \tilde{g}_n :

$$Y_l = \alpha + \beta \left(\sum_n s_{nl} g_n \right) + \left(\sum_a \tilde{\mu}_a s_{al} \right) + \epsilon_l \quad (29)$$

Intuitively, β is identified from the residual variation in g_n , after flexibly controlling for x_n with fixed effects. Formally, using the residual regression formula, the OLS estimator for β in (29) is given by:

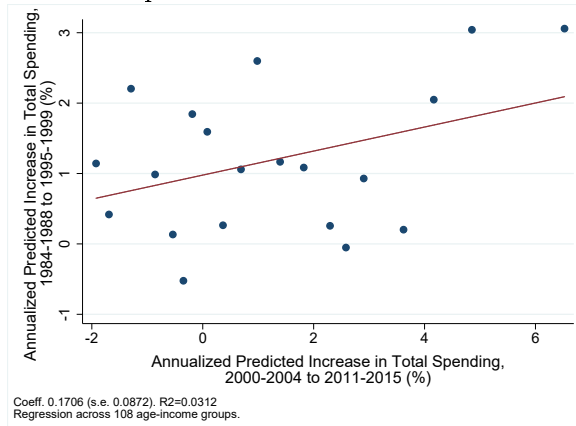
$$\hat{\beta} = \frac{\frac{1}{L} \sum_l \left[\overbrace{\left(\sum_n (s_{ln} - \bar{s}_n) \tilde{g}_n \right)}^{=Z_l - \hat{Z}_l} \cdot (Y_l - \hat{Y}_l) \right]}{\frac{1}{L} \sum_l \left[\sum_n (s_{ln} - \bar{s}_n) \tilde{g}_n \right]^2} = \beta + \frac{\frac{1}{L} \left[\sum_l \left(\sum_n (s_{ln} - \bar{s}_n) \tilde{g}_n \right) \epsilon_l \right]}{\frac{1}{L} \sum_l \left(\sum_n (s_{ln} - \bar{s}_n) \tilde{g}_n \right)^2} \quad (30)$$

where \hat{Z}_l and \hat{Y}_l are the best linear predictors of Z_l and Y_l , respectively, conditional on a constant and $\{s_{al}\}_{a=1}^A$. The second equality follows from the fact that $Y_l - \hat{Y}_l = \beta (\sum_n (s_{ln} - \bar{s}_n) \tilde{g}_n)$. This estimator of β with controls is consistent under conditions similar to those derived above, except that the conditions are now in terms of \tilde{g}_n instead of g_n .


D.2.3 Robustness

Figure D7: Falsification Tests

Panel A: Relationship between Actual and “Placebo” Growth Rates



Panel B: Falsification Test for New Products



Panel C: Placebo Test for Inflation for Continued Products




Table D6: Robustness Checks on Effects of Changes in Market Size

Panel A: Additional Outcomes

	Actual Spending Growth (%)	Share of Spending on Discontinued Products (pp)	Feenstra Ratio
Predicted Increase in Spending, Annualized (%)	1.031** (0.492)	1.3963*** (0.3377)	-0.00858*** (0.00301)
Age and Income Controls	Yes	Yes	Yes
Product Module Fixed Effects	Yes	Yes	Yes
R^2	0.48	0.48	0.40
Number of Observations	10,705	10,705	10,705
Number of Clusters	1,075	1,075	1,075

Panel B: Interaction By Herfindahl Index

	Share of Spending on New Products (pp)	Continued Products Inflation Rate (pp)
Annualized Predicted Increase in Spending × Herfindahl Index	-0.5434*** (0.1941)	-0.02396 (0.045003)
Age, Income and Herfindahl Controls	Yes	Yes
Product Module Fixed Effects	Yes	Yes
R^2	0.55	0.54
Number of Observations	10,705	10,705
Number of Clusters	1,075	1,075


Panel C: Robustness without Spending Weights

	Share of Spending on New Products (pp)	Continued Products Inflation Rate (pp)
Predicted Increase in Spending, Annualized (%)	2.4450*** (0.5027)	-0.4121*** (0.1171)
Age and Income Controls	Yes	Yes
Product Module Fixed Effects	Yes	Yes
R^2	0.53	0.50
Number of Observations	10,705	10,705
Number of Clusters	1,075	1,075


D.3 The Supply Response to Market Size Effects Implied by Changes in the Income Distribution

Figure D8: Growth of Number of Households across the Income Distribution

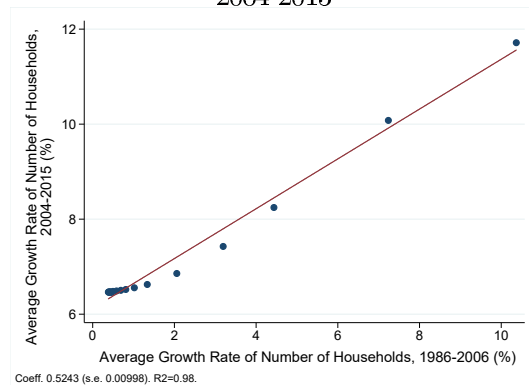
Panel A: Smoothed Growth Rates of Number of Households Across the Income Distribution (1986-2006)



Panel B: Raw Growth Rates of Number of Households Across the Income Distribution (1986-2006)



Panel C: Stability of Growth Rates of Number of Households Across the Income Distribution: 1986-2006 vs. 2004-2015



Notes: The data source is the Annual Social and Economic Supplement of the Current Population Survey. Panel A shows smoothed patterns of growth in the number of households across the income distribution, using a quartic polynomial with parameters estimated by OLS.

Table D7: Actual and Predicted Relationship between Mean Consumer Income, New Products and Inflation for Continued Products across Product Modules by Price Deciles

Panel A: With Product Module Fixed Effects


	Share of Spending on New Products (pp)		Continued Products Inflation Rate (pp)	
	Predicted	Actual	Predicted	Actual
Mean Consumer Income (per \$10,000)	1.0340*** (0.00846)	1.2364*** (0.1235)	-0.15938*** (0.000682)	-0.19128*** (0.028869)
Ratio of Slopes (Predicted/Actual)	83.63%		83.32%	
Product Module Fixed Effects	Yes	Yes	Yes	Yes
R^2	0.98	0.58	0.98	0.54
Number of Observations	10,750	10,750	10,750	10,750
Number of Clusters	1,075	1,075	1,075	1,075

Panel B: Without Product Module Fixed Effects

	Share of Spending on New Products (pp)		Continued Products Inflation Rate (pp)	
	Predicted	Actual	Predicted	Actual
Mean Consumer Income (per \$10,000)	1.0100*** (0.00714)	0.9650*** (0.1306)	-0.16153*** (0.001298)	-0.23138*** (0.0306006)
Ratio of Slopes (Predicted/Actual)	104.66%		69.81%	
Product Module Fixed Effects	No	No	No	No
R^2	0.97	0.01	0.97	0.01
Number of Observations	10,750	10,750	10,750	10,750
Number of Clusters	1,075	1,075	1,075	1,075

D.3.1 Dissimilarity Index

Figure D9: Dissimilarity Index



D.4 Additional Evidence


D.4.1 Change vs. Level of Market Size

Table D8: Do Product Innovations Follow Market Size or Change in Market Size?

	Share of Spending on New Products	
Lagged Change in Market Size	3.107*** (1.139)	1.901** (0.926)
Lagged Market Size	1.399 (1.439)	0.577 (1.269)
Product Group Fixed Effects	No	Yes
Weights	Yes	Yes

D.4.2 Changes in Markups

Figure D10: Changes in Wholesale Costs vs. Changes in Retailer Margins



D.5 Food Stamp Research Design


Using changes in food stamp policy across US states between 2000 and 2007, I estimate the causal effect of an increase in the level of spending per capita in a certain part of the product space on the introduction of new products and the rate of inflation (holding the number of consumers constant). I find a substantial effect.

D.5.1 Research Design

I rely on a novel research design based on changes in food stamp policy across US states between 2000 and 2007, which generates variation in per capita spending on food from low-income consumers. Between 2001 and 2007, the take-up rate for food stamps dramatically increased due to a series of policy changes that made

it easier for eligible individuals to enroll in the program. Ganong and Liebman (2016) document these trends, reproduced in Figure D11. They also document that the increase in take-up rate substantially varied across states, because different states adopted a different policy mix.⁷⁰ This policy variation generates variation in purchasing power for food products at the bottom of the income distribution and is plausibly exogenous to price dynamics. This addresses the endogeneity problem that better products get larger market shares and allows me to estimate the causal effect of an increase in per capita spending in a certain part of the product space on the inflation rate.

Figure D11: Changes in SNAP Take-up Rate and Total Enrollment over Time



Source: Ganong and Liebman (2015)

This identification strategy is a useful complement to the previous analysis based on changes in the number of consumers across the product space at the national level over time. First, it is interesting to examine whether variation in demand coming from changes in per capita spending generates similar effects to variation in demand coming from changes in the number of consumers. Second, the SNAP-based research design has a number of advantages from the point of view of identification: there is clearly no direct supply effect, the market size change occurs at the bottom of the distribution (thus breaking the usual collinearity between level of income and rate of growth in income), and the time frame and the location of the market size change are known very precisely. Third, these findings are of direct policy relevance (for a study of the short-run incidence effect of food stamp policy, see Hastings and Washington, 2010).

Thus, the research design is based on variation in changes in take-up rates across US states. I compare the difference between the inflation rates experienced by SNAP eligible and ineligible households between

⁷⁰For instance some states stopped requiring fingerprints from food stamp recipients, which facilitated the application process. Other states amended their vehicle policies, for instance excluding the value of all vehicles when determining eligibility for the program.

2004 and 2007 across states, running the following specification:

$$\pi_S^E - \pi_S^I = \alpha + \beta \Delta \tau_s^{SNAP} + \lambda X_s + \epsilon_s$$

Variation in the SNAP take-up rate induces variation in market size for manufacturers with *local* brand capital. Many UPCs are partly non-tradable because of the strength of local brand preferences (Bronnenberg, Dube and Gentzkow, 2012). The strength of local brand preferences varies across product groups. This provides an opportunity for a falsification test of the research design: inflation should respond to local changes in market size only in product groups for which brand preferences tend to be “local.”

