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

Xavier Jaravel, Stanford University

July 13, 2016

Abstract

Using detailed product-level data in the retail sector in the United States from 2004 to 2013, I find that product innovations disproportionately benefited high-income households due to the endogenous response of supply to shifts in demand, which amplified inequality. My analysis consists of two parts. In the *measurement* part, I show that annual quality-adjusted inflation was 0.66 percentage points lower for high-income households, relative to low-income households. This gap resulted from both lower inflation on continuing products and a faster increase in product variety for the high income. In the *mechanism* part, I use national and local changes in demand that are plausibly exogenous to supply factors — from shifts in the national income and age distributions over time, and from food stamp policy changes across states — to provide causal evidence that a shock to the relative demand for goods (1) affects the direction of product innovations, and (2) leads to a decrease in the relative price of the goods for which demand became relatively larger (i.e. the long-term supply curve is downward-sloping). Calibrations show that this channel explains most of the observed difference in quality-adjusted inflation rates across income groups. I find support for the external validity of these findings by using more aggregate data on the full consumption basket of American households back to 1953.

JEL codes: E31, I31, I32, O30, O31, O33

*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. I also thank Ufuk Akcigit, Adrien Auclert, David Autor, Matt Basilico, Shai Bernstein, Michael Best, Laura Blattner, Nick Bloom, Kirill Borusyak, Michael Boskin, Diego Comin, Gary Chamberlain, Varanya Chaubey, Gabriel Chodorow-Reich, John Coglianese, Tilman Dette, Mark Duggan, Max Eber, Josh Feng, Peter Ganong, Andy Garin, Matt Gentzkow, Robert Hall, Ben Hebert, Nathan Hendren, Leander Heldring, Erik Hurst, Simon Jaeger, Chad Jones, Max Kasy, Bill Kerr, David Koretz, David Laibson, Danial Lashkari, Jing Li, Hannes Mamberg, Leslie McGranahan, Sendhil Mullainathan, Emi Nakamura, Anna Paulson, Thomas Piketty, Dan Pollman, Benjamin Pugsley, Ananth Raman, Alexandra Roulet, Jann Spiess, Stefanie Stancheva, Jeremy Stein, Larry Summers, Colin Tomkins-Bergh, Chris Tonetti, Gabriel Unger, Heidi Williams, Alwyn Young and 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 for thoughtful discussions and comments. I thank Max Morales for excellent research assistance and I gratefully acknowledge financial support from the the HBS doctoral office, the Kauffman Foundation, the Stanford Institute for Economic Policy Research and the Washington Center for Equitable Growth.

1 Introduction

The rising level of nominal income inequality in the United States over the past thirty years has been a key topic of interest for academics and policymakers. The effect of skill-biased technical change in this process has been widely studied: innovations such as the diffusion of information and communications technologies have affected the relative price of skills and resulted in higher nominal income inequality. Much less attention has been paid to how price changes in the product market and the introduction of new products may differentially affect households at different points of the income distribution. Yet it is well-known that preferences are non-homothetic. Depending on their income level, households consume very different goods and services. Due to price changes in the product market over time, as well as changes in product variety, trends in real inequality may therefore differ from trends in nominal inequality. Product innovations may play a central role in this process 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 over the past ten years, a large sector accounting for over 25% of the US economy.

I investigate this question in two steps. First, in the *measurement* part of the paper, I show that in the retail sector over the past ten years the quality-adjusted price index of high-income households rose substantially *slower* than that of low-income households, which amplified inequality. This effect is large: real inequality in the retail sector increased 70% faster than nominal inequality during this period. To establish this, I build income-group-specific quality-adjusted price indices using detailed product-level data, in which I 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, in the *mechanism* part of the paper, 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 slower than that of the low-income. Specifically, my analysis shows that because demand from the high-income 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.[1] The retail sector is ideal to conduct this investigation because it accounts for a large share of US GDP, rich data is available, and the

[1] A particularly good example illustrating this idea 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.

notion of product (barcode) is well defined. I find support for the external validity of these findings, in other sectors beyond retail and over a longer time period, by using coarser data on the full consumption basket of American households back to the 1950s.

In the first part of the paper, I establish two new facts about inflation and increasing product variety across the household income distribution in the US retail sector from 2004 to 2013. I find that higher-income households experienced lower inflation and a faster increase in product variety than more modest households. The magnitude of these effects is large: over this 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 robust and hold before, during and after the Great Recession across product groups for a wide variety of price indices. They are based on detailed product-level data from the Nielsen Homescan Consumer Panel and Retail Scanner datasets, which are representative of the retail sector as a whole (which itself represents 40% of household expenditures on goods and 16% of household total expenditures). This analysis delivers a general methodological lesson for the measurement of inflation by statistical agencies: I show that the difference in inflation rates across income groups can be accurately measured only with product-level data. Indeed, 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[3] and suggest that trends in real inequality may be diverging from trends in nominal inequality. Collecting product-level data is key to accurately measure this divergence. This has important potential policy implications given the indexation of many government transfers.

In the second part of the paper, I examine whether the equilibrium response of supply to faster growth of demand from high-income consumers can explain the new facts on differential inflation and increase in product variety across the income distribution. It is a natural hypothesis 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, e.g. earning above $100,000 a year, has steadily increased - both because more and more households enter high-income brackets as the economy grows and because of rising inequality. I introduce a micro-founded model featuring monopolistic competition with variable elasticity of substitution preferences that differ across income groups, which generates a set of precise predictions that I take to the data and for which I find strong support. Intuitively, firms respond to changes in relative market size by skewing product introductions toward market segments that are growing faster. This process leads 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. In my data, product groups catering to higher-income households grow faster and have a higher rate of product introduction, as well as lower inflation on existing products. This

[2]As discussed in Section 4, 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.

[3]See Section 4 for a detailed discussion of how my results relate to the literature.

provides suggestive evidence in support of the theory but is not sufficient to establish causality from demand to supply.

To test the causal claim that increases in demand lead to a fall in inflation and an increase in product variety, I rely on two complementary identification strategies. First, I use shifts in the national age and income distribution between 2004 and 2013 to estimate the causal effect of changes in the number of consumers (market thickness) in a given part of the product space on inflation and product innovations. This research design is similar in spirit to Acemoglu and Linn (2004). Second, I introduce a novel research design exploiting variation in food stamp policy across US states between 2000 and 2007 to trace out the impact of changes in per capita spending on inflation and product innovation. As further discussed in Section 4, both research designs provide variation in demand plausibly exogenous to supply factors. They are complementary because the first is based on variation in the number of consumers and the second on variation in per capita spending, which in principle could have different effects on the equilibrium.[4]

Taken together, my results show that in response to growing demand the rate of introduction of new products increases and the equilibrium price of existing products falls, and that these effects are sufficiently strong to explain the divergence in price indices across income groups. According to my point estimates, a 1 percentage point increase in demand leads to a 10 basis point decline in inflation and a 35 basis point increase in spending on new products. In line with the model, the magnitude of the effect is similar regardless of whether the change in demand comes from a change in the number of consumers or in per capita spending. In simple calibrations based on historical changes in the US income distribution, I show that these effects are large enough to explain the new facts documented in the "measurement" part of the paper. In other words, these results suggest that absent the endogenous response of supply to market size effects, there would not have been a substantial difference in inflation nor in the rate of increase in product variety across income groups. This analysis has important implications for the endogenous growth literature, by providing evidence for endogenous product innovations across detailed product categories. It is also relevant for the trade literature and the debate on the role of markups in the gains from trade, because I present evidence that the gains from increased market size are largely due to a fall in markups (consistent with the model and variable elasticity of substitution preferences). More broadly, these results are relevant for policy, given that the effectiveness of any government transfer crucially depends on the equilibrium response of supply to market size.

Overall, this paper provides new evidence challenging the existing literature primarily in two respects. First, the literature suggests that households across the income distribution tend to experience similar inflation rates (e.g. McGranahan and Paulson, 2005), except during peculiar periods like the Great Recession (Argente and Lee, 2015). Second, theoretical work has focused on the "product cycle", the idea that innova-

[4]As discussed in Section 4, a wide class of models can generate the prediction that the quality-adjusted price falls in response to an increase in demand, because of the endogenous response of supply and general equilibrium effects. My preferred model relies on translog preferences with flexible preference parameters across income groups, which yields the prediction that changes in demand coming from changes in the number of consumers or from changes in per capita spending should lead to the same endogenous supply response. This prediction is borne out in the data, which rejects a broad class of monopolistic competition models using other preferences, for instance as in Zhelobodko et al. (2012).

tion 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 benefit everyone and, to a first-order approximation, they should lower all consumers' price index at approximately the same rate as they diffuse across the income distribution.[5] My findings suggest that market size effects and endogenously-increasing product variety may be a more important force than the product cycle, contributing to lower quality-adjusted inflation for higher-income households because market size grows faster at the top of the income distribution. More generally, this paper contributes to various strands of literature studying income inequality, price indices, technical change and monopolistic competition dynamics.[6]

The remainder of this paper is organized as follows: Section 2 describes the data; Section 3 presents the first main contribution of the paper, the measurement of quality-adjusted inflation and increasing product variety across the income distribution; Section 4 makes the second main contribution of the paper, establishing that increases in demand cause an increase in product variety and lower inflation on continuing products in equilibrium; Section 5 presents calibrations showing that historical changes in demand were large enough to cause the observed difference in quality-adjusted inflation across income groups through the endogenous response of supply; this section also provides evidence from the CPI and CEX data supporting the external validity of the findings in other sectors and in earlier time periods. A number of theoretical results, estimation details and robustness checks are reported in appendices.

2 Data Sources and Summary Statistics

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). With this data, I can track 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 16% of total household expenditures. Appendix B presents a detailed description of the data sources.

Three features of the data are particulalry useful for my analysis. First, product-level data is available

[5]Technically, because of the concavity of the utility function, the product cycle should result in lower quality-adjusted inflation for lower-income households.

[6]More precisely, this paper relates to at least seven strands of literature, which respectively examine nominal income inequality (Autor, Katz and Kruger (1998), Autor, Katz and Kearney (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), Comin, Lashkari and Mestieri (2015)), 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 (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)).

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. Intuitively, 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. The quantity and price data is also used to structurally estimate the relevant parameters of the demand system in Section 3. 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, the dataset offers a good measure of product innovations, defined as the introduction of new barcodes. Broda and Weinstein (2010) provide a detailed explanation regarding why it is reasonable to assume that all goods with different UPCs differ in some substantial way that might cause consumers to pay a different price for them and that it is rare for a meaningful quality change to occur that does not result in a change of UPC.[7] In other words, it is safe to assume that if the bar code changes, it is likely that some noticeable characteristic of the product has changed.[8] Similarly, it is possible to track products (barcodes) that are discontinued. Appendix Table 21 shows that creation is larger than destruction, i.e. new products tend to steal market shares from existing products. As will be discussed in Section 3, this displacement results in a bias in conventional price indexes that ignore the effects of changing quality.

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") are classified into 10 broad departments (dry grocery, general merchandise, health and beauty care, alcoholic beverages, deli, ...), 125 more detailed product groups (grooming aids, soup, beer, pet care, kitchen gadgets, ...) and 1,075 very detailed product modules (ricotta cheese, pet litter liners, bathroom scale, tomato puree, women's hair coloring, ...). When ranking product modules by mean consumer income, 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 filing, insecticide and frozen fruit drinks.

Finally, the data can be disaggregated at the level of 76 local markets, described in Appendix B. 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).

This high-quality scanner data thus makes it possible to measure at the most disaggregated level how demand, quantities and prices have evolved over time and across states for various income groups. Its main

[7] Broda and Weinstein (2010) make this point as follows: "Although it is difficult to enforce how a company uses a bar code, most industry experts strongly caution firms not to use the same bar code on more than one product. Doing so could cause confusion among retailers who would have trouble knowing what they were selling and for consumers whose receipts would not match their actual purchases. Similarly, firms typically do not use multiple UPCs for the same product because that makes it very difficult for retailers to reorder out of stock items. As a result, manufacturers tend to use other bar code systems for internal use and reserve the UPC for tracking products that are identical to the consumer. For example, changing the slogan on a Heinz ketchup bottle does not require a new bar code, but changing the size of the bottle does."

[8] Note 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 just because of a change in size or flavor: all of the results presented in the paper are similar when excluding these products from the definition of "new" products.

limitation is of course that it only covers the retail sector, which leaves the question of external validity entirely open: would similar patterns hold in the full consumption basket of American households?

Nonetheless, without any claim to external validity, this data is fruitful in two respects. First, it is ideal to make a series of methodological points about how to accurately measure income-group-specific inflation rates (see Section 3) and about how to distinguish between leading innovation models (see Section 4). Second, studying this data is intrinsically relevant because several US government transfers are indexed on the food CPI and the food-at-home CPI, such as the Supplemental Nutrition Assistance Program (also known as food stamps program) and the Child Nutrition Programs: the scanner data is perfectly suited to measure food-at-home inflation rates for various income groups, as well as the mechanisms that shape them.

Finally, by relying on economic theory and structural assumptions, as well as by augmenting the Nielsen data with additional price and spending data from BLS on sectors beyond retail, I am able to directly speak to the "external validity" question and offer estimates of income-group-specific inflation rates for the full consumption basket. Section 5 offers a detailed discussion of these issues.

2.1.2 Retailer Markup Data

To test specific predictions of the model in Section 4, 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 information 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.

2.1.3 Manufacturer Identifier Data

In order to measure manufacturer entry and competition, I have purchased 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 manufacturers 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. A second important feature of the data is that even the largest manufacturers are not close to being monopolists: the largest manufacturers in a product group on average has a market share of 22 percent. The model presented in Section 4 is consistent with these patterns.

2.1.4 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.[9]

2.2 Summary Statistics

Panel A of 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.

The product groups listed in Panel A 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.

[9] These 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.

Table 1: Summary Statistics

Panel A: 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, carbonated beverages, cereal, coffee, condiments, gravies, sauces, cookies, crackers, desserts, gelatins, syrup, flour, canned fruit, dried fruit, gum, james, jellies, juiced, canned juice, nuts, packaged milk, pasta, pet food, pickles, olive, prepared food, salad dressing, mayo, toppings, canned seafood, oil, snacks, non-carbonated soft drinks, soup, spices, seasoning, sugar, sweeteners, molasses, tea, canned vegetables, dried vegetables and grains	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, electric goods, photographic supplies, sewing notions, shoe care, soft goods, stationery, school supplies, sunglasses, toys and sporting goods	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

Table 1: Summary Statistics (*continued*)

Panel B: 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

A 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.[10] Moreover, the food industry has undergone a revolution in the past fifteen years with the rise of organic and natural food products, whose price relative to standard food products has been steadily decreasing.[11] The data is therefore ideal to study the dynamics of increasing product variety and how they may differentially benefit households across the income distribution.

Panel B of Table 1 compares aggregate spending share in the Nielsen scanner data compared 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).

3 Measuring Quality-Adjusted Inflation Across Income Groups

In this section, I compute quality-adjusted inflation rates across income groups, taking into account the welfare gains from increasing product variety. I start with a brief reminder on non-homothetic preferences and price indices. Second, I follow standard results in the literature to derive the exact price index in my preferred demand system. Third, I present results and robustness checks for inflation on existing goods across income groups. I show the relevance of these results for statistical agencies like BLS and discuss differences with the existing literature. Fourth, I document the difference in changes in product variety (due to both product creation and destruction) across income groups. Finally, I bring together the findings on inflation on overlapping products and on product creation and destruction to compute the full quality-adjusted inflation rate.

3.1 Nonhomothetic Preferences, Product Variety and Real 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. Figure 16 illustrates this idea. In this example, the "new" mapping is an upward shift of the "old" mapping (for instance because of productivity gains), but the shift is asymmetric and benefits

[10] Source: https://www.brewersassociation.org.

[11] For a detailed study by the US Department of Agriculture's Economic Research Service, see http://www.ers.usda.gov/amber-waves/2016-may/investigating-retail-price-premiums-for-organic-foods.aspx#.V3w-eesrLIW.

higher-income households relatively more. The shift takes into account changes in the quality-adjusted price of products as well as changes in the variety of available products for each nominal income level.