I set up a random effect model to identify in a data-driven way which product groups have strong brand preferences. Intuitively, local preferences must be strong for product groups in which I observe a lot of variation in the ranking of brands by market shares across different states. On the other hand, local preferences must be weak in product groups where the market shares of brands are very similar across states. The random effect model provides a way to conduct this comparison systematically and to handle noise efficiently. Formally, for each product group I write the market share of brand b in state s at time t as the sum of a “national preference” component λ_b , a “local preference” component μ_{bs} and a shock ϵ_{bst} . I then estimate the signal standard deviation of the “national preference” component, denoted $\hat{\sigma}_\lambda^2$, and the signal standard deviation of the “local preference” component $\hat{\sigma}_\mu^2$.⁷¹

$$\begin{aligned} s_{bst} &= \lambda_b + \mu_{bs} + \epsilon_{bst} \\ \hat{\sigma}_\epsilon^2 &= \text{Var}(s_{bst} - \bar{s}_{bs}) \\ \hat{\sigma}_\lambda^2 &= \text{cov}(\bar{s}_{bs}, \bar{s}_{b(s+1)}) \\ \hat{\sigma}_\mu^2 &= \text{Var}(s_{bst}) - \hat{\sigma}_\lambda^2 - \hat{\sigma}_\epsilon^2 \end{aligned}$$

Finally, I rank product groups according to the quantity $R = \frac{\hat{\sigma}_\mu^2}{\hat{\sigma}_\lambda^2}$. The product groups above median R are those where local preferences matter relatively more. The results I obtain from this procedure are very intuitive: sanitary protection, canning supplies, detergent, flour and deodorant are the five product groups for which local preferences are the weakest, while liquor, wine, beer, apparel and fresh meat are the five product groups with the strongest local preference component. I conduct the regression analysis across subsamples to check that the effect is driven by product group with a strong local brand component.

D.5.2 Results

I find a large effect, which can be summarized as follows: a 1 percentage point increase in spending per capita lowers the inflation rate by about 10 basis points. Consistent with my preferred model, the magnitude of this effect is similar to that of the effect of a change in the number of consumers documented in the previous subsection.

⁷¹This approach is similar to the model used in the teacher value-added literature, for instance in Kane and Staiger (2008).

Table D9: Results from SNAP Research Design

Panel A: Main Results				
	Actual Spending Growth for SNAP Eligible (%)		Difference in Continued Products Inflation Rate (pp)	
Change in Take-up Rate (pp), 2001-2007	0.2226*** (0.0770)		-0.0242*** (0.00791)	
Controls	Yes		Yes	
Weights	Yes		Yes	
<i>Standard errors clustered by 50 states</i>				
Panel B: Robustness				
	Difference in Continued Products Inflation Rate (pp)			
Change in Take-up Rate (pp), 2001-2007	-0.0242*** (0.00791)	-0.0195*** (0.00654)	-0.0151** (0.00776)	
2001 Take-up Rate	Yes	Yes	Yes	
Unemployment		Yes	Yes	
Population			Yes	
Employment growth			Yes	
Total Labor Force			Yes	
<i>Standard errors clustered by 50 states</i>				
Panel C: Local vs. National Brand Preferences				
	Difference in Continued Products Inflation Rate (pp)			
Change in Take-up Rate (pp), 2001-2007	-0.00828** (0.00398)	-0.01687*** (0.00523)	-0.000597 (0.00560)	
All product groups	X			
Top 50% by "local" preferences		X		
Bottom 50% by "local" preferences			X	
<i>Standard errors clustered by 50 states</i>				
Panel D: Food vs. Non-Food Products				
	Difference in Continued Products Inflation Rate (pp)			
Change in Take-up Rate (pp), 2001-2007	-0.0187** (0.00808)	-0.0115 (0.0203)	-0.0255*** (0.0110)	-0.0124 (0.0078)
Food product groups	X		X	X
Non-food product groups		X		
Top 50% by "local" preferences			X	
Bottom 50% by "local" preferences				X
<i>Standard errors clustered by 50 states</i>				

Table D9 shows these results in detail. Panel A summarizes the main results. A 10 percentage point increase in the take-up rate across states (which was the mean increase during this period) leads to a 2.2% increase in spending from SNAP-eligible households, and to a 24.2 basis point fall in inflation for these households, relative to SNAP-ineligible households. Panel B shows the robustness of this finding to the

inclusion of a series of controls, alleviating concerns about omitted variable biases.⁷²

Panel C repeats the analysis at the level of product groups and shows that the effect is driven by product groups with strong local brand preferences, consistent with the hypothesized mechanism. Finally, Panel D tests whether the effect is stronger for food products, which one would expect if recipients do not treat food stamps as fungible income.⁷³ Indeed, the effect is significant only in food categories, and within that set of products it is driven by the product categories with stronger local preferences (note that the point estimates for non-food products are not significant but are not precisely estimated zeroes).

D.5.3 Calibration: Changes in Spending per Capita across Income Groups

I calibrate the magnitude of the per capita spending channel by using the observed change over time in the average income of households making above \$100,000 and below \$30,000. Between 2004 and 2013, the average income of high-income households grew 0.93 percentage point faster than that of low-income households. By taking the ratio of the point estimates in Panel A of Table D9, I obtain that a 1 percentage point increase in spending per capita leads to a $\frac{24.2}{2.226} = 10.9$ basis point fall in inflation. Therefore, the annual inflation difference caused by rising inequality is equal to $0.93 \times 10.9 = 10.1$ basis points, which represents $\frac{10.1}{40.8} = 24.7\%$ of the benchmark inflation difference and $\frac{10.1}{66.0} = 15.3\%$ of the overall inflation difference.

Rising income inequality therefore has a sizable amplification effect on real inequality: the amplification factor is about one tenth. However, over the course of my sample this channel played a quantitatively less important role in lowering inflation for the high-income relative to the low-income compared with the “number of consumers” channel.

D.6 Alternative Mechanisms

I investigate various alternative explanations for the evidence. I first study in depth two mechanisms that may disproportionately benefit the poor: the product cycle and international trade. I show that although these mechanisms appear to indeed play a role and benefit the poor relatively more, they are quantitatively less important than other channels that disproportionately benefit the high-income. I then study a series of other possible mechanisms and find that they can’t be the primary drivers of the patterns found in the data.

The product cycle. I find that the difference in quality-adjusted inflation for the high- and low-income households is lower in product modules in which the “product cycle” is faster. Intuitively, if there is a high rate of product churn (a fast “product cycle”), then it is less easy for manufacturers to customize products and introduce new varieties, which will rapidly become outdated. Consumer electronics are a good example illustrating this idea: in that sector, the difference in quality adjusted inflation between low- and high-income households is close to 20 basis points, a third of the sample average. More broadly, I find that across product modules, a one standard deviation increase in the rate of “product churn” (measured as the sum of the

⁷²I have conducted a number of other falsification tests, not reported in this version of the draft but available upon request. In particular, I have compared inflation patterns for households in other parts of the income distribution (e.g. \$30k - \$100k) and found that they were not correlated with the increase in SNAP take-up rate.

⁷³On the fungibility of money and spending choices, see Shapiro and Hastings (2013).

share of spending on new products and the share of spending on products about to exit) is correlated with a 9.18 basis point decline ($t = 1.98$) in the difference in quality-adjusted inflation between low- and high-income households. These results provide partial support for the view that the product cycle tends to benefit “everyone” — but the dynamics of increasing product variety appear to matter more quantitatively.

International trade. Does trade with China disproportionately benefit the poor? This intuition is widespread and I do find support for it in the data, but this channel is not sufficient to outweigh the other forces at play that benefit the high-income relatively more. Matching HS6 code import data to Nielsen category by hand, I find that inequality in quality-adjusted inflation is lower in product modules with higher import penetration from China. Across product modules, a 10 percentage point increase in import penetration rank is correlated with a 6.23 basis point decline ($t = 2.03$) in the difference in quality-adjusted inflation between low- and high-income. In product modules above the median of import penetration, the difference in quality adjusted inflation between low- and high-income households is around 30 basis points, one half of the sample average. In other words, competitive dynamics from international trade tend to benefit the poor relatively more, but this effect does not outweigh the domestic competitive dynamics, which tend to disproportionately benefit the high-income.

Aggregate shocks. First, the various decompositions reported in Section 3 show that the results are not driven by broad shocks that would be specific to certain areas (Table D2) or to certain departments, product groups or product modules (Tables 3 and 2).

Online retail. The rise of online retail could have differentially benefited high- and low-income households. For instance, if higher-income households are more technology savvy, they might be more likely to use online platforms to search for products, which would increase their price elasticity and result in lower equilibrium markups. However, the inflation difference across product categories is not related to heterogeneity in exposure to online retail - in particular, it persists in categories that were very little affected by online retail during this period, such as food (Table 3).

Innovation dynamics independent of changes in market size. An alternative view of the innovation patterns is that product innovation may always be skewed towards the higher-income consumers, regardless of the underlying patterns of growing inequality. In other words, the patterns documented in Section 3 may be a steady state. By introducing flexible controls for the income distribution of consumers and for the quality distribution (price deciles) within a product module, Panel B of Appendix Table D6 shows that the estimated response of product innovations to market size is not confounded by static patterns related to income or quality. Moreover, I have not found empirical support for the predictions of a simple class of models that generate a steady-state difference in the inflation rates experienced by high- and low-income households - in these models, the equilibrium price elasticity of higher-income consumers should always be lower.⁷⁴

⁷⁴Intuitively, if high-income consumers are less price elastic and if the cost of increasing product variety is linear, in equilibrium we will observe a high flow of new products targeting higher income consumers. The equilibrium mechanism is that the high-end products have higher margins (because the high-income consumers are less price elastic) but have a shorter lifecycle (because they get displaced by other high-end product innovations).