This paper 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 real 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.[12]

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 the previous year. Such price indices are known as "exact price indices." This requires taking into account changes in product quality, 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 into a welfare metric the patterns of product creation and destruction. In order 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 on the set of products available in two consecutive years. The quality of a given product is assumed to be constant over time[13] 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

[12]It is also useful to note that given that the set of goods is not fixed, the difference in the rates of quality-adjusted inflation experienced between high- and low-income households could be permanent. If the set of goods were fixed, the divergence in inflation rates between goods should be bounded and eventually converge to 0, otherwise in the long run all consumers, regardless of their income level, would switch to the goods with slower price increases. But since there is entry and exit, quality-adjusted inflation may be permanently lower for one income group relative to another (e.g. at any point in time the price of the products catering to the high income may remain higher than that of the products catering to the low income, but in a quality-adjusted sense the price of the high-end products may be very low).

[13]This assumption is standard in the literature: see for instance Feenstra (1994) and Broda and Weinstein (2010). It appears to be reasonable even though advertising or the introduction of complementary goods may violate it. It can be tested, and I present the results of this test in the robustness test section below.

the utility function play 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). In addition to the exact price index derived in the following subsection, I consider the following price indices:

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

Another relevant 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.[14] Accordingly, I present robustness checks with and without chaining the indices.

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 product should lead to a decrease in the price index. Translating the increase in

[14] For example, suppose all the prices in the last period return to their initial levels in period 0, which implies that they must have flucutated 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.

product variety into welfare gains requires structural assumptions. The key assumption is that within a product module, varieties are horizontally differentiated. I use two standard frameworks: nested CES utility (see next subsection) and nested translog utility (in progress). Importantly, these structural assumptions do not matter quantitatively because the elasticity of substitution of products within a module is very high, such that the gains from increasing product variety are already reflected in the prices of existing products (this is shown formally in the next subsection). The key point is that the patterns of product creation and destruction matter in general equilibrium, but their welfare effect is almost entirely taken into account in the price changes of products existing across periods. I also provide bounds showing that the patterns of product creation and destruction in the data will lower the price index more for high-income households than low-income households, even under some violations of the structural assumptions.

3.3 Estimation Framework for the Nested CES Exact Price Index

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 nonhomotheticicities.

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, following Feenstra . Product groups are indexed by g and G is the set of all product groups. The elasticity of susbtitution across product groups is ρ. The elasticity of substitution between product groups is $\sigma = \rho/(\rho - 1)$. The upper level utility function is:

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

Composite consumption within a product group is given by:

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

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

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

where c_{ubgt} is the consumption of UPC u in product module m and product group g in period t. $\sigma_m = \rho_m/(\rho_m - 1)$ between UPCs within product module m. d_{umgt} is unobserved and reflects the quality of the UPC. So we want to estimate σ and two high-dimensional sets of elasticities of substitution, $\{\sigma_g\}_g$ and $\{\sigma_m\}_m$. We expect $\sigma_m > \sigma_g$ since there is more substitution across UPCs within a module than across modules within a group.

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

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

The minimum cost function at the product group level is:

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

And the overall price index is given by:

$$P_t = [\sum_g P_{gt}^{\sigma}]^{\frac{1}{\sigma}}$$

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:

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

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.

If we make the assumption that product quality is constant over time ($d_{umgt} = d_{umgt-1}$) and if we ignore the introduction of new products, given our assumption of a (nested) CES utility function and the results in Sato (1976) and Vartia (1976), the exact price index is:

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

where $I_{mg} = I_{mgt} \bigcap I_{mgt-1}$ 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. These weights are computed using cost shares in the two periods and are always bounded between the shares of spending in the t and $t-1$ (in other words the price index is bounded between the geometric Paasche and Laspeyres indices described in the previous subsection):

$$s_{umgt} = \frac{p_{umgt}x_{umgt}}{\sum_{u \in I_{mg}} p_{umgt}x_{umgt}}$$

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

As shown in Broda and Weinstein (2010), with change in varieties across periods the exact price index (quality-adjusted inflation) for product module m within product group g is then given by:

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

with

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

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. 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:

$$\frac{\lambda_{mgt}}{\lambda_{mgt-1}} = \frac{1 + Growth\ Rate\ of\ Spending\ on\ Overlapping\ Products_{gmt}}{1 + Growth\ Rate\ of\ Total\ Spending_{gmt}}$$

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, price changes in the set of existing products already take into account the entry of more varieties.[15]

In principle, we could use the result above 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 makes sense to define periods in years to prevent seasonal factors from driving product turnover. Thus, UPCs will be considered destroyed only if they were not purchased at any time during a yearlong period. Second, we need 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. As mentioned earlier, a key assumption is that the taste or quality parameters for common goods must remain constant in start and end years of the sample. In fact, it may vary over short horizons due to anything that might affect demand (e.g., marketing or fashion considerations). The reason for why immutable preferences over long time horizons must be assumed when deriving price indexes is that if the utility function is changing over time for either exogenous reasons (e.g., fashion) or endogenous reasons (e.g., marketing) then one cannot make sensible statements about how price changes affect welfare, nor can one derive exact price indexes because identical price vectors will yield different utility levels at different times. The choice of the time horizon also matters for the magnitude of the adjustment term for increasing product variety.

Thus, we need data on quantity and price for new products, discontinued products, and products existing across periods, which is readily available in the Nielsen data. We also need to estimate the two high-dimensional sets of elasticities of substitution, $\{\sigma_g\}_g$ and $\{\sigma_m\}_m$. The main challenge for estimation is that

[15] One can better understand the implications of the choice of time horizon by considering an examples of how the proposition captures the impact of different types of creation and destruction, quoting from 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."

we want to obtain a demand and supply equation using only information on prices and quantities. The insight of Feenstra (1994) as extended by Broda and Weinstein (2006) is that although we cannot identify supply and demand, the data does tell us something about the joint distribution of supply and demand parameters. Appendix C gives details about how to derive the estimation equations.[16]

3.4 Inflation Across Income Groups For Products Available in Consecutive Years

3.4.1 Results

Figure 1: Inflation Across Income Groups (Overlapping Products)

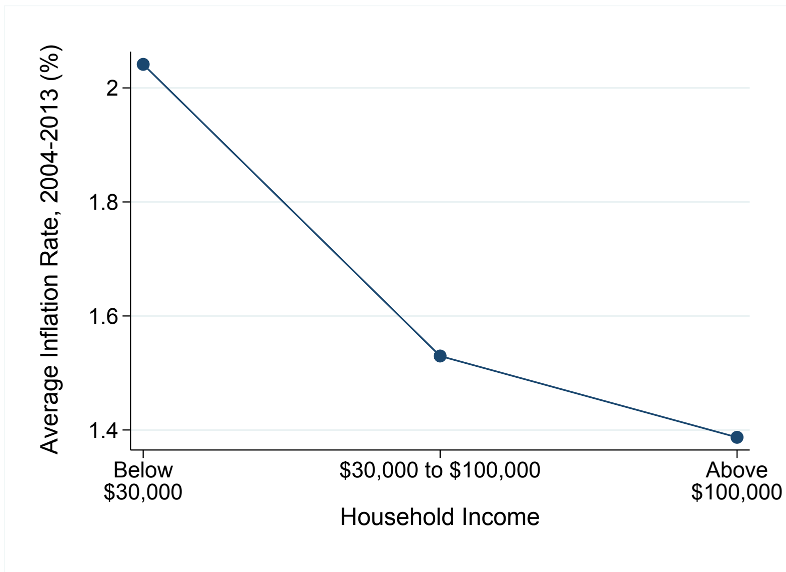

Figure 1 shows the average inflation between 2004 and 2013 on the set of overlapping products (defined as products that are available in consecutive years) for households across the income distribution. Inflation is computed using the exact price index for the nested CES utility function described in the previous subsection (without the adjustment for new and disappearing products, which is examined later in this section and does not affect the results). The inflation rate is about 0.65pp lower for households making more than $100,000 a year, relative to households making less than $30,000. As shown in Panel A of Figure 2, similar results are obtained when considering any of the price indices introduced in Subsection 3.2. In addition, Panel A of Figure 2 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 B 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,

[16] I have also estimated a demand system based on the translog expenditure function, following Feenstra and Weinstein (2015). The results are qualitatively similar to those presented here and are available from the author upon request.

relative to households making less than $30,000. Panel B of Figure 2 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 sytematically lower for the higher-income households.[17] Bootstrapped standard errors are very small, around 2 to 3 basis points, and are therefore omitted here.

Figure 2: Robustness of Inflation Difference between High- and Low-Income Households For Various Price Indices (Overlapping Products)

Panel A: For Various Price Indices

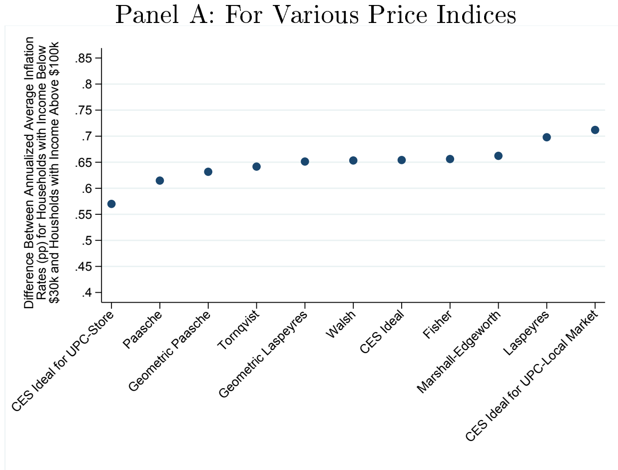

Panel B: Within Age Groups

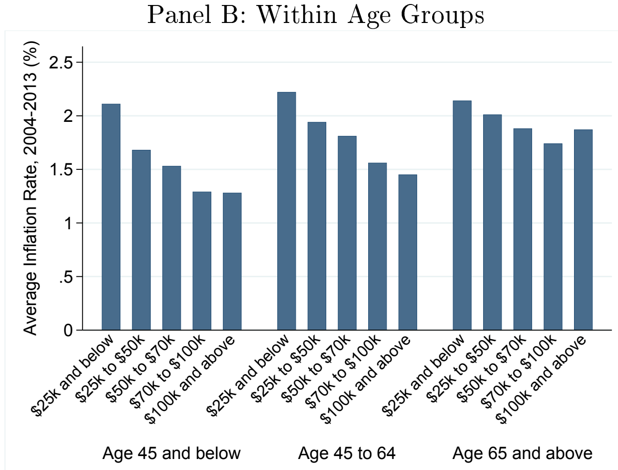

Table 2 shows the robustness of this result across subsamples. The difference between the inflation rates of high- and low-income households exists before, during and after the Great Recession,[18] and it is not driven by any single department. Appendix C presents various additional robustness tables and figures. First, Tables

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

[18]The difference inflation rates appears to be larger during the Great Recession. Argente and Lee (2015) 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.

14 and 15 describe the level of inflation for various cuts of the income distribution, various price indices and various periods. Figure 17 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 redefine products to be UPCs available in a given local market (Table 16) or UPCs available in a given store (Table 17) and show that the results continue to hold. Additional robustness checks are discussed at the end of this section.

Table 2: Robustness of Inflation Difference Across Income Groups (Overlapping Products) For Various Periods and Departments

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

3.4.2 Decompositions

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 \$100,00 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 bar codes, 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 and experienced the same inflation rate within an item category. The "within" component corresponds to the inflation difference 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, we can decompose the inflation difference between high- and low-income households as follows:

$$\pi^R - \pi^P \equiv \sum_G s_G^R \pi_G^R - \sum_G s_G^P \pi_G^P = \underbrace{\left(\sum_G s_G^R \pi_G - \sum_G s_G^P \pi_G \right)}_{Between} + \underbrace{\sum_G \overline{s_G} (\pi_G^R - \pi_G^P)}_{Within}$$

with s_G^i the share of spending of income group i on product grouping G andπ_G^i the inflation experienced by income group i in product grouping G. π_G and $\overline{s_G}$ denote the average inflation rate and the average spending shares for product grouping G.

Table 3 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 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 from spending patterns across the 1,025 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%.[19]

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. In Section 4, I examine whether local competition and changes in markups can explain these patterns.

[19] Note 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.

Table 3: Decomposition of the Inflation Difference Between High- and Low-Income Households

Aggregation Level (Broad to Narrow)	Decomposition	Inflation Difference pp	% of actual
Department	Between	0.06	8.6
Product Group	Between	0.14	21.4
Product Module	Between	0.28	42.8
UPC	Between	0.476	72.2
UPC-Local Market	Between	0.520	78.8
UPC-Store	Between	0.607	92.1

3.4.3 Relevance for the Methodology of Statistical Agencies

Table 3 means 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 reweight 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 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 have directly shown that the magnitude of this bias is large in the retail sector. Table 5 in Section 4 shows that a large share of the inflation difference across income groups could be captured by segmenting each of the detailed item categories 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 and beyond.[20]

[20] One would then need to infer the spending shares of various income groups along price deciles, which could be done for

3.4.4 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 extends well beyond the crisis and continues to hold even when substitution effects are ignored (indeed, Figure 2 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). In Section 4, I show that the magnitude of the inflation difference between high- and low-income households can be explained by the equilibrium response of supply to market size effects.[21]

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 evidence of differences in inflation rates across income groups and, in Section 4, I focus on a completely different explanatory mechanism.

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 as in Section 4.3.2. and to finally to relate these patterns to the dynamics of product creation and endogenous changes in markups, which are discussed in the remainder of the paper. My analysis shows that collecting product-level data is key to accurately measure the divergence of inflation rates across income groups - this methodological lesson is likely to apply to other sectors beyond retail.

3.5 Changes in Product Variety Across Income Groups

3.5.1 Results

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. Figure 3 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. Each dot on

instance by estimating "quality Engel curves" as in Bils and Klenow (2001).

[21] Note that both my results and the results of Argente and Lee (2015) 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.

the figure represents 5% of the data, which provides a non-parametric approximation to the conditional expectation function. For every $10,000 increase in the mean income of the consumers buying from a product module, the share of spending 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 product, 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 C Figure 19 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 C Figure 19 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 $\frac{\lambda_{mgt}}{\lambda_{mgt-1}} = \frac{1+Growth\ Rate\ of\ Spending\ on\ Overlapping\ Products_{gmt}}{1+Growth\ Rate\ of\ Total\ Spending_{gmt}}$ decreases with consumer income, which confirms that higher-income consumers benefit more from product innovations.

Figure 3: Product Variety Increases Faster In Product Modules Catering to Higher-Income Households

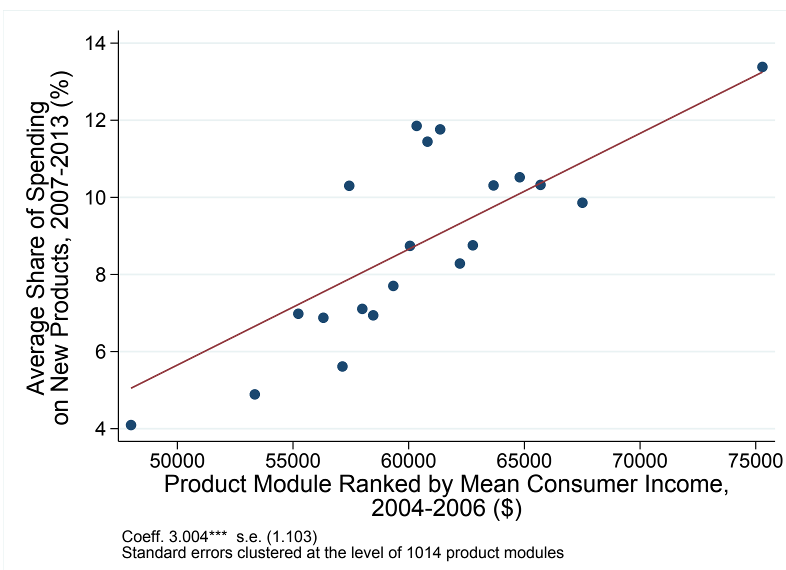

Similar results hold for other measures of "new products" - new UPCs relative to two, three or four years ago, as well as new brands. Moreover, I have examined patterns of product creation and destruction *within* product modules, which are similar to the patterns observed across modules: higher income households tend to benefit more from increasing product variety. The patterns within product modules are discussed in greater depth at the beginning of Section 4, where I show the role of the quality ladder: there are more product introductions at the top of the quality distribution within modules, where there are relatively more high-income consumers.