Household search behavior. Another possible channel for the results is that high-income consumers could have become more price elastic because their search behavior has changed. Such a channel would manifest itself primarily through within-UPC inflation difference between high- and low-income households, which Table 2 shows is not the case.

E Theory Appendix

E.1 Predictions from Competing Models

A variety of models can generate the key prediction that in general equilibrium the quality-adjusted price goes down when demand increases. There are three broad classes of such models: endogenous growth macro models with scale effects (e.g. Romer, 1990, Aghion and Howitt, 1992, and Acemoglu and Linn, 2004), trade models with free entry and endogenous markups through variable-elasticity-of-substitution preferences (e.g. Melitz, 2003, and Zhelobodko et al., 2012), and industrial organization models with free entry and endogenous markups through strategic interactions between firms (e.g. Sutton, 1991, and Berry and Reiss, 2006). Intuitively, in all of these models, when demand rises product variety increases through entry⁷⁵, and the price of continuing products decreases either because of a decrease in marginal cost⁷⁶ or because of a fall in markups.⁷⁷

Although their key prediction is similar, these models differ in important ways. First, it is important to establish whether quality-adjusted inflation is driven by the *level* of market size or by *changes* in market size. In most macro models, a permanent change in market size will have a *permanent* effect on the rate of economic growth: the returns from innovation are larger in bigger markets because the cost of innovation (assumed to be linear) can be spread out over more consumers, and therefore the level of innovation is always higher in bigger markets. Semi-endogenous growth models with decreasing returns to scale in the R&D production function (Jones, 1995) and models with endogenous markups and free entry offer a competing view, according to which an increase in market size will only have a *temporary* effect on the level of innovation. In other words, *changes* in market size are the relevant predictors of innovation, not the level of market size. Intuitively, endogenous changes in markups or the increased cost of innovation prevent scale effects from permanently raising the level of innovation. In Section 4.3.4, I conduct a direct test to distinguish between these competing views and I find support for the idea that changes in market size matter, rather than the level of market size.

Second, as previously mentioned, in some models the fall in inflation on continuing products results from a fall in markups, while in others it results from a fall in marginal cost. Using data on retailer markups and a double marginalization model that allows me to extrapolate these patterns to manufacturer markups, I provide suggestive evidence that most of the effect comes from changes in markups.

⁷⁵Recent work in the trade tradition models entry of products within multi-product firms, e.g. Mayer, Melitz and Ottaviano (2016).

⁷⁶In models in the macro tradition, the fall in marginal cost can either be exogenous or endogenous to the firm's decisions. Exogenous falls in marginal costs stem from increasing returns to scale (e.g. Matsuyama, 2002). In models with endogenous investment in marginal cost, the returns to marginal cost improvements increase with market size (e.g. Acemoglu and Linn, 2004).

⁷⁷In models in the trade tradition, markups fall because consumers move along their demand curves to a point with a higher price elasticity; while in models in the industrial organization tradition, markups fall because a larger market can sustain more firms and an increase in the number of firms reduces markups through strategic interactions.


Finally, “demand” is not a well-defined primitive object in any of these models. Rather, changes in demand in a given market could result from either a change in the *number of consumers* or from a change in *spending per capita*, respectively denoted L and E in the model introduced above. Depending on the model, variation in the number of consumers and variation in per capita spending could have different effects on the equilibrium.⁷⁸ Empirically, I find that these effects are in fact very similar.

In light of the results of the various tests reported in the remainder of this section, I develop my preferred model by relying on translog preferences with flexible preference parameters across income groups. In contrast with the other models mentioned above, my preferred model yields predictions in line with all aspects of the data: changes in market size drive the effect (rather than the level), falling markups are key, and changes in demand coming from changes in the number of consumers or from changes in per capita spending lead to the same endogenous supply response.

E.2 Intuitions

E.2.1 Reduced-Form Approach

Figure E1: Does the Price Fall When Demand Rises?



Because of nonhomothetic preferences and the endogenous price changes induced by changes in relative demand, changes in nominal inequality may overstate or understate changes in purchasing-power inequality. Consider Figure E1. When relative demand goes up, if the short-run supply curve is upward-sloping as in

⁷⁸For instance, in Zhelobodko et al. (2012) changes in spending per capita will only result in an impact on the equilibrium number of varieties, while the price of continuing products will be unaffected. In contrast, changes in the number of consumers will also lead to a fall in the price of continuing products.

standard price theory, then the equilibrium price should go up. However, supply may endogenously shift out due to the response of firms to market size effects. The price increase will at least be mitigated. As illustrated in Figure E1, the new equilibrium price could even be lower than the initial equilibrium price. This “price overshooting” is found to be relevant empirically in Section 4. In other words, the observed long-term supply curve is downward-sloping.⁷⁹

To investigate whether changes in nominal inequality overstate or understate changes in real inequality, the following concepts are useful:

- **Weak equilibrium (relative) bias** (“directed technical change”): when demand for a good becomes relatively more abundant, supply (technology, innovation, entrepreneurship, etc.) becomes endogenously biased towards this factor.
- **Strong equilibrium (relative) bias**: the relative supply curves for goods are downward-sloping.

Consider demand H for a high-quality good and demand L for a low-quality good. Endogenous technology A is a function of relative demand $\frac{H}{L}$. The equilibrium relative price is

$$\frac{p_H}{p_L} = f\left(\frac{H}{L}, A\left(\frac{H}{L}\right)\right)$$

There is weak equilibrium bias if:

$$\frac{\partial f}{\partial A} \frac{\partial A}{\partial H} < 0$$

There is strong equilibrium bias if:

$$\frac{\partial f}{\partial H} + \frac{\partial f}{\partial A} \frac{\partial A}{\partial H} < 0$$

where one could have $\frac{\partial f}{\partial H} > 0$, as in standard price theory.

The equations above and Figure E1 provide an intuitive reduced-form way of thinking about the effect of shifts in demand on the equilibrium price.

E.2.2 A Simple Microfoundation

Figure E1 provides an intuitive reduced-form way of thinking about the effect of shifts in demand on the equilibrium price. I now turn to providing a microfoundation for this effect, focusing on microfounded models of monopolistic competition with free entry.⁸⁰ The intuition for the effect of changes in market size on supply in monopolistic competition models is as follows: an increase in market size leads to more product entry, which puts downward pressure on the prices of existing products (pecuniary externality). Therefore, in such models innovation occurs entirely through product entry - there is no “process innovation” reducing the marginal cost of the existing products, whose price dynamics are determined by changes in markups.

⁷⁹The “observed” long-term supply curve is defined as the nexus of equilibrium points traced out by shifts in the demand curve. The concept of “observed” supply curve is useful in the context of monopolistic competition, where firms are not price takers and where the usual notion of “supply curve” is therefore not well defined.

⁸⁰This broad class of models is appealing for two reasons: the assumption of monopolistic competition is reasonable in retail, and these models nest the standard model of directed technical change (Acemoglu, 2002).

Within the class of monopolistic competition models with free entry of products, only some models are consistent with the “price overshooting” case illustrated in Figure E1. In particular, the CES model of Acemoglu (2002) does not allow for the possibility that the price goes down when demand goes up (see Appendix A for a detailed derivation). On the other hand, Melitz and Ottaviano (2008) is consistent with the strong equilibrium bias (see Appendix A for a derivation). In the rest of this section, I characterize the conditions under which “price overshooting” is possible using the general monopolistic competition model of Zhelobodko, Kokovin, Parenting and Thisse (2012). The key insight is that, in general equilibrium, the curvature of the utility function and variable markups drive the sign and magnitude of the response of the equilibrium price to changes in market size.

L consumers with additively separable preferences over varieties solve:

$$\max_{x_i \geq 0} U = \int_0^N u(x_i) di \quad s.t. \quad \int_0^N p_i x_i di = E$$

Consumer maximization yields

$$p_i(x_i) = \frac{u'(x_i)}{\lambda}$$

$$\lambda = \frac{\int_0^N x_i u'(x_i) di}{E}$$

Total quantity demanded is $q_i = Lx_i$. The monopolist takes the residual demand curve as given and solves:

$$\max \pi(q_i) = R(q_i) - C(q_i) \equiv \frac{u'(q_i/L)}{\lambda} q_i - V(q_i) - F$$

with $V(\cdot)$ the variable cost function and F the fixed cost. The optimal markup of the producer is therefore given by:

$$M^* = - \frac{x_i \cdot u''(x_i)}{u'(x_i)}$$

At the free entry equilibrium, $\pi(q_i^*) = 0$ and a mass N^* of firms satisfies labor market clearing⁸¹:

$$N^* = \frac{L \cdot E}{C(q_i^*)}$$

Therefore, the model delivers the following comparative statics:

$$\frac{dN^*}{dL} > 0 \quad \frac{dx_i^*}{dL} < 0 \quad \frac{dM_i^*}{dL} \begin{matrix} \leq \\ > \end{matrix} 0$$

The optimal markup is given by the inverse of the price elasticity of demand.⁸² This result is very general and holds regardless of the shape of the cost function $V(\cdot)$. It shows why the equilibrium response of prices to changes in market size crucially depends on the curvature of the utility function. The intuition for the comparative statics is as follows. When market size increases, new products enter the market. As a result, consumers start spreading out their expenditures across more products, due to taste for variety.

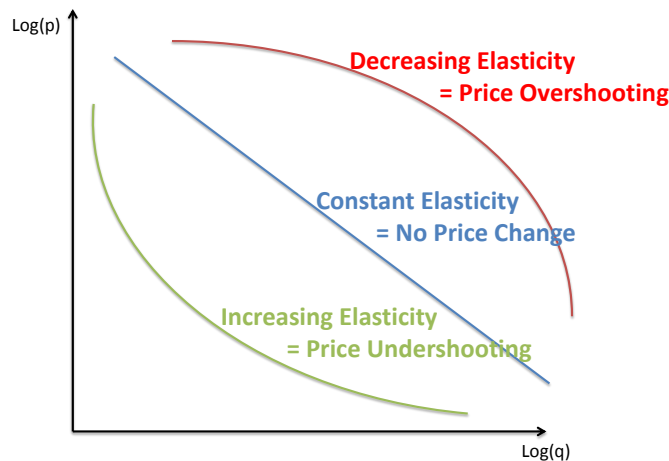
⁸¹A similar model can be solved by assuming that the sector is small relative to the total economy, which allows for ignoring some GE effects. See Mayer, Melitz and Ottaviano (2016).

⁸²The inverse of the price elasticity of demand is equal to the coefficient of relative risk aversion. Given our assumption of separable utility, it is also equal to the inverse of the elasticity of substitution between varieties.

Consequently, consumption per capita x_i for the existing products goes down, which induces a response of the optimal markup M^* . The equilibrium markup may increase, decrease or stay unchanged, depending on the properties of demand. Figure E2 shows this effect in log-log space. The blue curve corresponds to CES demand, as in Acemoglu (2002). Movements along the curve do not matter; the elasticity is constant. On the other hand, the red curve shows that when consumption per capita decreases (moving to the left along the curve), the price elasticity of demand goes up, i.e. the optimal markup goes down. Melitz and Ottaviano (2008) corresponds to this case. Conversely, as shown with the green curve, if the price elasticity of demand is increasing, the equilibrium price should go up in response to an increase in market size.

The market size comparative statics in the case of decreasing elasticity of substitution are in line with the stylized facts documented earlier: through market size effects and endogenous product entry, there should be a strong negative correlation between inflation and the share of spending on new products both *across* and *within* product modules. The main prediction of the model is of course that growing demand *causes* more product innovations and lower inflation. An additional prediction is that the inflation patterns on continuing products are driven by differences in changes in markups.⁸³ I test and find support for these predictions in the rest of this section.

Figure E2: The Equilibrium Response of Price to Changes in Market Size Depends on the Price Elasticity of Demand



⁸³Note that this speaks to an active debate in the trade literature about the source of the gains from trade and the role of variable markups and variable elasticity of substitution preferences. See in particular DeLoecker, Goldberg, Pavcnik and Khandelwal (2012), Feenstra and Weinstein (2016), and Mayer, Melitz and Ottaviano (2016).

E.3 General Equilibrium Model

E.3.1 Goals

The purpose of this model is to offer a *unified framework for the estimation of inflation across income groups and the estimation of the response of supply to market size effects*. In other words, the model features A. well-defined non-homothetic preferences giving rise to tractable income-group specific prices indices; B. closed-form solutions showing the general equilibrium response of supply to shifts in demand across the product space, in terms of endogenous product variety and endogenous markups.

Thus, this model provides a microfoundation for the measurement of inflation across income groups carried out in Section 3, as well as for the regression specifications and comparative statics used in Section 4 to characterize the response of supply to market size effects.⁸⁴

E.3.2 Setting and Main Assumptions

I consider a general equilibrium model with two types of agents differing in their productivity levels, denoted H for high productivity and L for low productivity. The numbers of agents of each type are denoted L^i , with $i = H, L$. These agents consume and produce in K different sectors in the economy. The number of products available in each sector is endogenous and denoted N_k .⁸⁵

In order to obtain tractable closed-form solutions, the following simplifying assumptions are made:

1. The model is static;
2. Firms in each sector are homogeneous, i.e. have the same marginal cost of production;
3. Consumer preferences are non-homothetic *across* sectors (i.e. different agents place different weights on the various sectors, depending on their income levels) but are homothetic *within* sectors (i.e. at the lowest level of aggregation, all agents have the same spending patterns);
4. High- and low-productivity agents enter the production function in a similar way across sectors, which implies that there is no feedback effect of shifting demand on wages across agent types.
5. High- and low-productivity agents pay the same price for each barcode (within sectors).

Each of these assumptions are relaxed in turn in extensions presented later in this appendix.

E.3.3 Consumers

Non-homothetic CES aggregator across sectors. Consumers of each type, indexed by $i = H, L$, maximize aggregate consumption C_i . As in Comin et al. (2016), C_i combines sector goods $\{C_{ik}\}_{k=1}^K$

⁸⁴The model can also be used to clarify the identification assumptions and the potential threats to identification discussed in Section 4.

⁸⁵In my model, product entry and firm entry are analogous, therefore N^K can be thought of as the equilibrium number of firms or the equilibrium number of sectors.

according to the implicitly defined function:

$$\sum_{k=1}^K \Omega_k^{\frac{1}{\sigma}} C_i^{\frac{\epsilon_k - \sigma}{\sigma}} C_{ik}^{\frac{\sigma-1}{\sigma}} = 1 \quad (\text{E1})$$

where σ is the elasticity of substitution and Ω_k 's are constant weights for the various sectors. Each sectoral good k is itself a consumption aggregator, described below, and is characterized by an income elasticity parameter ϵ_k . This is a generalization of the standard (homothetic) CES aggregator, which corresponds to the special case for which $\epsilon_k = 1$ for all sectors. Intuitively, as aggregate consumption C_i increases, the weight given to the consumption of good k varies at a rate controlled by the parameter ϵ_k . As a result, household i 's demand for sectoral good k features a constant elasticity in terms of aggregate consumption C_i , which is in turn determined by household income.⁸⁶

Note that, with only two household types, we can re-write the non-homothetic CES aggregator above as two income-group-specific standard (homothetic) CES aggregator, with income-group-specific sectoral weights $\tilde{\Omega}_{ik} = \Omega_k C_i^{\epsilon_k - 1}$. In other words, each household maximizes aggregate consumption C_i defined as:

$$\sum_{k=1}^K \tilde{\Omega}_k^{\frac{1}{\sigma}} C_i^{\frac{1-\sigma}{\sigma}} C_{ik}^{\frac{\sigma-1}{\sigma}} = 1 \quad \text{i.e.} \quad C_i = \left(\sum_{k=1}^K \tilde{\Omega}_{ik}^{\frac{1}{\sigma}} C_{ik}^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}} \quad \text{for } i = H, L$$

The standard CES results then apply for each household type.⁸⁷ The optimal allocation of expenditures across sectors is characterized by

$$C_{ik} = \tilde{\Omega}_{ik} \left(\frac{P_k}{\mathbb{P}_i} \right)^{-\sigma} C_i$$

where P_k is the sectoral price index and \mathbb{P}_i is the aggregate price index for a household of type i :

$$\mathbb{P}_i \equiv \frac{E_i}{C_i} = \left[\sum_{k=1}^K \tilde{\Omega}_{ik} P_k^{1-\sigma} \right]^{\frac{1}{1-\sigma}} \quad (\text{E2})$$

using the level of expenditures $E_i = \sum_{k=1}^K P_k C_{ik}$ and the expression for C_{ik} above.

Therefore, expenditure shares for each income group across sectors are given by:

$$s_{ik} = \frac{P_k \cdot C_{ik}}{E_i} = \frac{P_k \cdot \tilde{\Omega}_{ik} \left(\frac{P_k}{\mathbb{P}_i} \right)^{-\sigma} C_i}{C_i \cdot \mathbb{P}_i} = \tilde{\Omega}_{ik} \left(\frac{P_k}{\mathbb{P}_i} \right)^{1-\sigma} = \frac{\tilde{\Omega}_{ik} P_k^{1-\sigma}}{\sum_{k=1}^K \tilde{\Omega}_{ik} P_k^{1-\sigma}} \quad (\text{E3})$$

Translog preferences within sectors. Within a sector, consumers have the same translog preferences.

Let \tilde{N}_k be the total number of varieties (or firms) conceivably available in sector k and treat this number as fixed. Dropping the k subscripts for convenience and denoting by p_n the price of variety (or firm) n , within

⁸⁶Useful properties of this utility function include the fact that the elasticity of relative demand for two different goods with respect to aggregate consumption is constant: $\frac{\partial \log(C_i/C_j)}{\partial \log(C)} = \epsilon_i - \epsilon_j$; and the fact that the elasticity of substitution between goods of different sectors is uniquely defined and constant: $\frac{\partial \log(C_i/C_j)}{\partial \log(P_j/P_i)} = \sigma$. See Comin et al. (2016) for more details.

⁸⁷In the estimation carried out in Section 3, $\tilde{\Omega}_{ik}$ can be recovered directly from price and quantity data. The estimation framework also allows for elasticities of substitution σ to vary across income groups.

each sector the translog expenditure function is defined as:⁸⁸

$$\ln(E) = \ln(U) + \alpha_0 + \sum_{n=1}^{\tilde{N}} \alpha_n \ln(p_n) + \frac{1}{2} \sum_{n=1}^{\tilde{N}} \sum_{m=1}^{\tilde{N}} \gamma_{nm} \ln(p_n) \ln(p_m)$$

with $\gamma_{nm} = \gamma_{mn} \forall m, n$. The restrictions $\sum_{n=1}^{\tilde{N}} \alpha_n = 1$ and $\sum_{n=1}^{\tilde{N}} \gamma_{nm} = 0$ ensure that the expenditure function is homogeneous of degree one. Following the literature, I impose that all goods enter ‘‘symmetrically’’ into the expenditure function, i.e. $\alpha_n = \frac{1}{\tilde{N}}$, $\gamma_{nn} = \frac{-\gamma(\tilde{N}-1)}{\tilde{N}}$ and $\gamma_{nm} = \frac{\gamma}{\tilde{N}}$ for $m \neq n$, with $n, m = 1, \dots, \tilde{N}$. In the presence of unavailable goods, the expenditure function becomes complicated, involving their reservation prices. However, in the symmetric case defined above, Feenstra (2003) shows that the expenditure function can be simplified considerably, so that the reservation prices no longer appear explicitly. Specifically, imposing the symmetry restrictions and $\gamma > 0$ and assuming that only the goods $n = 1, \dots, N$ are available, Feenstra (2003) shows that the expenditure function can be written in a way such that reservation prices no longer appear:

$$\ln(E) = \ln(U) + a_0 + \sum_{n=1}^N a_n \ln(p_n) + \frac{1}{2} \sum_{n=1}^N \sum_{m=1}^N b_{nm} \ln(p_n) \ln(p_m)$$

where $a_n = \frac{1}{N}$, $b_{nn} = -\frac{\gamma(N-1)}{N}$ and $b_{nm} = \frac{\gamma}{N}$ for $m \neq n$, with $n, m = 1, \dots, N$, and $a_0 = \alpha_0 + \frac{1}{2} \frac{\tilde{N}-N}{\gamma N \tilde{N}}$. The second term appearing in a_0 reflects the welfare gains of increasing the number of available products - it shows that these gains are smaller and smaller as N approaches \tilde{N} , a key feature of translog (i.e. there are decreasing returns to increasing product variety, in contrast with the CES case).