3.5.2 Decompositions

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:

$$SSNP^R - SSNP^P \equiv \sum_G s_G^R SSNP_G^R - \sum_G s_G^P SSNP_G^P = \underbrace{\left(\sum_G s_G^R SSNP_G^R - \sum_G s_G^P SSNP_G^P\right)}_{Between} + \underbrace{\sum_G \overline{s_G}(\pi_G^R - \pi_G^P)}_{Within}$$

with s_G^i the share of spending of income group i on product grouping G and $SSNP_G^i$ the share of spending on new products for income group i in product grouping G. $SSNP_G$ and $\overline{s_G}$ denote the average share of spending on new products and the average spending shares for product grouping G. Table 4 shows that the difference beween 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 Table 3 and provides preliminary evidence that there is a tight connection between the inflation and product innovation patterns, which is further examined in Section 4.

Table 4: Decomposing the Difference in Shares of Spending on New Products between High- and Low-Income Households

Aggregation Level (Broad to Narrow)	Decomposition	Difference in Share of Spending on New Products (% of actual)
Department	Between	1.8
Product Group	Between	29.0
Product Module	Between	39.2

3.6 Quality-Adjusted Inflation Across Income Groups

Using the results from equation 2, I can bring together the previous facts on inflation for products available in consecutive years and on product creation and destruction. I find that between 2004 and 2013, on average annual quality-adjusted inflation was 66 basis point lower for households earning above $100,000 a year, relative to households earning below $30,000 a year.

Appendix C Table 23 shows the distribution of the estimated elasticities of substitution by income group. Two findings stand out. First, the elasticities tend to be slightly smaller for higher-income households, i.e. higher-income households are less price elastic in equilibrium.[22] This provides direct evidence of non-homotheticities (in addition, Figure 24 in Appendix C shows that the elasticites for high- and low-income consumers are not very strongly correlated). Second, the magnitude of the elasticities is very high. Using the optimal markup formula derived in Section 4, these magnitudes are consistent with the observed markups in the retail sector.[23] The high values of the elasticities mean that the “product variety” adjustment is very small: since the elasticities of substitution are very high, most of the welfare effects are captured by the inflation difference on goods that exist across consecutive years.[24] As a result, quality-adjusted inflation across income groups, reported in Appendix C Figure 23, looks virtually identical to inflation across income groups for overlapping products. Due to the high elasticities of substitution with 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.

The difference in annual quality-adjusted inflation between high- and low-income households is therefore very large, especially compared to the rise in nominal income inequality. According to IPUMS Census data, between 2004 and 2013, on average annual wage growth was 93 basis point faster for households earnings above $100,000, relative to households earning below $30,000. Therefore, in the retail sector price dynamics magnified inequality by $\frac{66}{93} = 70.9\%$.

The 66 basis point inflation difference between high- and low-income households I have measured in the retail sector between 2004 and 2013 is important for three reasons. First, my analysis delivered a methodological lesson which is likely to apply in other sectors: product-level data is needed to accurately measure the inflation difference between income groups. In particular, one cannot hope to learn much from the standard data sources produced and used by the Bureau of Labor Statistics, the CPI price series and CEX expenditure patterns. Yet these data are currently the basis for the indexation of government transfers. Second, the inflation difference across income groups measured in the Nielsen data, in particular in the food categories, is intrinsically relevant because several US government transfers are indexed on the food CPI, or food-at-home CPI, such as the Supplemental Nutrition Assistance Program (also known as food stamps program) and the Child Nutrition Programs. Third, based on the Nielsen data and using economic theory, 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. I carry out this exercise

[22]Note that the equilibrium elasticity of susbtitution depends on consumers' preference parameters, but also on the competitive environment if the elasticity of substitution is not constant. See Section 3 for a discussion of models of monopolistic competition with variable elasticities of substitution. I have estimated a demand system based on the translog expenditure function, which features decreasing elasticities of substitution, and have obtained qualitatively similar results.

[23]In retail (groceries and food) the margin is around 2.71%.
See http://pages.stern.nyu.edu/~adamodar/New_Home_Page/datafile/margin.html

[24]Indeed, from the derivation in subsection 3.2, quality adjusted inflation is given by $\pi_{mg} \equiv P_{mg} \cdot \left(\frac{\lambda_{mgt}}{\lambda_{mgt-1}}\right)^{\frac{1}{\sigma_m - 1}} \to P_{mg}$ as $\sigma_m \to \infty$

in Section 5.

3.7 Further Robustness Checks

Selection effects. A potential concern is that the inflation patterns described above could result from selection effects, for instance if 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.[25] Tables 19 to 21 show that such selection effects are in fact not at play in the data.[26]

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 3 and 8. 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 resticted 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 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 was an important force in the data, under the assumption described at the beginning of this section the nested CES demand system will 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.[27] 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

[25] See Pakes and Erickson (2011) for a discussion of such selection effects.

[26] Note that even if selection effects were at play, the nested CES structural demand system with new goods addresses these concerns by adjusting the price index when new varieties enter.

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

change in the quality-adjusted price. I address the concern that high-income households may be more like to be 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 product.

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 as 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 3 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 C Table 22 reports the results of such a regression and show 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 in price levels in any given year.[28] Figure 18 in Appendix C 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 B).

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-

[28]Note 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.

income households between 1999 and 2004.

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

Quarterly data. Table 18 shows that the results are very similar when repeating the analysis at a quarterly frequency.

4 The Equilibrium Response of Supply to Changes in Demand

In this section, I investigate the hypothesis that an important driver of the results presented in Section 3 - namely, the fact that higher-income household experience lower inflation and a faster increase in product variety from 2004 to 2013 - is differential income growth and directed product innovations. I first present a series of novel stylized facts showing that it is a very natural hypothesis to investigate. I then develop a theory showing how rising demand from high-income consumers may cause a shift in the direction of product innovations, and result in lower quality-adjusted inflation for the high-income. Intuitively, an increase in market size leads to more product entry, which puts downward pressure on the prices of existing products through increased competition. Next, I test the key channel of this theory by estimating the causal impact of a demand shock on price and innovation dynamics, using two complementary research designs providing variation in the number of consumers and in spending per capita across the product space, respectively. Finally, I present additional evidence allowing me to distinguish between different innovation models.

4.1 Preliminary Evidence

Four patterns in the data make it very natural to focus on the equilibrium reponse of supply - by which I primarily mean manufacturers - to changes in demand as a key mechanism. First, the price and product variety patterns *within* modules are closely linked, through the product quality distribution: there is more entry and lower inflation for higher-quality products. Second, I show that the inflation and product variety patterns also go hand in hand *across* product modules - the relationship is very strong, which pleads for a joint theory of price and product innovation dynamics. Third, I show that supply factors appear to be driving the patterns of increasing product variety. And fourth, I establish that retailer- and store-level dynamics are not driving these effects. Thus, I establish *ex-ante* in a simple data-driven way that it is natural to develop a theory of how manufacturers cater to the national market by directing their product innovations at growing segments of that market.

4.1.1 The Role of the Quality Ladder

Within product modules, the quality ladder plays an important role for the patterns of inflation and product introductions. Figure 4 documents that within product modules, there is more product entry (Panel A) and lower inflation (Panel B) 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, and it is not subject to mean reversion because the price of the UPC in both the start and the end periods are used to classify the UPC across price deciles. Prices are adjusted for the weight of the item in order to provide a more accurate measure of the unit price.

Appendix D Figure 25 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 4, which confirms that mean reversion is not driving the results.

Figure 4: Inflation and New Products across the Product Quality Distribution, within Product Modules

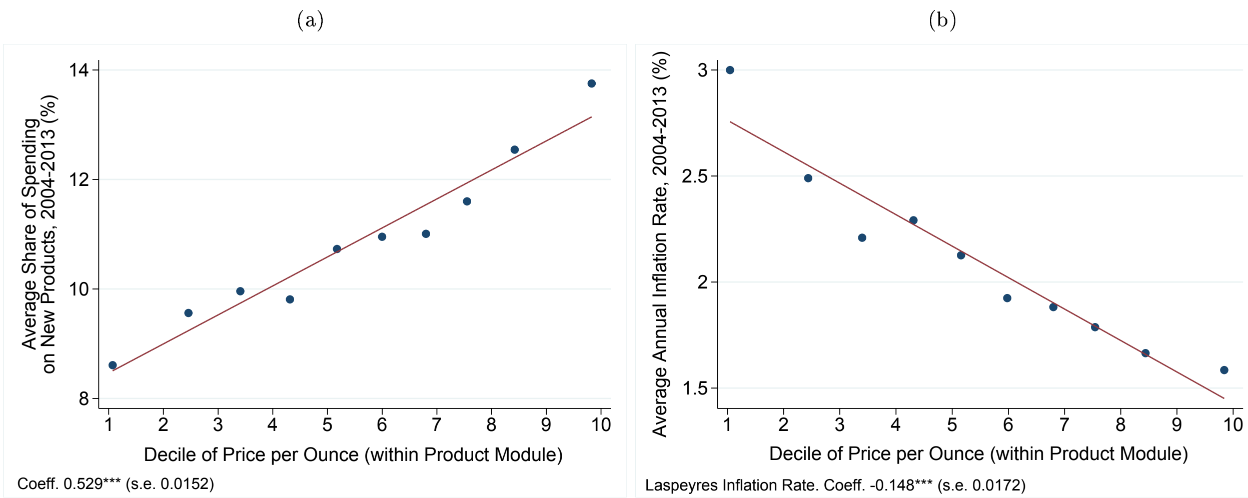

Table 5 shows that differences in the spending patterns of high- and low-income households across price deciles within product modules explain more than 85% of the inflation difference between high- and low-income households that exists across UPCs. In other word, the decomposition shows that the inflation 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.

Table 5: Decomposition of the Inflation Difference Between High- and Low-Income Households Relative to Across-UPC Benchmark

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

4.1.2 The Need for a Joint Theory of Inflation and Increase in Product Variety

The negative correlation between inflation and the share of spending on new products is also a key feature of the data *across* product modules. Figure 5 shows this relationship across product modules.

Figure 5: The Negative Relationship Between Inflation and Share of Spending on New Products across Product Modules

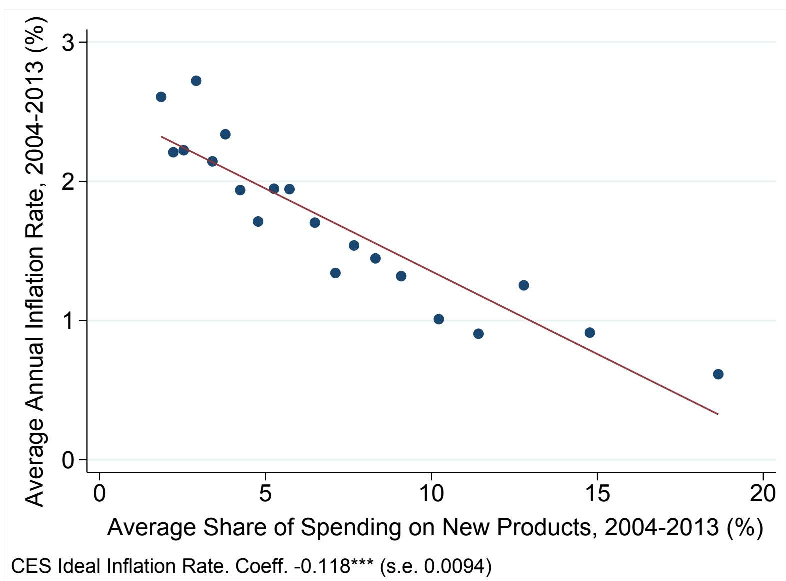

A simple decomposition exercise shows that the relationship between inflation and product innovations across modules can explain a large fraction of the inflation patterns across income groups documented in Section 3.[29] As previously mentioned, for any product grouping G, we can write the inflation difference

[29]This is similar in spirit to the reweighting technique introduced in DiNardo, Fortin and Lemieux (1996).

between income groups as:

$$\pi^R - \pi^P \equiv \sum_G s_G^R \pi_G^R - \sum s_G^P \pi_G^P = \underbrace{\left(\sum_G s_G^R \pi_G - \sum s_G^P \pi_G\right)}_{Between} + \underbrace{\sum_G \overline{s_G}(\pi_G^R - \pi_G^P)}_{Within}$$

with s_m^i the share of spending of income group i on product grouping G,π_G^i the inflation experienced by income group i in product grouping G, and with π_G and $\overline{s_G}$ denoting the average inflation rate and the average spending shares for product grouping G. We can now decompose the "between" component further and examine how much of the inflation difference across categories is explained by (or predicted by) differences in shares of spending on new products across categories:

$$\left(\sum_G s_G^R \pi_G - \sum s_G^P \pi_G\right) = \left(\hat{\pi_G^R} - \hat{\pi_G^P}\right) + R$$

with

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

where X_G is share of spending on new products in G. $\hat{\beta}$ is the OLS estimate of β. This procedure calibrates 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 5), 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 5). Table 6 shows that for the various levels of aggregation, around half of the inflation difference between high- and low-income households can be explained by differences in patterns of product innovations.[30] These results provide strong support for the idea that low inflation and high product introductions go hand in hand.

Table 6: How Much of the the Difference in Inflation Between High and Poor is Explained by Patterns of Product Innovations?

Aggregation Level (Broad to Narrow)	Share of Rich-Poor Inflation Difference Explained
Department	40.9
Product Group	58.3
Product Module	51.3

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

The relationships described so far between inflation and new products, both within and across modules, are only correlations and should not be interpreted as causal. But they provide transparent evidence on the pervasive nature of the relationship between inflation and product innovations and on its relevance for understanding changes in real inequality.

4.1.3 The Role of Supply Effects

Do the product variety patterns across income groups come from supply or demand? As shown on Figure 3, 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:

$$ShareSpendingNewProducts_{HM}$$

$$= \alpha + \beta ProductModuleIncomeRank_M + \alpha_H + \epsilon_{HM}$$

where α_H is a household fixed effect and $ProductModuleIncomeRank_M$ 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 7, 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 3. 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 7: New Products Target Higher-Income Consumers

	$ShareSpendingNewProducts_{hm}$	
$ProductModuleIncomeRank_M$	2.79***	2.82***
	(1.024)	(1.031)
$ProductModuleIncomeRank_M \times HighIncome_H$		-0.24***
		(0.063)
$ProductModuleIncomeRank_M \times LowIncome_H$		0.11*
		(0.058)
$Household\ Fixed\ Effects$	Yes	Yes

Standard errors clustered by product modules.

4.1.4 Retailers vs. Manufacturers

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 8 shows 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 store. 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 8: 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

4.2 Theory

4.2.1 Intuition: Tracing Out the Observed Long-Term Supply Curve

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 real inequality. Consider Figure 6. When relative demand goes up, if the short-run supply curve is upward sloping as in 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, which will at least mitigate the price increase and, as illustrated in Figure 6, could potentially result in a new equilibrium price that is lower than the initial equilibrium price. This "price overshooting" case proves to be relevant empirically, as shown in the rest of this section. In other words, the observed long-term supply curve is downward sloping.[31]

Figure 6: Does the Price Fall When Demand Rises?

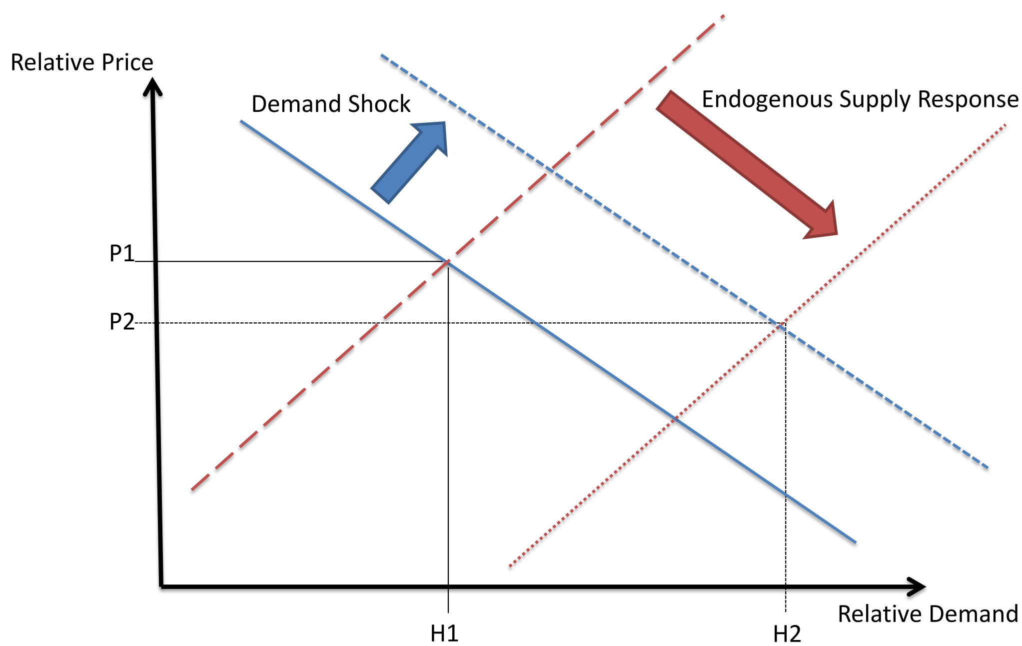

To investigate whether changes in nominal inequality overstate or understate changes in real inequality, the following concepts are useful:

- **Weak equilibrium (relative) bias** ("directed technical change"): when demand for a good becomes

[31] The observed long-term supply curve is defined as the nexus of equilibrium points traced out by shifts in the demand curve.

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

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

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

The equations above and 6 provide an intuitive reduced-form way of thinking about the effect of shifts in demand on the equilibrium price. In the next subsection, I discuss a specific microfounded model that generates more precise predictions.

4.2.2 A Simple Microfoundation

I focus on microfounded models of monopolistic competition with free entry of products. This broad class of models is appealing for three reasons: the assumption of monopolistic competition is reasonable in retail (the Herfindahl index for most product groups is below 0.20), these models nest the standard model of directed technical change (Acemoglu, 2002), and they generate rich product-level predictions and counterfactuals. 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.

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 6. 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:

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

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 $V(.)$ is 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 clearing:[32]

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

Therefore, the model delivers the following comparative statics:

$$\frac{dN^*}{dL} > 0 \qquad \frac{dx_i^*}{dL} < 0 \qquad \frac{dM_i^*}{dL} \lesseqgtr 0$$

Therefore, the optimal markup is given by the inverse of the price elasticity of demand.[33] This result is very general and holds regardless of the shape of the cost function $V(.)$. 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. Consequently, consumption per capita x_i for the existing products goes down, which induces a responses of the optimal markup M^*. The equilibrium markup may increase, decrease or stay unchanged, depending on the properties of demand. Figure 7 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.

[32] A 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).

[33] This term 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.

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.[34] I test and find support for these predictions in the rest of this section.

Because preferences are nonhomothetic, the curvature of the utility function may differ for consumers in different income groups and the equilibrium response of price to market size may be different in product modules catering to different consumer segments. This framework allows for rich counterfactuals to answer the question: what would have been the difference in inflation and rates of product introduction across income groups absent the endogenous response of supply to market size effects? The framework is based on homothetic utility functions within product modules, but as previously discussed I separately estimate these utility functions for different groups of consumers across the income distribution, which effectively allows for nonhomotheticities by letting the parameters of the utility function vary freely with the level of income.[35]

Figure 7: The Equilibrium Response of Price to Changes in Market Size Depends on the Price Elasticity of Demand

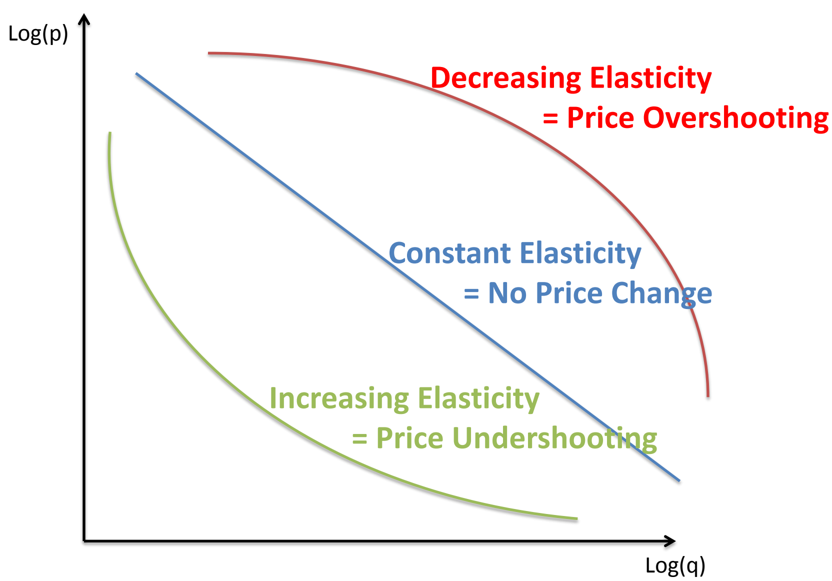

[34]Note that this speaks to an active debate in the trade literature about the source of the gains from trade and the role of variable markups and variable elasticity of substitution preferences. See in particular DeLoecker, Goldberg, Pavcnik and Khandelwal (2012), Feenstra and Weinstein (2016), and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

[35]Section 3 shows how to nest the various sub-utility functions for each product module into one aggregate utility function.

4.2.3 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 model 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 entry[36] and the price of continuing products decreases either because of a decrease in marginal cost[37] or because of a fall in markups.[38]

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, 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.

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 *market thickness* (the number of consumers) or from a change in *spending per capita,* respectively denoted L and E in the model introduced above. Depending on

[36]Some recent work in the trade tradition models entry of products within multi-product firms, e.g. Mayer, Melitz and Ottaviano, 2016.

[37]In models in the macro tradition, the fall in marginal cost can either be exogenous, in models with increasing returns to scale (e.g. Matsuyama, 2002), or endogenous, in models with endogenous investment in marginal cost improvements, where the returns increase with market size (e.g. Acemoglu and Linn, 2004).

[38]In 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.

the model, variation in the number of consumers and variation in per capita spending could have different effects on the equilibrium.[39] 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. In this draft, this model is reported at the end of Appendix A.

4.3 The Causal Effect of Changes in Market Thickness on Product Innovations and Inflation

4.3.1 Motivating Evidence and Identification Challenge

The key causal channel in the model is that growing demand causes more product entry, and in turn lower prices on existing products due to a fall in markups. Figure 8 shows that product modules catering to higher-income households indeed have both higher growth and lower inflation. Moreover, Table 7 provided early evidence that supply factors play an important role in differential product introductions across the product space.

However, these facts alone do not establish that the endogenous response of supply to demand is a key channel. The equilibrium relationship between price and quantity across product modules does not identify the causal effect of demand, because of reverse causality (demand might be following supply) 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). In the remainder of this section, I build a predictor of (potential) demand that is plausibly orthogonal to supply factors. Specifically, I consider changes in market size across product modules over time at the national level driven by changes in the age and income distributions. In robusness checks, I use variation in market size both over time and across local markets within the US.

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

Figure 8: Product Modules Catering to Higher Income Households Have Faster Growth and Lower Inflation

(a) Growth Across Product Modules (b) Inflation Across product Modules

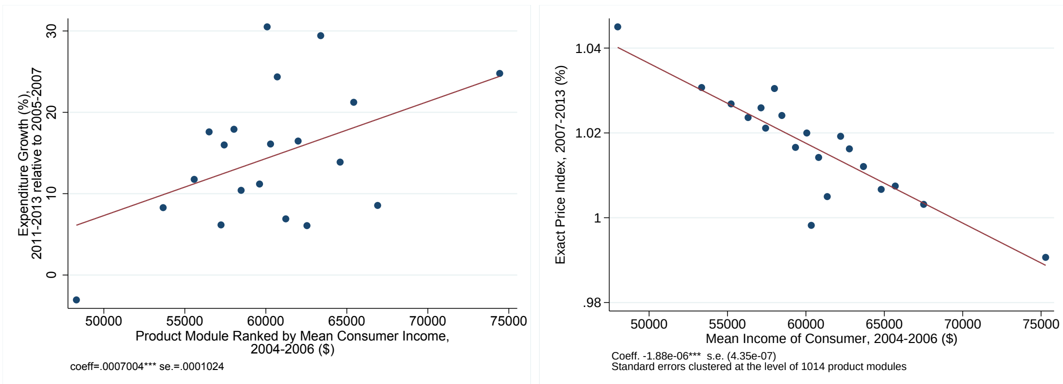

4.3.2 Research Design

A major difficulty in any investigation of the impact of market size on innovation is the endogeneity of market size: better products will have larger markets. A strategy to overcome this problem is to exploit variations in market size driven by US demographic changes, which should be exogenous to other, for example scientific, determinants of innovation and entry of new products. To estimate potential market size, I construct age-income profiles of users for each product module $\times$ price decile, and then compute the implied market size from aggregate demographic changes given these (time invariant) income-age profiles. This identification strategy is similar to Acemoglu Linn (2004). Using this strategy, 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 very different nature. More importantly, this paper is the first to examine the causal effect of changes in market size on the price of existing products, as well as on the aggregate price taking into account the welfare gains from increased product variety. I find that prices go down when demand goes up, i.e. the observed supply curve is downward sloping.

The predictor of market size is built as follows. At the beginning of the sample (2004-2006), I compute per capita expenditures $E_{MG}^{T_0}$ in product module $\times$ price decile M for fifteen age-income groups G I consider.[40] Then, I predict (potential) demand at time t as:

$$D_{Mt} = \sum_G E_{MG}^{T_0} P_{Gt}$$

Thus, the spending profiles are kept constant and the variation in predicted demand comes entirely from changes in age-income group size P_{Gt}. To implement this design, I compute growth of demand based on the

[40] Specifically, I consider the interaction of three age groups - below 45, between 45 and 65, and above 65 - and five income groups - annual household income below 25k, 25k to 45k, 45k to 60k, 60k to 100k, and above 100k.

change in the size of each age-income group in 2011-2013 relative to 2004-2006.

The identification assumption is that the direct effect of changes in age-income group size on the equilibrium price was only through demand. For instance, if the 20-year old were better at creating new products targeting people in the same age group, the identification assumption would be violated. As a robustness check, I repeat the analysis for older households, who are closer to retirement age. I find similar point estimates, which suggests that direct supply effects are not driving the results.

Figure 9: Changes in Market Size from Changes in the Age Distribution

(a) Changes in Age Distribution, 2005 to 2013

(b) Example of a Product with a Steep Age Profile, Diapers

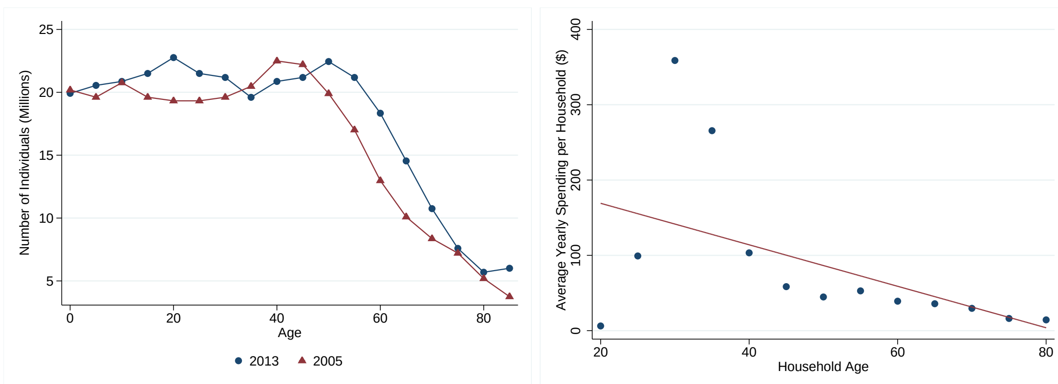

4.3.3 Results

I first present the results with a series of binned scatter plots, where each dot represents 10% of the data. I then show the results in a regression table, with standard errors clustered by product module. Figures 10 and 11 below 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. These results lend strong support to the hypothesis that supply endogenously responds to changes in market size.

Table 9 shows that the relationships between predicted market size growth, product innovations and inflation are significant at the 1% level. The interpretation of the magnitudes is as follows: a one percentage point increase in the growth of demand[41] causes a 0.35 percentage point increase in the share of spending on new products and a 0.11 percentage point decline in the inflation rate on goods that are available across years. Figure 26 in Appendix D shows the relationship between predicted and actual spending growth, which is also strong. Column 3 of 9 confirms that the relationship between predicted and actual growth of total spending is significant at the 5% level.[42] The point estimates are precisely estimated and are used for a

[41] Where growth of demand is measured as the predicted growth in total spending in a product module - price decile, given changes in the age and income distributions.

[42] The 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.

calibration reported in Section 5.

Figure 10: Higher Market Size Leads to Product Entry

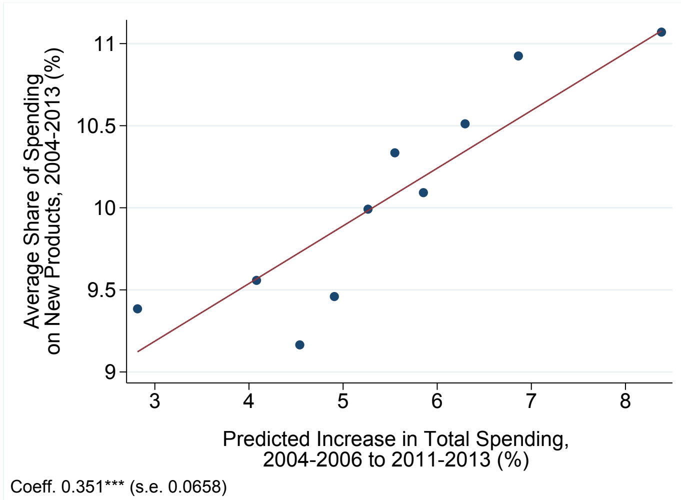

Figure 11: Higher Market Size Leads to Lower Inflation (Overlapping Goods)

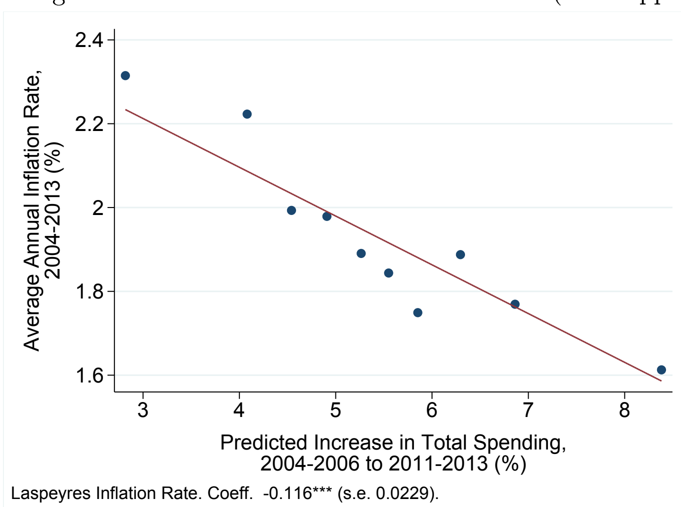

Table 9: Causal Effects of Changes in Market Size

	Share of Spending on New Products (pp)	Overlapping Goods Inflation Rate (pp)	Actual Spending Growth (%)
Predicted Increase in Spending (%)	0.351***	-0.116***	1.031**
	(0.0658)	(0.0229)	(0.492)
Product Module Fixed Effects	Yes	Yes	Yes
Spending Weights	Yes	Yes	Yes
Sample Restricted to Positive Spending Growth	Yes	Yes	Yes
Number of Observations	9,089	9,089	9,089
Number of Clusters	1,006	1,006	1,006

Standard errors clustered by product modules.