By Shephard’s lemma, the spending share on each variety n is given by:

$$s_n = \frac{1}{N} + \gamma \left(\overline{\ln(p)} - \ln(p_n) \right) \quad (\text{E4})$$

with $\overline{\ln(p)} = \frac{1}{N} \sum_{m=1}^N \ln(p_m)$. Thus, a 1% increase in the price of a product, holding the overall mean price fixed, lowers its expenditure share by γ percentage points.

Therefore, the elasticity of demand for each product n is given by:

$$\epsilon_n = 1 - \frac{d \ln(s_n)}{d \ln(p_n)} = 1 + \frac{(N-1) \cdot \gamma}{N \cdot s_n} \quad (\text{E5})$$

i.e. the elasticity of demand is decreasing in expenditure share, therefore it is increasing in price. The elasticity of demand is the key feature of preferences which determines firms’ optimal markups in equilibrium.

Labor supply. Labor is supplied inelastically. High-productivity households are endowed with l^H effective units of labor, as against l^L effective units of labor for low-productivity households. The wage for one effective unit of labor is the numeraire.

Budget constraint. Each household type is subject to the budget constraint:

$$\sum_{k=1}^K \sum_{n=1}^{N_k} c_{ink} \cdot p_{nk} = l^i \quad \forall i \quad (\text{E6})$$

⁸⁸See Diewert (1974) and Feenstra (2003).

Two-stage budgeting. In addition to the budget constraint, equations E1, E3, E4 and E5 above completely characterize the demand side of the model. I have shown how the non-homothetic CES-aggregator can be re-written as separate income-group-specific homothetic CES aggregators across sectors. The translog utility functions within sectors are also homothetic, which makes it possible to rely on two-stage budgeting: households first allocate their expenditures across sectors according to E3, and then within sector according to E4. Their equilibrium demand elasticity is given by E5 and is the key equilibrium object that governs optimal markups.

E.3.4 Producers

Homogeneous firms/varieties within sectors. Within each sector k , symmetric firms (or, alternatively, varieties) enter until profits are brought to zero. To obtain tractable closed-form solutions, the cost structure is assumed to be homogeneous across firms within a sector: firms have the same marginal cost, produce varieties of the same quality, and incur the same entry cost. Therefore, firm/variety subscripts n can be dropped in what follows.

Labor demand. All firms/varieties produce using the same production function. The quantity produced by a single firm/variety is given by $q_k = Z_k l_k$, where l_k is labor demand for production and Z_k is a productivity factor specific to sector k . Moreover, all firms pay the same sunk “entry cost” equal to f_k effective units of labor. The required amount of labor for entry per firm is therefore $\frac{f_k}{Z_k}$. Thus, the total amount of labor required by all firms in sector k is:

$$L_k = \underbrace{N_k \cdot \frac{q_k}{Z_k}}_{\text{production}} + \underbrace{N_k \cdot \frac{f_k}{Z_k}}_{\text{entry}}$$

Optimal markups. Firms are monopolistically competitive. Therefore, the same formula as described in Section 4 applies. In each sector k , firms charge an optimal markup equal to $\frac{1}{\epsilon_{nk}} - 1$. Using the fact that $s_{nk} = \frac{1}{N_k}$ by symmetry, from E5 we have that the optimal markup for product n in sector k is given by:

$$\mu_{nk} - 1 = \mu_k - 1 = \frac{1}{\epsilon_k - 1} = \frac{1}{(N_k - 1) \cdot \gamma_k} \quad (\text{E7})$$

In contrast with the CES case, the optimal markup is decreasing in the number of products.

Free entry. Firms enter sector k until profits are brought to zero via decreasing markups. Using the fact that the wage is the numeraire, this can be written as:

$$\pi_k = (\mu_k - 1) \frac{q_k}{Z_k} - \frac{f_k}{Z_k} = 0 \quad \forall k \quad (\text{E8})$$

where q_k is the total quantity produced by the firm in market k in equilibrium.

E.3.5 Equilibrium (Proof of Proposition 1)

The equilibrium objects are as follows:⁸⁹

1. Equilibrium consumption levels for each variety n and each household type i in each sector k ; by symmetry, $c_{ink} = c_k^i$
2. Equilibrium production levels for each variety n in each sector k ; by symmetry, $q_{nk} = q_k$
3. Equilibrium prices for each variety n in each sector k ; by symmetry, $p_{nk} = p_k$
4. Equilibrium number of varieties N_k in each sector k
5. Optimal markup for each variety n in each sector k ; by symmetry, $\mu_{nk} = \mu_k$
6. Total labor demand L_k for each sector k

The equilibrium conditions are as follows:

1. Optimal allocation of spending across sectors: equation E3
2. L^L budget constraints for the low-income and L^H budget constraints for the high-income: equation E6
3. Optimal markups in each sector k : equation E7
4. Zero-profit condition in each sector k : equation E8
5. Equilibrium of supply and demand for each variety n in each sector k ; by symmetry:

$$q_k = L^H \cdot c_k^H + L^L \cdot c_k^L$$

6. Equilibrium of supply and demand for labor:

$$\underbrace{\sum_k \frac{N_k}{Z_k} (L^H \cdot c_k^H + L^L \cdot c_k^L + f_k)}_{\text{total demand for effective labor units}} = \underbrace{l^H \cdot L^H + l^L \cdot L^L}_{\text{total endowment of effective labor units}}$$

Intuitively, in equilibrium households maximize their utility subject to their budget constraint (conditions 1 and 2), no existing firm can increase its profit by changing its output (condition 3), no firm can enter any sector and make positive profits (condition 4), the product market clears (condition 5), and the labor market clears (condition 6).

We can now solve the model. From E8 and the equilibrium of supply and demand for each variety,

$$(\mu_k - 1) (L^H \cdot c_k^H + L^L \cdot c_k^L) = f_k \quad \forall k$$

where $\mu_k = \frac{1}{(N_k - 1) \cdot \gamma_k}$ by E7. Using the equilibrium spending shares across sectors from E3 and symmetry within sector, the equilibrium spending on each variety within each sector k is given by:

$$e_k \equiv p_k \cdot (L^H \cdot c_k^H + L^L \cdot c_k^L) = L^H \cdot \frac{l^H \cdot s_{Hk}}{N_k} + L^L \cdot \frac{l^L \cdot s_{Lk}}{N_k} \quad \forall k$$

⁸⁹Note that the analysis can directly be extended to an arbitrary number of income groups.

The optimal price is the optimal markup over the marginal cost, i.e. $p_k = \mu_k \frac{1}{Z_k} = \frac{1+(N_k-1)\cdot\gamma_k}{(N_k-1)\cdot\gamma_k} \frac{1}{Z_k} \quad \forall k$.
 Plugging this back in E8, we can write:

$$\begin{aligned}
 (\mu_k - 1) \frac{e_k}{p_k} &= f_k \\
 \frac{1}{(N_k - 1) \cdot \gamma_k} \frac{\frac{L^H \cdot l^H \cdot s_{Hk} + L^L \cdot l^L \cdot s_{Lk}}{N_k}}{\frac{1+(N_k-1)\cdot\gamma_k}{(N_k-1)\cdot\gamma_k} \frac{1}{Z_k}} &= f_k \\
 \frac{(L^H \cdot l^H \cdot s_{Hk} + L^L \cdot l^L \cdot s_{Lk}) \cdot Z_k}{N_k(1 + (N_k - 1) \cdot \gamma_k)} &= f_k \\
 \gamma_k N_k^2 + (1 - \gamma_k)N_k - \frac{(L^H \cdot l^H \cdot s_{Hk} + L^L \cdot l^L \cdot s_{Lk}) \cdot Z_k}{f_k} &= 0
 \end{aligned} \tag{E9}$$

Therefore, in equilibrium:

$$N_k^* = \frac{(\gamma_k - 1) + \sqrt{(1 - \gamma_k)^2 + 4\gamma_k \frac{(L^H \cdot l^H \cdot s_{Hk} + L^L \cdot l^L \cdot s_{Lk}) \cdot Z_k}{f_k}}}{2\gamma_k} \tag{E10}$$

This also gives us the equilibrium price of each variety, given that the optimal markup over marginal cost is entirely determined by the equilibrium number of varieties:

$$p_k^* = \frac{1 + (N_k^* - 1) \cdot \gamma_k}{(N_k^* - 1) \cdot \gamma_k} \frac{1}{Z_k} \tag{E11}$$

Following Feenstra (2003) and using symmetry of all varieties, the aggregate price index within each sector is given by:

$$\ln(P_k^*) = \alpha_{0k} + \frac{1}{2} \frac{1 - \frac{N_k^*}{N_k}}{\gamma_k N_k^*} + \frac{1 + (N_k^* - 1) \cdot \gamma_k}{(N_k^* - 1) \cdot \gamma_k} \frac{1}{Z_k} \tag{E12}$$

Note that welfare goes up (i.e. P_k^* goes down) as the equilibrium number of varieties N_k^* increases because of two forces. First, consumers love variety, which is captured by the term $\frac{1}{2} \frac{1 - \frac{N_k^*}{N_k}}{\gamma_k N_k^*}$ (note that this term decreases at a decreasing rate as N_k^* increases, because the product space gets filled and there are decreasing returns to increasing product variety). Second, an increasing number of varieties leads to lower markups, which is reflected by the term $\frac{1+(N_k^*-1)\cdot\gamma_k}{(N_k^*-1)\cdot\gamma_k} \frac{1}{Z_k}$.

It can be checked that the equilibrium solution above satisfies the labor market clearing condition:

$$\begin{aligned}
 \sum_k \frac{N_k}{Z_k} (L^H \cdot c_k^H + L^L \cdot c_k^L + f_k) &= \sum_k \frac{N_k}{Z_k} \left(\frac{f_k}{(\mu_k - 1)} + f_k \right) \\
 &= \sum_k \frac{N_k \cdot f_k}{Z_k} (\gamma_k N_k - \gamma_k + 1) \\
 &= \sum_k \left(\frac{f_k}{Z_k} \gamma_k N_k^2 - \frac{f_k}{Z_k} (1 - \gamma) N_k \right) \\
 &= \sum_k (L^H \cdot l^H \cdot s_{Hk} + L^L \cdot l^L \cdot s_{Lk}) \\
 &= l^H \cdot L^H + l^L \cdot L^L
 \end{aligned}$$

where the first line follows from E8 and the fourth line follows from E9. This completes the derivation of the equilibrium.⁹⁰

E.3.6 Comparative Statics

The model delivered simple closed-form solutions for the equilibrium number of varieties, the price of each variety, the sectoral price index, and the (non-homothetic) aggregate price index for each household type. The regression specifications in Section 4 examine the equilibrium response of the sectoral price index (where a sector is defined as a product module by price decile) and the total number of varieties to changes in the number of consumers and per capita spending. This can be shown directly in the model by taking comparative statics with respect to L^i , the number of consumers of type i , and l^i , nominal spending per capita for consumers of type i .

Equation E10 shows that changes in the number of consumers (L^i) and in spending per capita (l^i) have the same effect on the equilibrium, consistent with the result found in Section 4. From equations E10, E11, E12 and E2, the comparative statics of interest are:

$$\begin{aligned}\frac{dN_k^*}{dL^i} &= \frac{dN_k^*}{dl^i} > 0 \quad \forall k, i \\ \frac{dp_k^*}{dL^i} &= \frac{dp_k^*}{dl^i} < 0 \quad \forall k, i \\ \frac{dP_k^*}{dL^i} &= \frac{dP_k^*}{dl^i} < 0 \quad \forall k, i \\ \frac{d\mathbb{P}_i^*}{dL^i} &= \frac{d\mathbb{P}_i^*}{dl^i} < 0 \quad \forall i\end{aligned}$$

Thus, when either the number of consumers of a certain type or spending per capita from consumers of that type increase, the total number of varieties increase, the price of each variety decreases, the sectoral price index decreases, and the aggregate price index decreases. The observed supply curve is downward-sloping, in line with the regression results and the intuition presented in Section 4.⁹¹

E.3.7 Proof of Proposition 2

Consider two periods, $t-1$ and t . Represent the change in the income distribution between these two periods by a set $\{g_i\}_{i=1}^I$ of growth rates in the number of households with income (productivity) l^i , i.e such that $L_t^i = (1 + g_{it})L_{t-1}^i$. For each income group i , define the welfare-relevant market size effect implied by changes

⁹⁰Technically, the conditions derived above are only necessary conditions. To show that there are also sufficient conditions for existence and unicity of the equilibrium, we only need to show that N_k^* is uniquely determined. s_{ik} is monotonically increasing in N_k^* , which establishes existence.

⁹¹Note that we could also include transfers in the model, for instance to speak directly to the specifications used for the SNAP research design in Section 4. With transfers, the budget constraint becomes:

$$\sum_{k=1}^K \sum_{n=1}^{N_k} c_{ink} \cdot p_{nk} = l^i + T^i$$

where T^i denotes government transfers, such that $L^H \cdot T^H + L^L \cdot T^L = 0$. It follows immediately that changes in T^i have exactly the same effects on the equilibrium as changes in l^i : the comparative statics are identical.

in the income distribution as:

$$\widetilde{g}_{it} = \sum_{k=1}^K s_{ik(t-1)} \cdot \left(\sum_{j=1}^I s_{kj(t-1)} \cdot g_{jt} \right)$$

where $s_{kj(t-1)}$ denotes the share of j in total spending in k at $t-1$.

By (E2), changes in the price index across the income distribution are given by:

$$\begin{aligned} \pi_t^i &\equiv \log(\mathbb{P}_t^i) - \log(\mathbb{P}_{t-1}^i) \\ &= \frac{1}{1-\sigma} \log \left(\sum_{k=1}^K \widetilde{\Omega}_{ik} P_{kt}^{1-\sigma} \right) - \frac{1}{1-\sigma} \log \left(\sum_{k=1}^K \widetilde{\Omega}_{ik} P_{k(t-1)}^{1-\sigma} \right) \\ &= \frac{1}{1-\sigma} \log \left(\frac{\sum_{k=1}^K \widetilde{\Omega}_{ik} P_{kt}^{1-\sigma}}{\sum_{k=1}^K \widetilde{\Omega}_{ik} P_{k(t-1)}^{1-\sigma}} \right) \\ &= \frac{1}{1-\sigma} \log \left(\underbrace{\sum_{k=1}^K \frac{\widetilde{\Omega}_{ik} P_{kt}^{1-\sigma}}{\sum_{k=1}^K \widetilde{\Omega}_{ik} P_{k(t-1)}^{1-\sigma}}}_{=s_{ik(t-1)}} \left(\frac{P_{kt}}{P_{k(t-1)}} \right)^{1-\sigma} \right) \end{aligned} \quad (\text{E13})$$

The fall in the price index is larger for households who were spending more on parts of the product space where the sector price index declined faster, i.e. $\frac{P_{kt}}{P_{k(t-1)}}$ is low.

Next, let's show that the fall in $\frac{P_{kt}}{P_{k(t-1)}}$ is governed by the term $\sum_{j=1}^I g_{jt} \cdot s_{jk(t-1)}$. Note that if the spending shares do not change much across periods, i.e. $s_{jk(t-1)} \approx s_{jkt}$, then $\sum_{j=1}^I g_{jt} \cdot s_{jk(t-1)}$ gives the percentage change in spending in sector k between the $t-1$ and t (since total spending at $t-1$ is given by $L^H \cdot l^H \cdot s_{Hk} + L^L \cdot l^L \cdot s_{Lk}$, or $(\sum_i L_t^i \cdot l^i \cdot s_{ikt})$ with more than two groups). From (E10), (E11) and (E12), we know that the equilibrium sectoral price index is entirely determined by spending, i.e. we can write:

$$\begin{aligned} P_{k(t-1)} &= f_k \left(\sum_i L_t^i \cdot l^i \cdot s_{ik(t-1)} \right) \\ P_{kt} &= f_k \left(\sum_i (1 + g_{it}) \cdot L_t^i \cdot l^i \cdot s_{ik(t-1)} \right) \end{aligned}$$

Log-linearizing,

$$\frac{P_{kt}}{P_{k(t-1)}} - 1 \approx \epsilon_k \cdot \underbrace{\left(\sum_{j=1}^I g_{jt} \cdot s_{kj(t-1)} \right)}_{\text{growth of demand}}$$

where ϵ_k is the (semi-)elasticity of the sectoral price index to a change in market size, which in general depends on γ_k , $N_{k(t-1)}$, Z_k , f_k and initial demand $(\sum_i L_t^i \cdot l^i \cdot s_{ikt})$. Assume that this elasticity is negative and similar across sectors: $\epsilon_k = \epsilon < 0 \forall k$, in line with the evidence in Sections 4 and D.5. Intuitively, this assumption means that supply responds in a similar way to proportional changes in market size across the

product space. Therefore,

$$\begin{aligned} \left(\frac{P_{kt}}{P_{k(t-1)}} \right)^{1-\sigma_i} &\approx \left(1 + \epsilon \cdot \left(\sum_{j=1}^I g_{jt} \cdot s_{kj(t-1)} \right) \right)^{1-\sigma_i} \\ &\approx 1 + (1 - \sigma_i) \cdot \epsilon \cdot \left(\sum_{j=1}^I g_{jt} \cdot s_{kj(t-1)} \right) \end{aligned} \quad (\text{E14})$$

where the second line follow from a first-order Taylor expansion. Equation (E14) shows that sectoral price indices fall disproportionately in parts of the product space that grow faster.

Next, to simplify the analysis and in line with the results in Table C2, assume that $\sigma_i = \sigma_m = \sigma$.⁹² Then,

$$\begin{aligned} \widetilde{g}_{it} > \widetilde{g}_{mt} &\iff \sum_{k=1}^K s_{ik(t-1)} \cdot \left(\sum_{j=1}^I g_{jt} \cdot s_{kj(t-1)} \right) > \sum_{k=1}^K s_{mk(t-1)} \cdot \left(\sum_{j=1}^I g_{jt} \cdot s_{kj(t-1)} \right) \\ &\iff 1 + (1 - \sigma) \cdot \epsilon \cdot \sum_{k=1}^K s_{ik(t-1)} \cdot \left(\sum_{j=1}^I g_{jt} \cdot s_{kj(t-1)} \right) < 1 + (1 - \sigma) \cdot \epsilon \cdot \sum_{k=1}^K s_{mk(t-1)} \cdot \left(\sum_{j=1}^I g_{jt} \cdot s_{kj(t-1)} \right) \\ &\iff \sum_{k=1}^K s_{ik(t-1)} \cdot \left(1 + (1 - \sigma) \cdot \epsilon \sum_{j=1}^I g_{jt} \cdot s_{kj(t-1)} \right) < \sum_{k=1}^K s_{mk(t-1)} \cdot \left(1 + (1 - \sigma) \cdot \epsilon \sum_{j=1}^I g_{jt} \cdot s_{kj(t-1)} \right) \\ &\iff \sum_{k=1}^K s_{ik(t-1)} \cdot \left(\frac{P_{kt}}{P_{k(t-1)}} \right)^{1-\sigma} < \sum_{k=1}^K s_{mk(t-1)} \cdot \left(\frac{P_{kt}}{P_{k(t-1)}} \right)^{1-\sigma} \\ &\iff \frac{1}{1 - \sigma} \log \left(\sum_{k=1}^K s_{ik(t-1)} \cdot \left(\frac{P_{kt}}{P_{k(t-1)}} \right)^{1-\sigma} \right) < \frac{1}{1 - \sigma} \log \left(\sum_{k=1}^K s_{mk(t-1)} \cdot \left(\frac{P_{kt}}{P_{k(t-1)}} \right)^{1-\sigma} \right) \\ &\iff \pi_{it} < \pi_{mt} \end{aligned}$$

where the fourth line follows from (E14) and the sixth line follows from (E13). This completes the proof of Proposition 2.