Table 10 shows the robustness of these results. Panel A runs a falsification test in the set of product modules - price deciles that experienced negative spending growth during the period 2004-2013. The model does not predict a significant relationship between change in market size and entry of new products or inflation in this subsample, and indeed I do not find any. Panel B addresses the potential concern that some of the relationship between predicted demand and innovation and inflation could be spuriously driven by a differential increase in supply across the product space. It shows that the results are very similar when considering product module - price deciles that cater to consumers above the age of fifty. In other words, the result is not driven by young consumers, for whom direct supply effects are more likely to exist.

A variety of additional robustness checks are reported in Appendix D. Panel A of Appendix D Table 24 shows that the points estimates are very stable when including flexible controls for the initial (2004-2006) age and income distributions in each product module - price decile. Specifically, I control (linearly) for the 10th, 25th, 50th, 75th and 90th percentiles, as well as the mean, of both the age and income distributions. Panel B also shows stability of the point estimates when introducing fixed effect for each price decile within a product module and when omitting product module fixed effects. Panel B of Appendix D Table 24 shows that the results are similar when using truncated weights. Statistical significance at the 1% level is maintained in all specifications, with standard errors clustered by product modules.

In Appendix D, as an additional robustness check I use time variation in the household age and income distributions in seventy-six local markets tracked by Nielsen within the US between 2004 and 2013. I compare inflation patterns across product module - local market cells with increasing or decreasing predicted market size. I again find that inflation is lower when predicted demand increases - the point estimates are very similar to those obtained from the analysis at the national level and robust to the inclusion of various fixed effects.

Table 10: Robustness of Causal Effects of Changes in Market Size

Panel A: Falsification Test in Product Module - Deciles with Negative Spending Growth

	Share of Spending on New Products (pp)	Overlapping Goods Inflation Rate (pp)
Predicted Increase in Spending (%)	-1.093	0.162
	(1.148)	(0.108)
Product Module Fixed Effects	Yes	Yes
Spending Weights	Yes	Yes
Sample Restricted to Negative Spending Growth	Yes	Yes
Number of Observations	632	632
Number of Clusters	305	305

Standard errors clustered by product modules.

Panel B: The Effect Is Not Driven by Young Consumers

	Share of Spending on New Products (pp)	Overlapping Goods Inflation Rate (pp)
Predicted Increase in Spending (%)	0.306***	-0.113***
	(0.075)	(0.021)
Product Module Fixed Effects	Yes	Yes
Spending Weights	Yes	Yes
Sample Restricted to Positive Spending Growth	Yes	Yes
Number of Observations	6,571	6,571
Number of Clusters	926	926

Sample restricted to product modules - price deciles with mean consumer age above 50.
Standard errors clustered by product modules.

4.3.4 Further Evidence to Distinguish between Models

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 change in spending using the observed change in the size of the various age-income groups in each state. I find that the fall in inflation is entirely predicted by the *increase* in spending, rather than by the initial level of spending.

Table 11: Lower Inflation is Caused by Increases in Market Size

	Overlapping Goods Inflation Rate (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 FE	Yes
Weights	Yes

Standard errors clustered by product modules

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 product. A first-order Taylor expansion yields a convenient additive expression for the log price change:

$$p_{it} = m_{it} + c_{it}$$
$$\Delta^t log(p_{it}) \approx \Delta^t log(c_{it}) + \Delta^t \frac{m_{it}}{c_{it}}$$

I can then run the following regression across product modules, with store-year fixed effects to absorb rent and labor costs:

$$\Delta^t log(p_{it}) = \beta I_i + \lambda_{st} + \epsilon_{it}$$
$$\Delta^t log(c_{it}) = \widetilde{\beta} I_i + \widetilde{\lambda_{st}} + \widetilde{\epsilon_{it}}$$
$$\Delta^t \frac{m_{it}}{c_{it}} = \bar{\beta} I_i + \bar{\lambda_{st}} + \bar{\epsilon_{it}}$$

with I_i income rank of module. Note that $\beta \approx \widetilde{\beta} + \bar{\beta}$. We can use this relationship[43] to answer the following question: do prices rise more slowly for high-income consumers because retailer margins decline more quickly or because wholesale costs rise more slowly?

As shown in Figure 12 and Table 12, changes in retailer margins account for about half of the differential inflation between high- and low-income households. This number can be thought of as a lower bound[44] on the total share of changes in markups in the overall inflation difference, because wholesalers (and in turn manufacturers) themselves have a markup. I have checked that this relationship is robust across years and when using other specifications. These results provide support for the prediction of the model that variable markups are a key channel.[45]

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

[44] See Appendix A for a formal double marginalization model making that point.

[45] Variable markups are often studied in the macro literature in the context of short-run business cycle fluctuations. The fact that markups explain a large share of the difference in inflation between high- and low-income households does not mean that

Figure 12: Changes in Wholesale Costs vs. Changes in Retailer Margins

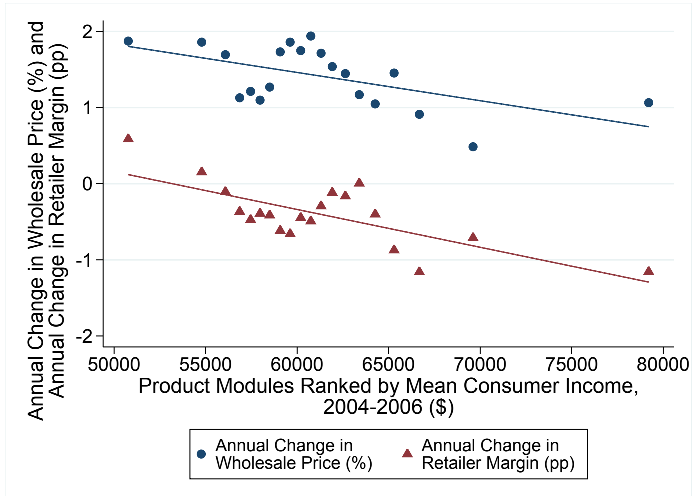

Table 12: Changes in Wholesale Costs vs. Changes in Retailer Margins

	Log Price Change	Log Wholesale Cost Change	Retailer Margin Change (pp)
$ProductModuleIncomeRank_M$	-0.777***	-0.341***	-0.448***
	(0.188)	(0.103)	(0.212)
Spending Weights	Yes	Yes	Yes
Store-Year Fixed Effects	Yes	Yes	Yes
Number of Observations	6,002,235	6,002,235	6,002,235
Number of Clusters	628	628	628

Standard errors clustered by product modules.

4.4 The Causal Effect of Changes in Per Capita Spending on Product Innovations and Inflation

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.

4.4.1 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 from low-income consumers on food. Between 2001

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.

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 for the program. Ganong and Liebman (2016) document this effect, reproduced in Appendix D Figure 31. They also document that the increase in take-up rate substantially varied across states, because different states adopted a different policy mix.[46] 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.

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 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. UPCs can be thought of as are partly non-tradable because of the strength of local brand preferences (Bronnenberg, Dube and Gentzkow, 2012). However, 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

[46] For 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.

estimate the signal standard deviation of the "national preference" component, denoted $\hat{\sigma}^2_\lambda$, and the signal standard deviation of the "local preference" component $\hat{\sigma}^2_\mu$:[47]

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

Finally, I rank product groups according to the quantity $R = \frac{\hat{\sigma}^2_\mu}{\hat{\sigma}^2_\lambda}$. 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.

4.4.2 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 market thickness documented in the previous subsection.

Table 13 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.[48]

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

[48]I 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.

Table 13: Results from SNAP Research Design

Panel A: Main Results

	Actual Spending Growth for SNAP Eligible (%)	Difference in Overlapping Goods Inflation Rate (pp)
Change in Take-up Rate (pp), 2001-2007	0.2226***	-0.0242***
	(0.0770)	(0.00791)
Controls	Yes	Yes
Weights	Yes	Yes

Standard errors clustered by 50 states

Panel B: Robustness

	Difference in Overlapping Goods Inflation Rate (pp)		
Change in Take-up Rate (pp), 2001-2007	-0.0242***	-0.0195***	-0.0151**
	(0.00791)	(0.00654)	(0.00776)
2001 Take-up Rate	Yes	Yes	Yes
Unemployment		Yes	Yes
Population			Yes
Employment growth			Yes
Total Labor Force			Yes

Standard errors clustered by 50 states

Panel C: Local vs. National Brand Preferences

	Difference in Overlapping Goods Inflation Rate (pp)		
Change in Take-up Rate (pp), 2001-2007	-0.00828**	-0.01687***	-0.000597
	(0.00398)	(0.00523)	(0.00560)
All product groups	X		
Top 50% by "local" preferences		X	
Bottom 50% by "local" preferences			X

Standard errors clustered by 50 states

Panel D: Food vs. Non-Food Products

	Difference in Overlapping Goods Inflation Rate (pp)			
Change in Take-up Rate (pp), 2001-2007	-0.0187**	-0.0115	-0.0255***	-0.0124
	(0.00808)	(0.0203)	(0.0110)	(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

Standard errors clustered by 50 states

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.[49] 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).

4.4.3 Additional Evidence from Variation in the Rate of Growth of Inequality Across US States

The findings from the SNAP research design show that inflation goes down when spending per capita increases in a given product category. However, the variation used in this research design is specifically about the lower tail of the income distribution. Could it be the case that inflation responds very differently to changes in income per capita at other points of the income distribution?

I provide suggestive evidence that the effect is similar in other parts of the income distribution by exploiting variation in the rate of inequality growth 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 sates. 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 13 shows this result.

Figure 13: The Inflation Difference Between High- and Low-Income Increases as Inequality Increases Faster

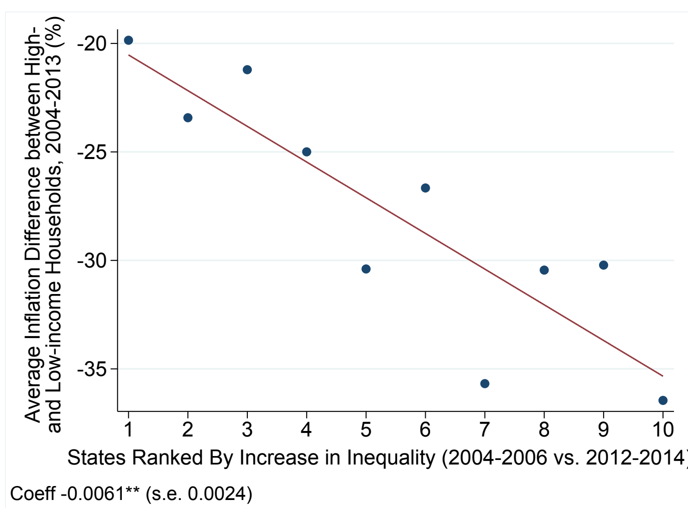

[49] On the fungibility of money and spending choices, see Shapiro and Hastings (2013).

4.5 Alternative Mechanisms

I have investigated 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 disproportionatley 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 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 (Appendix Table 8) or to certain departments, product groups or product modules (Tables 2 and 3).

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 2).

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 24 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.[50]

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 3 shows is not the case.

Other mechanisms. In ongoing work, I use Nielsen TDLink data to characterize changes in the competitive environment of retailers and I document how competitive dynamics differ across areas depending on the density of high- and low- income households. Finally, I test the predictions of models featuring dynamic pricing and increasing returns to scale.

5 Calibrations, External Validity, and Implications

This section brings together the previous results using simple calibrations showing that historical changes in relative demand explain most of the inflation difference across income groups. It then presents new evidence supporting the external validity of these results in a broader sample of goods and over a long time period. Finally, it discusses implications for public policy, our understanding of inequality, and our understanding of innovation dynamics.

5.1 Calibrations

In Section 4, the market thickness research design and the SNAP research design both delivered the result that increases in market size, whether from more consumers or more spending per capita, lead to lower quality-adjusted inflation. I now present simple calibrations combining the historical variation in purchasing power across income groups with the point estimates for the equilibrium response of inflation to changes in market size.

Over the course of the sample, and more broadly over recent decades, demand from high-income consumers

[50] Intuitively, 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).

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. This process is sometimes referred to as "structural change." 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).

I rely on public-use micro data from the Census (IPUMS) to calibrate changes in demand from high- and low-income consumers. I then use the point estimates from Section 4 to infer the implied effect on inflation for these consumers, which I find to be very large. The calibrations show that, given the historical changes in the nominal income distribution, the "endogenous supply response" channel explains almost all of the observed difference in inflation rates across income groups.

5.1.1 Changes in Market Thickness across Income Groups

The "market thickness" channel refers to the fact that more and more consumers became high-income earners, which induced a supply reponse. To calibrate the quantitative importance of this channel, I use the observed change over time in the number of households making above \$100,000 and below \$30,000.

Between 2004 and 2013, on average the number of high-income households (earning above \$100,000) grew 3.12pp faster than the number of low-income households (making below \$30,000). Multiplying this number by the point estimate in the second column of Table 9 implies that historical changes in market thickness across income groups caused an annual inflation difference of $3.12 \times 11.6 = 36.1$ basis points, which represents $\frac{36.1}{40.8} = 88.4\%$ of the benchmark inflation difference, at the product module$\times$price decile level. The 40.8 basis point benchmark inflation difference is taken from the fourth row of Table 5: it is the relevant benchmark because the regressions were all conducted at the product module$\times$price decile level, a level at which the inflation difference between high- and low-income is already attenuated because of aggregation bias. Moreover, the implied annual inflation difference of 36.1 basis point represents $\frac{36.1}{66.0} = 54.6\%$ of the overall inflation difference between high- and low-income households documented in Section 3. Thus, the response of inflation to changes in market thickness is sufficiently large to explain most of the overall difference in inflation rates across income groups, and almost all of the relevant benchmark inflation difference.

5.1.2 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 pp faster than that of low-income households. By taking the ratio of the point estimates in Panel A of Table 13, I obtain that a 1pp 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 "market thickness" channel.

Taken together, the calibrations show that the endogenous supply reponse to changes in market thickness and spending per capita across income groups explains close to the entirety of the inflation difference between high- and low-income households.[51]

5.2 External Validity

Whether the findings documented so far apply to other sectors beyond retail and over longer horizons remains an open question.[52] To make progress on this issue, I follow two complementary approaches. First, 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.

Second, I use more detailed CEX share data and CPI price series to characterize the sign and magnitude of the inflation difference between high- and low-income households for the full consumption basket, going back to 1953. Due to data limitations, this analysis is of course much coarser than the previous analysis based on the Nielsen sample, but it provides striking and transparent evidence supporting the external validity of my findings.

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

Assume that households' utility function is CES with $\sigma^i > 1$ over an aggregator for Nielsen goods, denoted N_i, and an aggregator for outside goods, denoted O_i. In other words, Nielsen goods are 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).

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 where 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 $\sigma^i > 1$, this implies that the relative price of the high-income consumption basket was declining even faster for outside goods, relative

[51] In fact, adding up the calibrated numbers for the relevant inflation difference benchmark, I get that the supply response channel explains $88.4 + 24.7 = 113\%$ of the inflation difference, i.e. it "over-explains" it. This is not surprising, given that other forces are likely to lower the inflation of lower-income households, for instance the international trade channel, whose relevance was documented at the end of Section 4.

[52] See the debate between Moretti (2013) and Diamond (2016) about price changes in the housing sector for college and high-school graduates. To my knowledge, there have been very few attempts at developing non-homothetic price indices and McGranahan and Paulson (2005) remains the main reference. Handbury (2013) developed a non-homothetic price indices across cities but did not have panel data to study inflation dynamics.

to the low-income consumption basket.

Formally,

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

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,

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

Appendix A provides a formal proof. 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.

5.2.2 A Direct Measure of the Full-Basket Inflation Difference between High- and Low-Income Household using CPI and CEX Data

I use BLS and CEX data to probe the external validity of the two core findings of the paper: inflation is lower for the high-income, and this is largely due to the response of supply to market size effects.

I proceed in two steps. First, I collect CPI price series on 48 CEX expenditure categories going back to 1953, which cover the full consumption basket. These categories are listed in Appendix B and are matched by hand across the CPI and CEX surveys. Second, I build price indices for the consumption baskets of college graduates and high-school dropouts, using expenditure shares fixed at 1980-1985 levels (which are observed in the CEX data). I focus on inflation patterns across education groups instead of income groups for two reasons: education is much better measured than income in the early CEX surveys (and education is potentially a good proxy for permanent income); and over a long time horizon changes in relative market size are much easier to measure across education groups than across income groups (it is well documented that the number of college graduates and the college premium started increasing in the 1970s, e.g. in Autor, Katz and Krueger, 1998).