E.3.8 Implied elasticity of the price of continuing products to changes in product variety

Given E10 and E11, the elasticity of inflation on continued products to product introductions is:

$$\eta = \frac{dp_k^*}{dN_k^*} \frac{N_k^*}{p_k^*} = -\frac{N_k}{N_k - 1} \cdot \frac{1}{1 + (N_k - 1) \cdot \gamma_k}$$

Hence, for large N_k ,

$$\eta \rightarrow -\frac{1}{1 + \frac{1}{\mu_k}}$$

where μ_k is the markup.

According to the Census Annual Retail Trade Survey, retail markups are about 35%. In the empirical analysis reported in Tables 5 and D6, I find that a 1 percentage point increase in demand leads to a 2.7

⁹²This assumption can be relaxed. We only need the two elasticities to not be “too different”, namely they must satisfy:

$$\frac{\log \left(\sum_{k=1}^K s_{ik(t-1)} \cdot \left(1 + (1 - \sigma_i) \cdot \epsilon \sum_{j=1}^I g_j \cdot s_{jk(t-1)} \right) \right)}{\log \left(\sum_{k=1}^K s_{mk(t-1)} \cdot \left(1 + (1 - \sigma_m) \cdot \epsilon \sum_{j=1}^I g_j \cdot s_{jk(t-1)} \right) \right)} < \frac{1 - \sigma_i}{1 - \sigma_m} < \frac{\sum_{k=1}^K s_{mk(t-1)} \cdot \left(\sum_{j=1}^I g_j \cdot s_{jk(t-1)} \right)}{\sum_{k=1}^K s_{ik(t-1)} \cdot \left(\sum_{j=1}^I g_j \cdot s_{jk(t-1)} \right)}$$

percentage point increase in spending on new products, a 1.4 percentage point increase in product exit, and a 40 basis point decline in inflation on continued products. The implied elasticity of the price of continued products to product introductions is $\frac{0.40}{2.7-1.4} = 0.30$. By comparison, the implied elasticity from the model and using the retail markup from the Census Annual Retail Trade Survey is $\frac{1}{1+0.37} = 0.27$.

E.3.9 Extensions

In this section, I present a number of extensions of the model. I first show the results when the lower nest in the demand system is CES instead of translog. I then relax the main simplifying assumptions of the model, introducing 1. dynamics, 2. heterogeneous firms, 3. non-homotheticities within sectors, 4. feedback effects of shifting demand on the relative income of the various consumer types.

Model with CES preferences in lower nest

When preferences within a sector are CES (instead of translog), consumer's demand elasticity becomes independent of the number of varieties and is denoted $\theta_k > 1$ for each sector k . Accordingly, firms' optimal markup is given by:

$$\mu_k - 1 = \frac{1}{\theta - 1}$$

The equilibrium is then characterized by:

$$\begin{aligned} N_k^* &= \frac{(L^H \cdot l^H \cdot s_{Hk} + L^L \cdot l^L \cdot s_{Lk}) \cdot Z_k}{\theta_k \cdot f_k} \\ p_k^* &= \frac{1}{\theta - 1} \frac{1}{Z_k} \\ P_k^* &= N_k^{\frac{1}{1-\theta}} p_k^* \end{aligned}$$

The comparative statics of interest are:

$$\begin{aligned} \frac{dN_k^*}{dL^i} &= \frac{dN_k^*}{dl^i} > 0 \quad \forall k, i \\ \frac{dp_k^*}{dL^i} &= \frac{dp_k^*}{dl^i} = 0 \quad \forall k, i \\ \frac{dP_k^*}{dL^i} &= \frac{dP_k^*}{dl^i} < 0 \quad \forall k, i \\ \frac{d\mathbb{P}_i^*}{dL^i} &= \frac{d\mathbb{P}_i^*}{dl^i} < 0 \quad \forall i \end{aligned}$$

Thus, because of constant markups, CES delivers the prediction that inflation on continuing varieties should not respond to changes in the number of consumers or spending per capita. All welfare effects are through changes in the number of varieties, which is not in line with the results presented in Section 4.

Endogenous savings

To relax the restriction that the model is static, see Bilbiie, Ghironi and Melitz (2012).

Heterogeneous firms

To relax the assumption of homogeneous firms, see Rodriguez-Lopez (2010).

Non-homotheticities within sector

To relax the assumption that consumers have similar preferences (and, in particular, similar elasticities) within sector, see Hottman, Redding and Weinstein (2016).

Multi-product firms

To introduce multi-product firms in the model, see Mayer, Melitz and Ottaviano (2015) and Hottman, Redding and Weinstein (2016).

Feedback effects of shifting demand on the relative income of the various consumer types

To introduce feedback effects, assume that the two types of labor are not perfectly substitutable and that the production function is: $q_k(\omega) = (Z_k^H l_k^H(\omega))^\alpha (Z_k^L l_k^L(\omega))^{1-\alpha}$, where ω denotes a variety. The derivation below shows that this implies that the effective wage faced by the firm is a Cobb-Douglas mix of the wages of the high-productivity and low-productivity agents, which are denoted W^H and W^L respectively.

The cost minimization problem of the firm is:

$$\min_{l_k^H(\omega), l_k^L(\omega)} W^H l_k^H(\omega) + W^L l_k^L(\omega) \quad s.t. \quad (Z_k^H l_k^H(\omega))^\alpha (Z_k^L l_k^L(\omega))^{1-\alpha} = 1$$

This yields:

$$l_k^L(\omega) = \frac{\left(\frac{1-\alpha}{\alpha}\right)^\alpha}{(Z_k^H)^\alpha (Z_k^L)^{1-\alpha}} \left(\frac{W^H}{W^L}\right)^\alpha$$

$$l_k^H(\omega) = \frac{\left(\frac{\alpha}{1-\alpha}\right)^{1-\alpha}}{(Z_k^H)^\alpha (Z_k^L)^{1-\alpha}} \left(\frac{W^L}{W^H}\right)^{1-\alpha}$$

So at the optimum the cost of producing one unit is:

$$\begin{aligned} C &= W^H l_k^H(\omega) + W^L l_k^L(\omega) \\ &= \frac{1}{(Z_k^H)^\alpha (Z_k^L)^{1-\alpha}} \left[\left(\frac{\alpha}{1-\alpha}\right)^{1-\alpha} \left(\frac{W^L}{W^H}\right)^{1-\alpha} W^H + \left(\frac{1-\alpha}{\alpha}\right)^\alpha \left(\frac{W^H}{W^L}\right)^\alpha W^L \right] \\ &= \frac{1}{(Z_k^H)^\alpha (Z_k^L)^{1-\alpha}} \left[\left(\frac{\alpha}{1-\alpha}\right)^{1-\alpha} (W^L)^{1-\alpha} (W^H)^\alpha + \left(\frac{1-\alpha}{\alpha}\right)^\alpha (W^H)^\alpha (W^L)^{1-\alpha} \right] \\ &= \frac{\theta}{(Z_k^H)^\alpha (Z_k^L)^{1-\alpha}} (W^H)^\alpha (W^L)^{1-\alpha} \end{aligned}$$

$$\text{with } \theta = \left(\frac{\alpha}{1-\alpha}\right)^{1-\alpha} + \left(\frac{1-\alpha}{\alpha}\right)^\alpha$$

E.4 Double Marginalization with Monopolistic Retailers and Manufacturers

This notes solves for the optimal markups of the retailer and the manufacturer under the assumption that all products are measure 0 (i.e. there is no cross price effects at either the retailer or manufacturer levels, and all products can be thought of as monopolistic competitors). The only relevant feature of the production process is that there are two levels: products are monopolistically supplied by manufacturers to retailers, which in turn supply these products monopolistically directly to consumers.

E.4.1 Setting

There are L consumers (changes in L will represent changes in market size), there is a representative retailer and a representative manufacturer (keeping track of the number of retailers or manufacturers doesn't matter since products are measure 0). The equilibrium is solved for by backward induction.

E.4.2 Consumer problem and optimization

L consumers with additively separable preferences over varieties solve

$$\max_{x_i \geq 0} U = \int_0^N u(x_i) di \quad s.t. \quad \int_0^N p_i^R x_i di = E$$

Maximization yields

$$\begin{aligned} p_i^R(x_i) &= u'(x_i)/\lambda \\ \lambda &= \frac{\int_0^N u(x_i) di}{E} \end{aligned}$$

Total quantity demanded is $q_i = Lx_i$

E.4.3 Retailer problem and optimization

Retailer bears fixed cost $F^R > 0$ and variable cost $V^R(q_i) = q_i \cdot p_i^M$ (i.e. the marginal cost is given by the price charged by the manufacturer) and charges price p_i^R to the consumer. In monopolistic competition, the retailer solves

$$\max_{q_i \geq 0} \pi^R(q_i) = R^R(q_i) - C^R(q_i) = \frac{u'(q_i/L)}{\lambda} q_i - V^R(q_i) - F^R$$

At the optimum,

$$u'\left(\frac{q_i}{L}\right) + \frac{q_i}{L} u''\left(\frac{q_i}{L}\right) = \lambda V'(q_i) \quad (\text{E15})$$

which can be re-written as the optimal equilibrium retailer markup:

$$M_i^{R*} = \frac{p_i^R - p_i^M}{p_i^R} = -\frac{x_i \cdot u''(x_i)}{u'(x_i)}$$

E.4.4 Manufacturer problem and optimization

In monopolistic competition, the manufacturer solves:

$$\max_{q_i \geq 0} \pi^M(q_i) = R^M(q_i) - C^M(q_i) \equiv p_i^M q_i - c_i^M q_i - F^M$$