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 therefore 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.[53]

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

To probe the external validity of the core findings of the paper, I use the CPI and CEX data to ask two questions: is inflation lower for the consumption basket of college graduates, relative to high-school dropouts, over a long horizon? and is the difference getting larger after the 1970s, when demand from college graduates starts rising faster, in the broader context of increasing inequality?

The answer to both questions is a resounding yes. Figure 14 establishes this by plotting the price index of high-school dropouts relative to college graduates from 1953 to 2015. Relative to high-school dropouts, average annual inflation for college graduates was 10 basis points lower during 1953-1970 and 25 basis points 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 in the exercise I conducted in Section 3 using the data of McGranahan and Paulson (2005).

Figure 14: Inflation Difference between High-School Dropouts and College Graduate for Full Consumption Basket

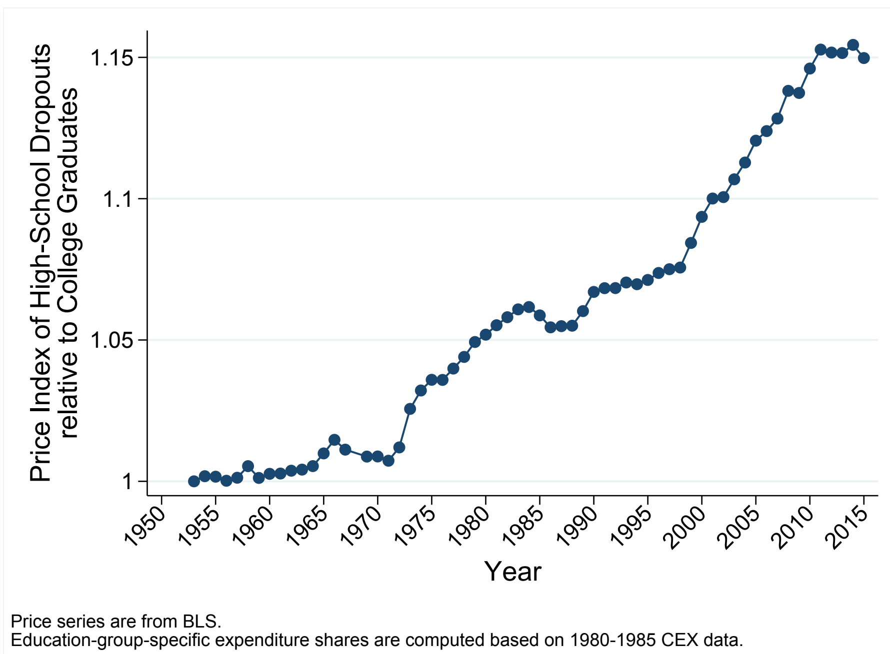

These results are robust and are not driven by any single broad product category. In ongoing work, I test the robustness of the results further by considering other base years for the spending shares, as well as other education or income groups.

5.3 Implications

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. evidence on external validity above). Several government transfers are indexed on food-at-home CPI (e.g. food stamps); many others are indexed on the full-basket CPI (e.g. Social Security), and so are income poverty thresholds[54] and tax brackets. In order 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 at 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.[55]

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 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, 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 reponse 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.

Inequality. The two mains lessons of this paper regarding inequality are that real inequality is increasing faster than commonly thought, at least in the retail sector but probably also beyond,[56] and that changes in nominal inequality have an amplification effect, because of the response of supply to changes in relative market size (cf. Figure 13). 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 no speak to that part of the

[54]Following Orshansky (1962), poverty is measured according to an "absolute" scale 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.

[55]Indeed, Table 5 showed largely result from inflation difference across the quality distribution, at least in retail.

[56]As previously discussed, extrapolating the inflation difference found in Nielsen to the broad economy, which is justified by the structural exercise presented above, means that real inequality increased 70% faster than nominal inequality! Note that I have carried out the analysis by considering a specific notion of inequality, comparing a group of high-income consumers earning above $100,000 per year to a group of low-income households earning below $30,000 per year. But similar patterns exist when considering other cuts of the income distribution: the only limitation of my data is that I cannot study top income households, e.g. the top 1% or 0.1% of the income distribution.

income distribution, where quality-adjusted consumption is very difficult to measure.[57]

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" — a new flavor, a new size, etc. —, 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 impotant 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 it is the case for retail, and a similar logic might apply in other sectors, as suggested by the evidence I presented on external validity.

6 Conclusion

In this paper, I have shown that quality-adjusted inflation substantially varies across income groups in the retail sector. The current methodology of statistical agencies like the Bureau of Labor Statistics cannot capture this variation, which exists primarily at the product level rather than across broad item categories. Furthermore, I have established that product introductions and prices endogenously respond to changes in market size in a way that magnifies the welfare effects of changes in nominal inequality. As shown in simple calibrations, the endogenous response of supply to changes in market size over the past decade can explain most of the observed difference in inflation rates across income groups during this period. Finally, more aggregate data on the full consumption basket of American households back to 1953 supports the external validity of these findings.

This paper opens up several directions for future research. Could a similar analysis be conducted with suitable micro-data beyond the retail sector and beyond the United States? How should one adjust optimal redistributive taxation formulas (e.g. as in Mirrlees, 1971) to take into account the endogenous response of supply to changes in market size? These and other extensions await further research.

[57]Indeed, 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).

BIBLIOGRAPHY

Acemoglu, Daron, 2002, "Directed Technical Change," The Review of Economic Studies, Vol. 69, No. 4.

Acemoglu, Daron, 2007, "Equilibrium Bias of Technology," Econometrica, Econometric Society, vol. 75(5), pages 1371-1409.

Argente, David and Munseob Lee, 2015, "Cost of Living Inequality During the Great Recession," Working Paper.

Aguiar, Mark and Erik Hurst, "Consumption versus Expenditure," Journal of Political Economy, 2005, 113 (5), 919–948.

Aguiar, Mark and Erik Hurst, "Life-Cycle Prices and Production," American Economic Review, 2007, 97 (5), 1533–1559.

Beraja, Martin, Erik Hurst, and Juan Ospina, "The Aggregate Implications of Regional Business Cycles," Working Paper, 2015.

Broda, Christian and David E Weinstein, "Globalization and the Gains From Variety," Quarterly Journal of Economics, 2006, 121 (2), 541–585.

Broda, Christian and David E Weinstein, "Product Creation and Destruction: Evidence and Price Implications," American Economic Review, 2010, 100 (3), 691–723.

Broda, Christian and John Romalis, "TheWelfare Implications of Rising Price Dispersion," 2009.

Chevalier, Judith A and Anil K Kashyap, "Best Prices: Price Discrimination and Consumer Substitution," NBER-Working Paper No. 20768, 2014.

Coibion, Olivier, Yuriy Gorodnichenko, and Gee Hee Hong, "The Cyclicality of Sales, Regular and Effective Prices: Business Cycle and Policy Implications," American Economic Review, 2015, 105 (3), 993–1029.

Diewert, W Erwin, "Exact and Superlative Index Numbers," Journal of Econometrics, 1976, 4 (2), 115–145.

Feenstra, Robert C, "New Product Varieties and the Measurement of International Prices," American Economic Review, 1994, 84 (1), 157–177.

Gopinath, Gita, Pierre-Olivier Gourinchas, Chang-Tai Hsieh, and Nicholas Li, "International Prices, Costs and Mark-Up Differences." American Economic Review, 2011, 101 (6): 2450–86.

Handbury, Jessie, "Are Poor Cities Cheap for Everyone? Non-Homotheticity and the Cost of Living Across US Cities," 2013.

Hobijn, Bart and David Lagakos, "Inflation Inequality in the United States," Review of Income and Wealth, 2005, 51 (4), 581–606.

McGranahan, Leslie and Anna Paulson, "Constructing the Chicago Fed Income Based Economic Index-Consumer Price Index: Inflation Experiences by Demographic Gropu: 19832-2005," Federal Reserve Bank of Chicago Working Paper.

Pisano, Luigi and Andrea Stella, "Price Heterogeneity and Consumption Inequality," Working Paper.

Sato, Kazuo, "The Ideal Log-Change Index Number," Review of Economics and Statistics, 1976, 58 (2), 223–228.

Stroebel, Johannes and Joseph Vavra, "House Prices, Local Demand, and Retail Prices," NBERWorking Paper No. 20710, 2015.

Zhelobodko, Evgeny, Sergey Kokovin, Mathieu Parenti and Jacques-Francois Thisse, 2012, "Monopolistic Competition: Beyond the Constant Elasticity of Substitution," Econometrica, Vol. 80, No. 6.

Appendix A

More on Theory

A. The Strong Equilibrium Bias in a Model Following Acemoglu (2002)

Consumers

Individual Demands

Preferences are different for the two groups:

$$U_{poor} = a + u(x_L) = a + \frac{1}{1-\alpha_L} x_L^{1-\alpha_L}$$

$$U_{rich} = a + v(x_R) = a + \frac{1}{1-\alpha_R} x_R^{1-\alpha_R}$$

where a is the numeraire. The FOC for the composite high-quality and low-quality goods are:

$$u'(x_L) = x_R^{-\alpha_L} = p_L$$

$$v'(x_R) = x_R^{-\alpha_R} = p_R$$

Nominal inequality is measured by w_{rich}/w_{poor}, while real inequality is measured by $\frac{w_{rich}}{p_H}/\frac{w_{poor}}{p_L}$.

Aggregate Demands

Normalize the total mass of consumers to 1, with share λ of rich types (if desired, we can introduce another scaling factor to study market size effect due to the total number of consumers). Aggregate demands are given by:

$$D_L = (1-\lambda) p_L^{-\frac{1}{\alpha_L}}$$

$$D_H = \lambda p_H^{-\frac{1}{\alpha_H}}$$

Producers

Final Producer

Final producer just uses capital and combines "varieties $x(v,i)$" to produce two types of goods, H or L (profits are thrown away). With $i = L/H$, we can write the problem as:

$$max_{|x(v,i)| \in N(i)} \frac{p(i)}{1-\epsilon_i} \left(\int_0^{N(i)} x(v,i)^{1-\epsilon_i} dv \right) - \int_0^{N(i)} p^x(v,i) x(v,i) dv$$

Note that the returns to scale are decreasing. The optimal choice is:

$$x(v,i) = \left(\frac{p(i)}{p^x(v,i)} \right)^{\frac{1}{\epsilon_i}}$$

We denote by $\sigma_i = \frac{1}{\epsilon_i}$ the elasticity of substitution between machines.

Intermediate Producers

The intermediate monopolist has a patent and chooses the optimal price (we consider one period only here, but easy to extend since problem is separable). The cost of production of a machine is ψ_i units of the final good. The value of a patent (of a variety) for the intermediate good in sector i (lasting one period) is:

$$V(v,i) = (p^x(v,i) - \psi_i)x(v,t)$$

where the optimal price chosen by the monopolist maximizes (taking demand as given):

$$max_{p^x(v,i)}\ (p^x(v,i) - \psi_i)\left(\frac{p(i)}{p^x(v,i)}\right)^{\frac{1}{\epsilon_i}}$$

Hence the optimal choice:

$$p^x(v,i) = \frac{\psi_i}{1-\epsilon_i}$$

The value function at the optimum is:

$$V(v,i) = p(i)^{\frac{1}{\epsilon_i}}(\frac{1-\epsilon_i}{\psi_i})^{\frac{1-\epsilon_i}{\epsilon_i}}\epsilon_i$$

Aggregate Supply

Total quantity supplied in equilibrium is given by:

$$S_i^* = \frac{1}{1-\epsilon_i}\left(\int_0^{N(i)} x(v,i)^{1-\epsilon_i}dv\right) = \frac{1}{1-\epsilon}\left(\int_0^{N(i)}\left(\frac{p(i)}{p^x(v,i)}\right)^{\frac{1-\epsilon_i}{\epsilon_i}}dv\right)$$

$$= \left(\frac{1}{1-\epsilon_i}\right)^{\frac{1}{\epsilon_i}}\left(\frac{p(i)}{\psi_i}\right)^{\frac{1-\epsilon_i}{\epsilon_i}}N(i)$$

Solving for the equilibrium

With exogenous varieties

$$p_H = \left[\left(\frac{\lambda}{N_H}\right)(1-\epsilon_H)^{\frac{1}{\epsilon_H}}(\psi_H)^{\frac{1-\epsilon_H}{\epsilon_H}}\right]^{\frac{1}{\frac{1-\epsilon_H}{\epsilon_H}+\frac{1}{\alpha_H}}}$$

$$p_L = \left[\left(\frac{1-\lambda}{N_L}\right)(1-\epsilon_L)^{\frac{1}{\epsilon_L}}(\psi_L)^{\frac{1-\epsilon_L}{\epsilon_L}}\right]^{\frac{1}{\frac{1-\epsilon_L}{\epsilon_L}+\frac{1}{\alpha_L}}}$$

Relative market size *increases* the relative price, because supply is fixed. So real inequality when prices are endogenous is “lower” than real inequality when prices are exogenous.

With endogenous varieties

Interior solution

At an interior solution (there is research on both kinds of goods), the no arbitrage condition between two types of inventions requires:

$$\eta V(v,H) = \eta V(v,L)$$

The equilibrium ratio of varieties and the equilibrium prices are given by:

$$\frac{(N_H)^{\frac{1}{\epsilon_H}}}{(N_L)^{\frac{1}{\epsilon_L}}} = \frac{(\lambda)^{\frac{1}{\epsilon_H}}}{(1-\lambda)^{\frac{1}{\epsilon_L}}} \left(\frac{\kappa_H}{\kappa_L}\right)^{\zeta} \frac{(\psi_H)^{\frac{1-\epsilon_H}{\epsilon_H^2}}}{(\psi_L)^{\frac{1-\epsilon_L}{\epsilon_L^2}}} \frac{(1-\epsilon_H)^{\left(\frac{1}{\epsilon_H}\right)^2}}{(1-\epsilon_L)^{\left(\frac{1}{\epsilon_L}\right)^2}}$$

$$p_H = (\frac{\psi_H}{1-\epsilon_H})^{1-\epsilon_H} \left(\frac{1}{\epsilon_H}\right)^{\epsilon_H}$$

$$p_L = (\frac{\psi_L}{1-\epsilon_L})^{1-\epsilon_L} \left(\frac{1}{\epsilon_L}\right)^{\epsilon_L}$$

$$V^* = 1$$

Corner solutions and the strong equilibrium bias

$$\begin{cases} V(v,H) > V(v,L) \\ N_H = N + \bar{N_H} \;\&\; N_L = \bar{N_L} \end{cases}$$

$$p_L = \left[\left(\frac{1-\lambda}{\bar{N_L}}\right)(1-\epsilon_L)^{\frac{1}{\epsilon_L}} (\psi_L)^{\frac{1-\epsilon_L}{\epsilon_L}}\right]^{\frac{1}{\frac{1-\epsilon_L}{\epsilon_L}+\frac{1}{\alpha_L}}}$$

Consider a relative demand shock (compared with previous periods, where the steady state relative demand is embodied in the steady state relative number of varieties).

Assume that all research is allocated to the high quality good:

$$\begin{aligned} V(v,H) \quad &> \quad V(v,L) \\ (Profitability\ Condition) \\ \iff \frac{\left(\frac{\lambda}{N+\bar{N_H}}\right)}{\left(\frac{1-\lambda}{\bar{N_L}}\right)} &> \left[\frac{(\frac{1-\epsilon_L}{\psi_L})^{\frac{1-\epsilon_L}{\epsilon_L}}\epsilon_L}{(\frac{1-\epsilon_H}{\psi_H})^{\frac{1-\epsilon_H}{\epsilon_H}}\epsilon_H}\right]^{\zeta} \frac{(1-\epsilon_L)^{\frac{1}{\epsilon_L}}(\psi_L)^{\frac{1-\epsilon_L}{\epsilon_L}}}{(1-\epsilon_H)^{\frac{1}{\epsilon_H}}(\psi_H)^{\frac{1-\epsilon_H}{\epsilon_H}}} \end{aligned}$$

There is “overshooting” of the relative price (strong equilibrium bias) if the new relative price is smaller than the old one:

$$\frac{p_H}{p_L} < \frac{\bar{p_H}}{\bar{p_L}} \qquad (Price\ Overshooting\ Condition)$$

$$\iff \frac{\frac{\lambda}{N+\bar{N_H}}}{\frac{1-\lambda}{\bar{N_L}}} < \left[\frac{(\psi_H)^{\frac{1-\epsilon_H}{\epsilon_H}}}{(\psi_L)^{\frac{1-\epsilon_L}{\epsilon_L}}} \cdot \frac{(1-\epsilon_L)^{\frac{1-\epsilon_L}{\epsilon_L}}}{(1-\epsilon_H)^{\frac{1-\epsilon_H}{\epsilon_H}}} \cdot \frac{\epsilon_L}{\epsilon_H}\right]^{\zeta} \frac{(1-\epsilon_L)^{\frac{1}{\epsilon_L}}(\psi_L)^{\frac{1-\epsilon_L}{\epsilon_L}}}{(1-\epsilon_H)^{\frac{1}{\epsilon_H}}(\psi_H)^{\frac{1-\epsilon_H}{\epsilon_H}}}$$

which cannot be satisfied at the same time as the profitability condition.

B. The Strong Equilibrium Bias in a Model Following Melitz and Ottaviano (2008)

- Preferences for λ high-income consumers and $1-\lambda$ low-income consumers

$$U_i = q_{0i}^c + \alpha \int_{\omega \in \Omega_i} q_{\omega i}^c d\omega - \frac{1}{2}\gamma \int_{\omega \in \Omega_i} (q_{\omega i}^c)^2 d\omega - \frac{1}{2}\eta (\int_{\omega \in \Omega_i} q_{\omega i}^c d\omega)^2$$

$$\implies p_{\omega i} = \alpha - \gamma q_{\omega i}^c - \eta Q_i^c \ ; \ Q_i = \int_{\omega \in \Omega_i} q_{\omega i}^c d\omega$$

- To enter the differentiated sector, a firm must incur a sunk entry cost of f_E units of labor
  - Then the firm's unit labor requirement of cost c is drawn from a cumulative distribution function $G(c)$ with support on $[0, c_M]$
  - The zero-profit cost cutoff (c_{Di}) is a sufficient statistic that determines firm outcomes as a function of their cost draw:

$$\begin{aligned} p_i(c) &= \frac{1}{2}(c_{Di} + c) \quad (prices) \\ \mu_i &= p_i(c) - c = \frac{1}{2}(c_{Di} - c) \quad (markups) \\ r_i(c) &= \frac{L_i}{4\gamma}[(c_{Di})^2 - c^2] \quad (revenues) \\ \pi_i(c) &= \frac{L_i}{4\gamma}[c_{Di} - c^2] \quad (profits) \end{aligned}$$

- Under the assumption that productivity $\frac{1}{c}$ is Pareto distributed with lower bound $\frac{1}{c_M}$ and shape parameter k, the (closed economy) cost cutoff is given by:

$$c_{Di} = (\frac{\gamma \phi}{L_i})^{\frac{1}{k+2}}$$

  - So the cost cutoff falls (meaning the average productivity is higher) when varieties are closer substitutes (lower γ), when there is a better distribution of cost draws (lower c_M), when sunk costs fall (lower f_E) and in bigger markets (higher L_i). These comparative statics induce an increase in the "toughness of competition" in the form of a larger number of varieties consumed (higher N_i) and lower average prices (lower $\bar{p_i}$).

- This implies that the relative price decreases with (relative) market size
  - Intuition: firms are "locked in" and the long-run supply curve is downward sloping because of entry
  - Note that even in the short run relative prices will never mitigate the increase in inequality
    * Intuition: the marginal cost of production is constant: $p = \frac{\gamma}{L}q + \psi$
    * Can generate price effects mitigating increases in inequality by introducing specialized labor

C. A Model with Endogenous Markups, Endogenous Marginal Cost Improvements, and Translog Preferences

Notation

Households have preferences over a continuum of differentiated goods in the set Δ. This set includes the total number of actual, old, and potential (not yet invented) goods and has a measure of $\hat{N}$. Let $\Delta^{'}$, with measure N, be the subset of Δ that contains the set of goods that are actually availabe for purchase at Home.

Preferences and Demand

The symmetric translog expenditure function, defined over the continuum of available goods, is given by:

$$ln(E) = ln(U) + a + \frac{1}{N}\int_{i\in\Delta'} ln(p_i)di + \frac{\gamma}{2N}\int_{i\in\Delta'}\int_{j\in\Delta'} ln(p_i)\left(ln(p_j) - ln(p_i)\right)djdi$$

where $a = \frac{1}{2\gamma N}$. γ is always positive a high γ implies high substitutability (or low differentiation). The intuition for the various terms is as follows:

By Shephard's lemma we get the share of expenditures on good i:

$$s_i = \gamma ln(\frac{\hat{p}}{p_i})$$

$$\hat{p} = exp\left(\frac{1}{\gamma N} + \overline{ln(p)}\right)$$

$$\overline{ln(p)} = \frac{1}{N}\int_{i\in\Delta'} ln(p_i)di$$

where $\hat{p}$ is the maximum price a firm can set.

The demand of the representative household, with income I, is then given by:

$$q_i = s_i\frac{I}{p_i} \tag{3}$$

Profit Maximizing Price

The monopolist takes the residual demand curve as given, and we assume constant marginal cost for the production of good i.

$$argmax_{p_i}\ \ p_iq_i - mc_iq_i$$

The optimal price is such that

$$p_i = \left(1 + ln(\frac{\hat{p}}{p_i})\right)mc_i$$

Can solve either with the approximation $ln(\frac{\hat{p}}{p_i}) \approx \frac{\hat{p}}{p_i} - 1$ or using the Lambert function Ω (the inverse funtion of $f(\Omega) = \Omega e^{\Omega}$[58]). The exact solution for p_i is therefore

$$p_i = \Omega\left(\frac{\hat{p}}{mc_i}e\right)mc_i \tag{4}$$

[58]Its key properties are: $\Omega'(x) > 0$, $\Omega''(x) < 0$, $\Omega(0) = 0$, and $\Omega(e) = 1$.

We can write the markup of product i as

$$\mu_i = \Omega(\frac{\hat{p}}{mc_i}e) - 1$$

which is strictily decreasing with the marignal cost and reaches 0 when $mc_i = \hat{p}$.

It can also be shown that

$$s_i = \gamma\mu_i \tag{5}$$

i.e. the market share and the markup of producer i are directly proportional.

Production

Firms are indexed by their productivity φ_i. A firm with productivity φ produces with marginal cost $mc_i = \frac{1}{\varphi_i}$. Below we will solve for cases where φ_i is constant vs. not.

To enter the market, all firms have to pay a fixed cost f_E.

Equilibrium

We start by expressing all equilibrium quantities as a function of the minimum price $\hat{p}$, which we then derive as the cutoff productivity level.

Equilibrium quantities

The first equation that comes out of the supply side of the model is equation (2), which by plugging in the expression for markup and for the marginal cost we can write as $p_i = (1+\mu_i)\frac{1}{\varphi_i}$ (marginal benefit = marginal cost for producer). We now combine this with the equation that comes out of the demand side of the model, equation (1): $q_i = s_i\frac{I}{p_i} = \gamma\mu_i\frac{I}{p_i}$ (where the second equality comes from equation (3)). Therefore,

$$p_i = \gamma\mu_i\frac{I}{q_i} = (1+\mu_i)\frac{1}{\varphi_i}$$

$$q_i = I\gamma\frac{\mu_i}{1+\mu_i}\varphi_i$$

Equilibrium profits

Profits are simply given by

$$\pi_i = p_iq_i - mc_iq_i - f_E = \gamma\mu_iI - I\gamma\frac{\mu_i}{1+\mu_i} - f_E = \gamma I\frac{\mu_i^2}{1+\mu_i} - f_E$$

Equilibrium productivity cutoffs

Homogeneous marginal costs

In the case with homogeneous marginal cost $\frac{1}{\varphi}$, firms enter until $\pi_i = 0$, i.e. equilibrium markups are pinned down by:

$$\frac{\mu_i^2}{1+\mu_i} = \frac{f_E}{\gamma I} \tag{6}$$

Now we would like to shows that if market size goes up either because of a change in the number of consumers or in spending per capita, markups must fall, which can happen only if the number of firms increase.

This is easy to see from equation (4), which implicitly defines the equilibrium number of firms. Recall that the expression for markups is:

$$\mu^* = \Omega(\frac{\hat{p}}{mc}e) - 1 = \Omega(\varphi \cdot exp\left(\frac{1}{\gamma N^*} + p^*\right) \cdot e) - 1$$

It looks like we have two unknowns and just one equation. But in fact, the equilibrium price is pinned down by the number of firms. To see this, go back to equation (2):

$$p^* = \Omega\left(\frac{\hat{p}}{mc_i}e\right) mc_i = \Omega\left(\varphi \cdot exp\left(\frac{1}{\gamma N^*} + p^*\right) \cdot e\right) \frac{1}{\varphi}$$

This equation shows that p^* is a function of N^*(and not of any of the other variables like income or market size that we are interested in for the comparative static exercise). By the implicit function theorem, we can show that:

$$\frac{dp^*}{dN^*} > 0$$

A more transparent approach is to use the approximation $ln(\frac{\hat{p}}{p_i}) \approx \frac{\hat{p}}{p_i} - 1$, which delivers:

$$p^* \approx \frac{\hat{p}}{p^*} mc_i = \frac{exp\left(\frac{1}{\gamma N} + ln(p^*)\right)}{p^*} \frac{1}{\varphi}$$

i.e.

$$p^* \approx \frac{exp\left(\frac{1}{\gamma N}\right)}{\varphi}$$

which is clearly declining in N.

Heterogeneous marginal costs

In progress.

Extension with Endogenous Process Innovations

In progress.

D. 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.

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 entire equilibrium is solved for by backward induction.*

Consumer problem and optimization

L consumers with additively separable preferences over varieties solve

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

Maximization yields

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

Total quantity demanded is $q_i = Lx_i$

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,

$$u'(\frac{q_i}{L}) + \frac{q_i}{L} u''(\frac{q_i}{L}) = \lambda V'(q_i) \tag{7}$$

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

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

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 (1) and differentiating by p_i^M using the implicit function theorem::

$$u'(\frac{q_i}{L}) + \frac{q_i}{L}u^{''}(\frac{q_i}{L}) = \lambda p_i^M$$

$$u''(\frac{q_i}{L})\frac{dq_i}{dp_i^M}\frac{1}{L} + \frac{dq_i}{dp_i^M}\frac{1}{L}u^{''}(\frac{q_i}{L}) + \frac{q_i}{L^2}u^{'''}(\frac{q_i}{L})\frac{dq_i}{dp_i^M} = \lambda$$

Let's assume that $u^{'''}(\frac{q_i}{L}) = 0$, i.e. 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,

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

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:

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

The remainder of Appendix A is available from the author upon request.

Appendix B

More on 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 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.

Description of the CPI and CEX data used to measure the full-basket inflation difference between high- and low-income housholds going back to 1953. 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.

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 15: Map of the 76 Local Markets Tracked in the Nielsen Datasets

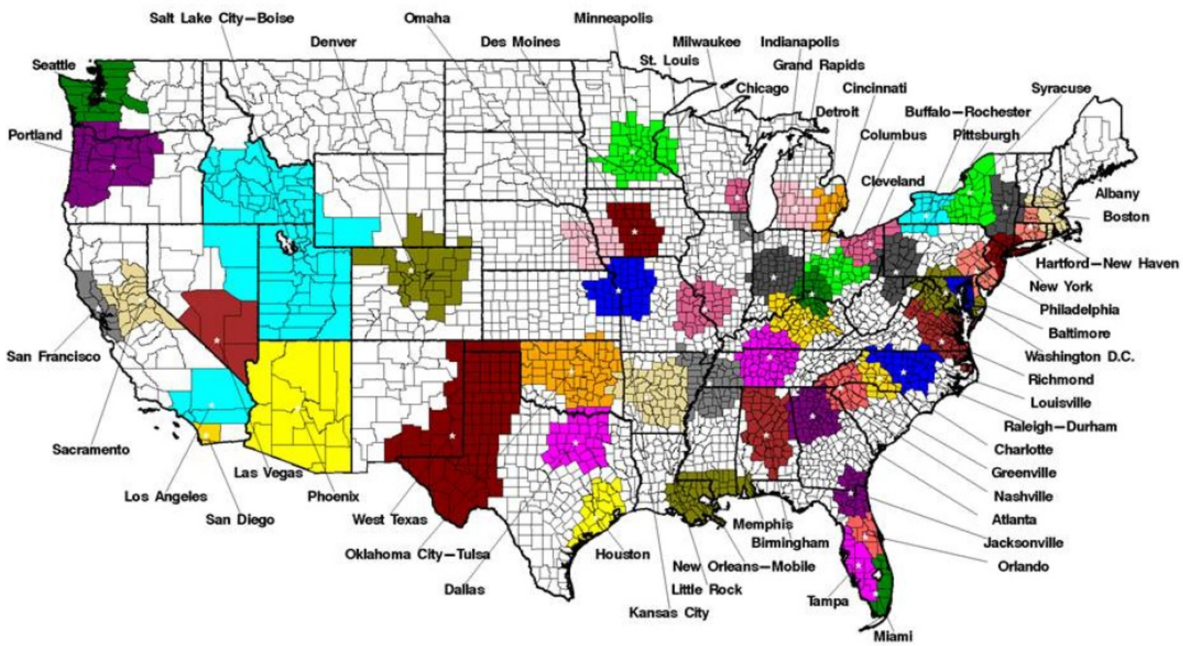

The remainder of Appendix B is available from the author upon request.