At the optimum,

$$\frac{dp_i^M}{dq_i} q_i + p_i^M = c_i^M$$

which can be re-written as the equilibrium manufacturer markup:

$$M_i^{M*} = \frac{p_i^M - c_i^M}{p_i^M} = -\frac{dp_i^M}{dq_i} \frac{q_i}{p_i^M}$$

To solve the optimal manufacturer markup, we just need to find out the equilibrium value of $\frac{dp_i^M}{dq_i}$. We do this starting from (E15) and differentiating by p_i^M using the implicit function theorem. Assuming that $u'''(\frac{q_i}{L}) = 0$ and simplifying, the condition becomes:

$$\frac{dq_i}{dp_i^M} = \frac{\lambda \cdot L}{2 \cdot u''(\frac{q_i}{L})} < 0$$

Substituting in the expression $\lambda = u'(x_i)/p_i^R$ from the consumer maximization problem:

$$\frac{dp_i^M}{dq_i} = \frac{2 \cdot u''(\frac{q_i}{L}) \cdot p_i^R}{u'(\frac{q_i}{L}) \cdot L}$$

Therefore,

$$M_i^{M*} = -\frac{dp_i^M}{dq_i} \frac{q_i}{p_i^M} = -\frac{2 \cdot u''(x_i)}{u'(x_i)} \cdot \frac{q_i}{L} \cdot \frac{p_i^R}{p_i^M} = 2 \cdot M_i^{R*} \cdot \frac{p_i^R}{p_i^M}$$

From the retailer markup we have:

$$\frac{p_i^R}{p_i^M} = \frac{1}{1 + \frac{x_i \cdot u''(x_i)}{u'(x_i)}} = \frac{1}{1 - M_i^{R*}} > 1$$

Therefore,

$$M_i^{M*} = 2 \cdot \frac{M_i^{R*}}{1 - M_i^{R*}}$$

So the markup charged by the manufacturer is larger than the markup charged by the retailer. Note that the manufacturer markup always responds more than the retailer markup to a change in market size:

$$\frac{dM_i^{M*}}{dL} = \frac{2}{(1 - M_i^{R*})^2} \cdot \frac{dM_i^{R*}}{dL}$$

Expressing this as elasticities:

$$\epsilon^{M_i^{M*}} = \frac{1}{1 - M_i^{R*}} \cdot \epsilon^{M_i^{R*}}$$

F External Validity Appendix

This appendix presents additional proofs and empirical evidence speaking to the “external validity” (i.e. beyond retail) of the findings presented in the main text.

F.1 A Lower Bound for the Full Basket Inflation Difference between High- and Low-Income Households: Structural Extrapolation from Nielsen Data.

Assume that for each income group $i = Rich, Poor$, households’ utility function is CES with $\sigma^{Poor} \geq \sigma^{Rich} > 1$ over an aggregator for Nielsen goods, denoted N_i , and an aggregator for outside goods, denoted O_i . In other words, Nielsen goods are assumed to be on average substitutes for goods outside of the Nielsen sample (e.g. food-at-home is in N_i and food-away-from-home is in O_i) and the elasticity of substitution is assumed to be weakly larger for low-income households (intuitively, as their income increases households become less price elastic).

Using CEX data and matching the Nielsen spending categories to CEX categories by hand, I find that during the 2000s, the share of spending on Nielsen product groups for high-income households declined at a rate 0.086 basis points *faster* than for low-income households ($t = 1.99$). This means that high-income households were substituting away from Nielsen goods relative to low-income households, in spite of the lower inflation they were enjoying for this set of goods. Under the assumption that $\sigma^{Poor} \geq \sigma^{Rich} > 1$, this implies that the relative price of the high-income consumption basket was declining even faster for outside goods, relative to the low-income consumption basket.

Formally, for each income group utility is given by:

$$U_i = \left[a_i (N_i)^{\frac{\sigma^i - 1}{\sigma^i}} + (1 - a_i) (O_i)^{\frac{\sigma^i - 1}{\sigma^i}} \right]^{\frac{\sigma^i}{\sigma^i - 1}}$$

with N goods covered by Nielsen and O the outside good. For each income group i , utility maximization yields the familiar formulas for the spending shares S_N^i and S_O^i , sectoral price index P_N^i and P_O^i , and overall price index Π^i . Then,

$$\Delta S_N^{Rich} < \Delta S_N^{Poor} \implies (\Delta \Pi^{Poor} - \Delta \Pi^{Rich}) > \underbrace{(\Delta P_N^{Poor} - \Delta P_N^{Rich})}_{=66bp}$$

In ongoing work, I study the robustness of these results by making adjustment to spending patterns that account for income-group-specific reporting biases in the CEX of the kind documented by Aguiar and Bils (2015). I also repeat the exercise by keeping the income distribution fixed over time within each income group, in order to ensure that the differential evolution of spending shares is not driven by non-homotheticity patterns. Thus, based on the inflation patterns in Nielsen data, basic spending shares from the CEX and economic theory, I show how one can interpret the 66 basis point inflation difference found in the Nielsen

data as a *lower bound* for the full consumption basket inflation difference between high- and low-income households, during the relevant sample period.

Proof. The utility function is CES over Nielsen goods and other goods, with $\sigma^i > 1$ and i indexing household type.

$$U_i = \left[a_i (N_i)^{\frac{\sigma^i-1}{\sigma^i}} + (1-a_i) (O_i)^{\frac{\sigma^i-1}{\sigma^i}} \right]^{\frac{\sigma^i}{\sigma^i-1}}$$

For each income group, the share of spending on each good is given by:

$$\begin{aligned} S_N^i &= a_i^{\sigma^i} \left(\frac{P_N^i}{\Pi^i} \right)^{1-\sigma^i} \\ S_O^i &= (1-a_i)^{\sigma^i} \left(\frac{P_O^i}{\Pi^i} \right)^{1-\sigma^i} \end{aligned}$$

where Π^i is the income-group-specific aggregate price index corresponding to the cost of a unit of utility:

$$\Pi^i = \left(a_i^{\sigma^i} (P_N^i)^{1-\sigma^i} + (1-a_i)^{\sigma^i} (P_O^i)^{1-\sigma^i} \right)^{\frac{1}{1-\sigma^i}}$$

Therefore,

$$\Delta \log(S_N^i) = (1-\sigma^i) (\Delta \log(P_N^i) - \Delta \log(\Pi^i))$$

i.e. income group i substitutes toward good N if and only if the rate of inflation is smaller for good N relative to the full consumption basket (and the degree of substitution is higher if σ^i is higher). Thus,

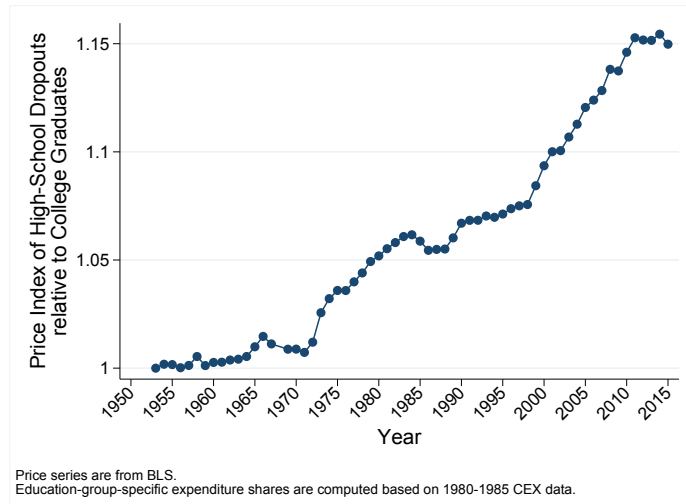
$$\begin{aligned} \Delta \log(S_N^{Rich}) &< \Delta \log(S_N^{Poor}) \\ \iff (1-\sigma^{Rich}) (\Delta \log(P_N^{Rich}) - \Delta \log(\Pi^{Rich})) &< (1-\sigma^{Poor}) (\Delta \log(P_N^{Poor}) - \Delta \log(\Pi^{Poor})) \\ \iff \frac{(1-\sigma^{Rich})}{(1-\sigma^{Poor})} (\Delta \log(P_N^{Rich}) - \Delta \log(\Pi^{Rich})) &> (\Delta \log(P_N^{Poor}) - \Delta \log(\Pi^{Poor})) \end{aligned}$$

Assuming $\sigma^{Poor} \geq \sigma^{Rich} > 1$, we have $\frac{(1-\sigma^{Rich})}{(1-\sigma^{Poor})} < 1$, hence:

$$\implies \Delta \log(\Pi^{Poor}) - \Delta \log(\Pi^{Rich}) > \underbrace{\Delta \log(P_N^{Poor}) - \Delta \log(P_N^{Rich})}_{=66bp}$$

F.2 Long-Run Inflation Inequality across Education Groups

Figure F1: Full-Basket Inflation Inequality across Education Groups in the Long Run



Notes: This figure reports the relative price index of high-school dropouts relative to college graduates over time. The relative price index is normalized to one in 1953. Education-group-specific price indices are built using CPI and CEX data as described in Section 5.2.

F.3 Relative TFP Growth and Patents for High- and Low-Income Households over a Long Time Horizon


In this section, I follow Boppart and Weiss (2016) to provide evidence that technical change has disproportionately benefited high-income households over the long run by using TFP data from the NBER-CES database and patent data from the USPTO, in conjunction with expenditure patterns on goods across income groups.

I proceed in three steps, following the approach of Boppart and Weiss (2016), who report similar results across education groups. First, I convert Consumer Expenditure Survey UCCs to national account PCEs, which yields a dataset with income-group-specific expenditure shares across about 230 product categories. Second, I use the I-O tables to convert the final commodities into industry value-added. Finally, I link the different industries to the NBER-CES database for TFP and to USPTO technology classes.

The results are shown in Figure F2. A clear pattern emerges: over more than five decades, TFP and patents have been biased in favor of the high-income.

Figure F2: Relative TFP Growth and Patents for High- and Low-Income Households over Long Time Horizon

Panel A: Relative TFP



Panel B: Relative Granted Patents