Appendix C

Estimation of Quality-Adjusted Inflation and Further Robustness Checks

Nominal and Real Inequality

Figure 16: The Mapping Between Nominal Income and Utility

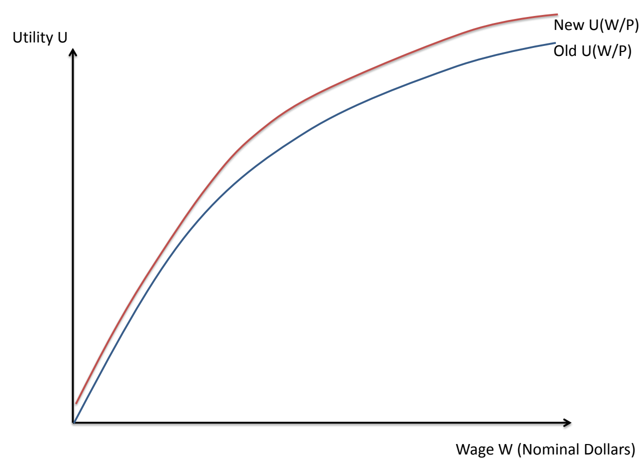

Estimation Equations

Estimating the elasticities: given the formula reported in the main text, we only need to estimate the group-specific and 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:

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

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):

$$ln(c_{umgt}) = \alpha ln(p_{umgt}) + \chi_{mg}$$

$$ln(c_{umgt}) - ln(E_{mgt}) = \alpha ln(p_{umgt}) - ln(E_{mgt}) + \chi_{mg}$$

$$ln(s_{umgt}) = \alpha ln(p_{umgt}) - ln(E_{mgt}) + \chi_{mg}$$

Differencing over time:

$$\begin{aligned} ln(s_{umgt}) - ln(s_{umg(t-1)}) &= \Delta ln(s_{umgt}) \\ &= \alpha \left[ln(p_{umgt}) - ln(p_{umg(t-1)})\right] + ln(E_{mgt}) - ln(E_{mg(t-1)}) \end{aligned}$$

so

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

The fixed effect corresponds to the change in total expenditures in the module (which is observed). In practice there will be estimation error, 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:

$$\Delta^k ln(s_{umgt}) = (1-\sigma_m)\Delta^k ln(p_{umgt}) + \epsilon^k_{umgt} \tag{8}$$

$$\Delta^k ln(p_{umgt}) = \frac{1}{\alpha}\Delta^k ln(s_{umgt}) + \delta^k_{umgt} \tag{9}$$

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

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 $E_t[\epsilon^k_{umgt}\delta^k_{umgt}] = 0$.

$$v_{umgt} = \epsilon^k_{umgt} \times \delta^k_{umgt}$$

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

This can be written as:

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

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

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

If you rewrite $\alpha = \frac{1+\omega_m}{\omega_m}$ to match the notation in Broda and Weinstein (2006), this yields:

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

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

$$E_t\left[\left(\Delta^k ln(p_{umgt})\right)^2\right] = \theta_1 E_t\left[\left(\Delta^k ln(s_{umgt})\right)^2\right] - \theta_2 E_t\left[\Delta^k ln(s_{umgt})\Delta^k ln(p_{umgt})\right] \tag{10}$$

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

Average Inflation Rate of Various Income Groups According to Various Price Indices

Table 14: 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 14: Average Annual Inflation Rates Across Three Income Groups (*Continued*)

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

	Income $<$ \$30k	Income $\in$ [\$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 $\in$ [\$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 14: Average Annual Inflation Rates Across Three Income Groups (*Continued*)

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

	Income $<$ \$30k	Income $\in$ [\$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 15: Average Annual Inflation Rates Across Four Income Groups

Panel A: Full Sample (Percentage Points)

	Income $<$ \$25k	Income $\in$ [\$25k-\$50k]	Income $\in$ [\$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 $\in$ [\$25k-\$50k]	Income $\in$ [\$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 16: 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 17: 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 18: 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 17: Inflation Difference Between Various Income Groups For Various Price Indices (Fixed Basket)

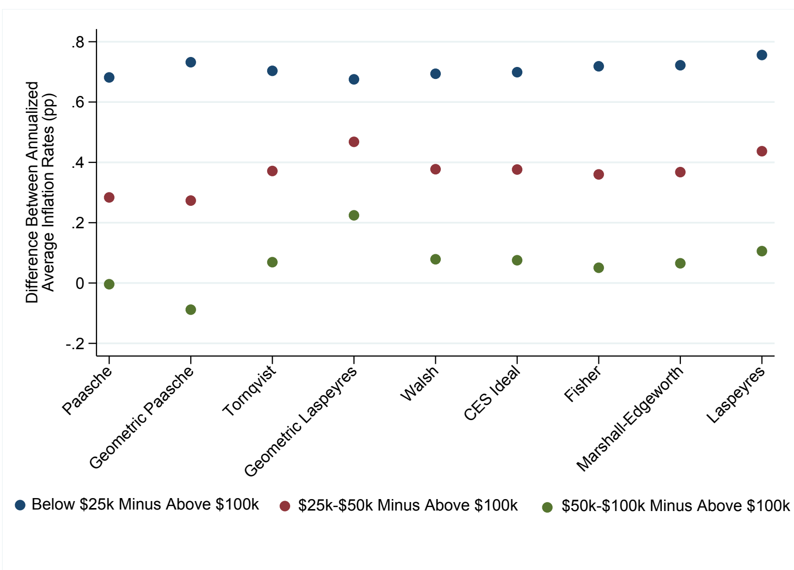

Is Differential Inflation (Fixed Basket) Across Income Groups Driven by a Selection Effect?

Table 19: 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 20: 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 21: Share of spending on new and discontinued products across income groups

	Share of Spending on Products...	
Household Income	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%

Differences in Prices Paid for Same Products for Rich and Poor

Figure 18: The Distribution of Average Unit Prices Paid is the Same Across Income Groups (Reweighting by Spending Shares)

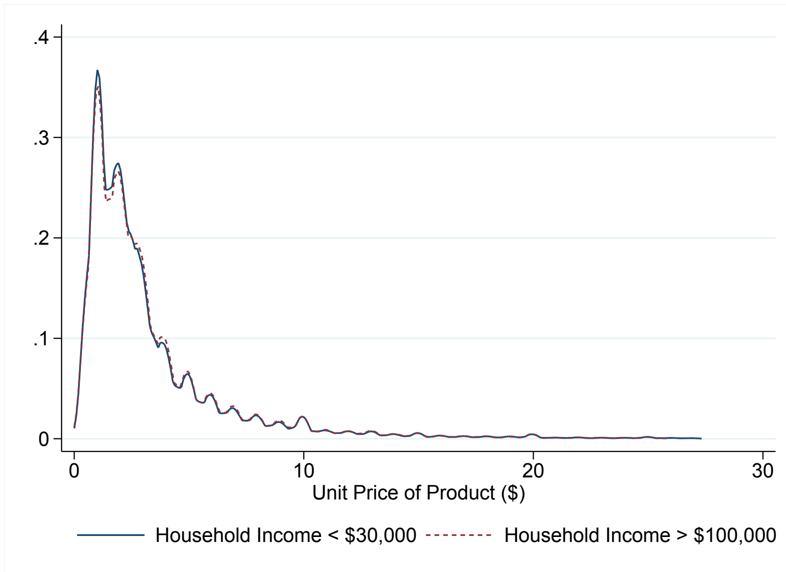

Table 22: Differences in Price Level Paid for Same UPC by High- and Low-Income Households ($)

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

Further Robustness Checks on Increase in Product Variety Across Income Groups

Figure 19: Robustness Checks on Increase in Product Variety across Income Groups

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

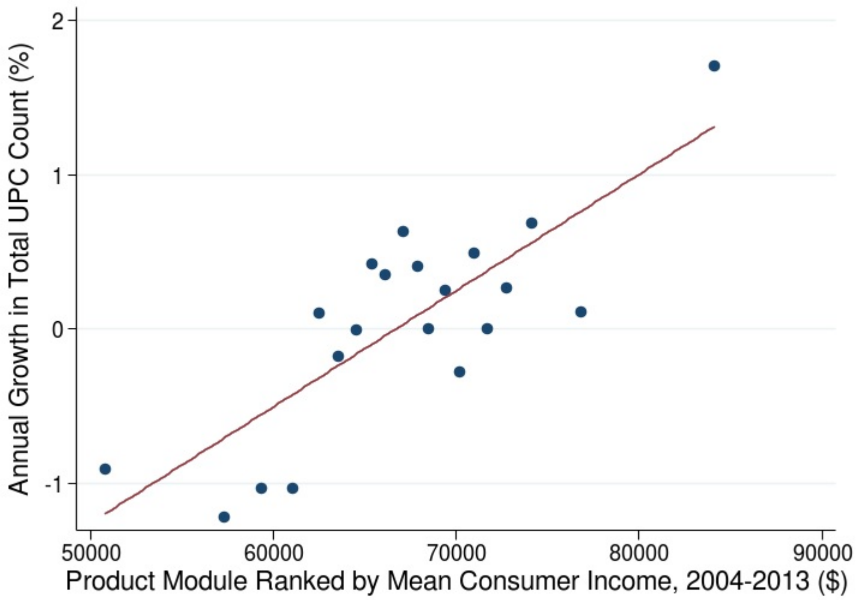

Panel B: Feenstra Ratio across the Product Space

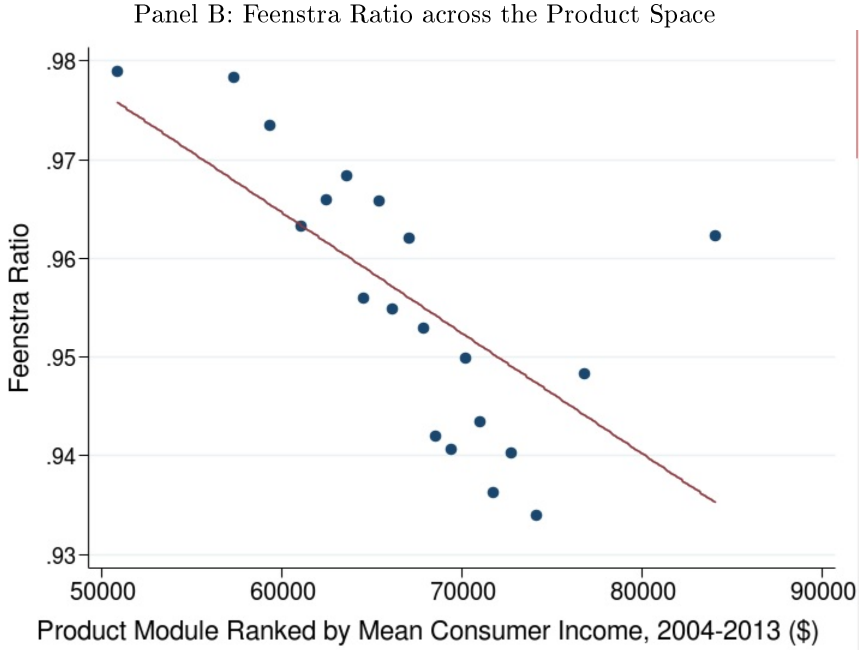

Figure 20: The Positive Relationship Between Share of Spending on New Products and Mean Consumer Income

Panel A: Across Product Modules

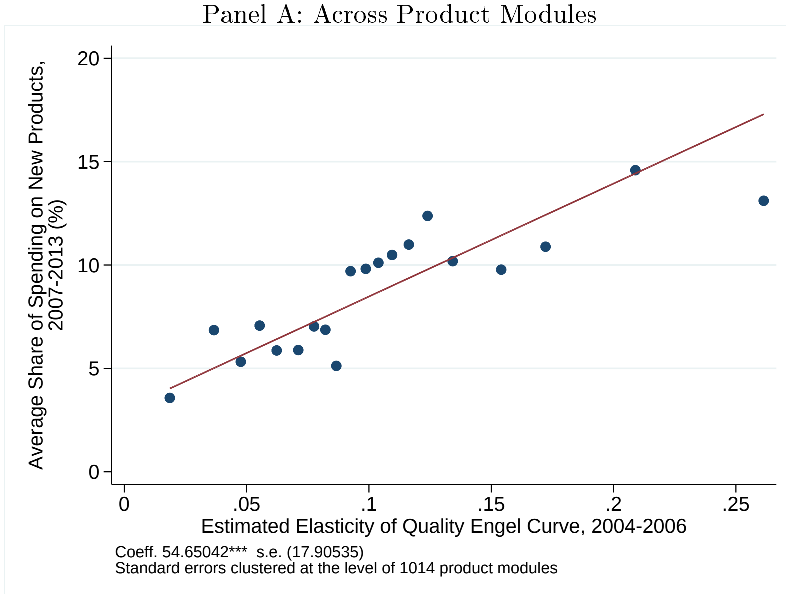

Panel B: Across Product Modules with Product Group Fixed Effects

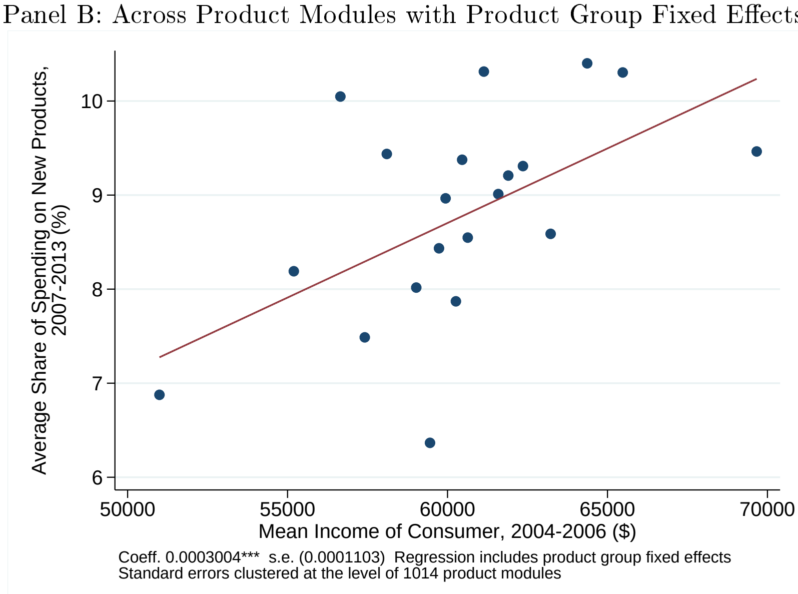

Figure 21: The Relationship Between Share of Spending on New Products and Mean Income Depends on the Quality Engel Curves Elasticity

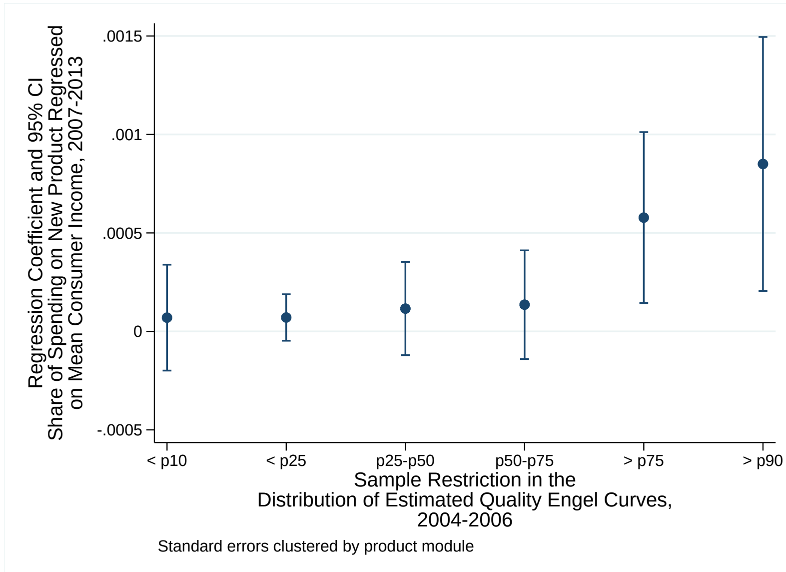

Figure 22: The Relationship Between Share of Spending on New Products and Representative Consumer Mean Income, Controlling for Household Fixed Effects

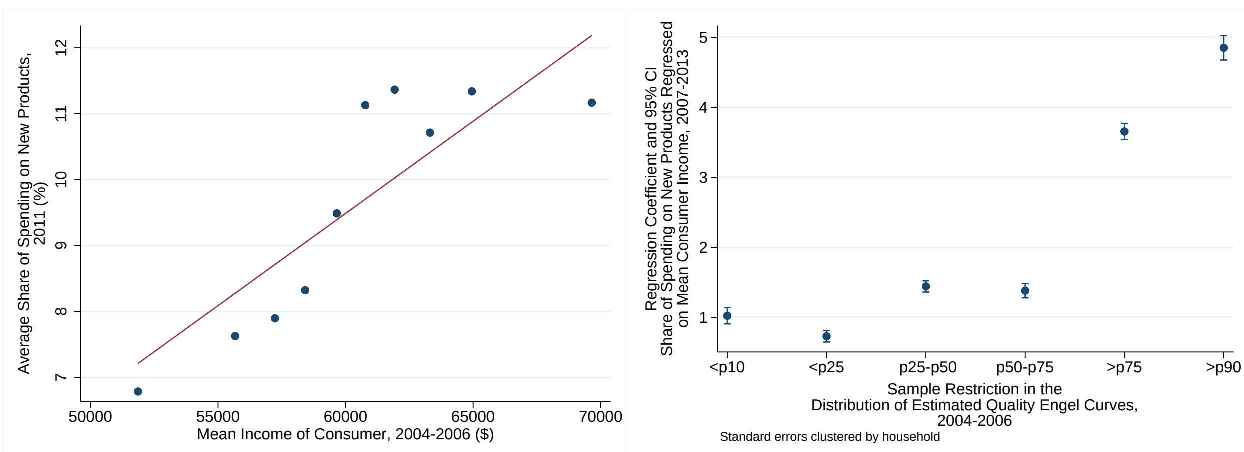

Quality-Adjusted Inflation

Table 23: Distribution of Estimated Module-Level Elasticites of Substitution For Three Income Groups

		Income < \$30k	Income ∈ [\$30k-\$100k]	Income > \$100k
Percentile	10th	9.08	9.30	8.56
	25th	13.91	13.44	13.47
	50th	21.45	19.87	19.78
	75th	35.30	32.74	34.25
	90th	65.98	56.75	71.27

Figure 23: Quality-Adjusted Inflation across Income Groups

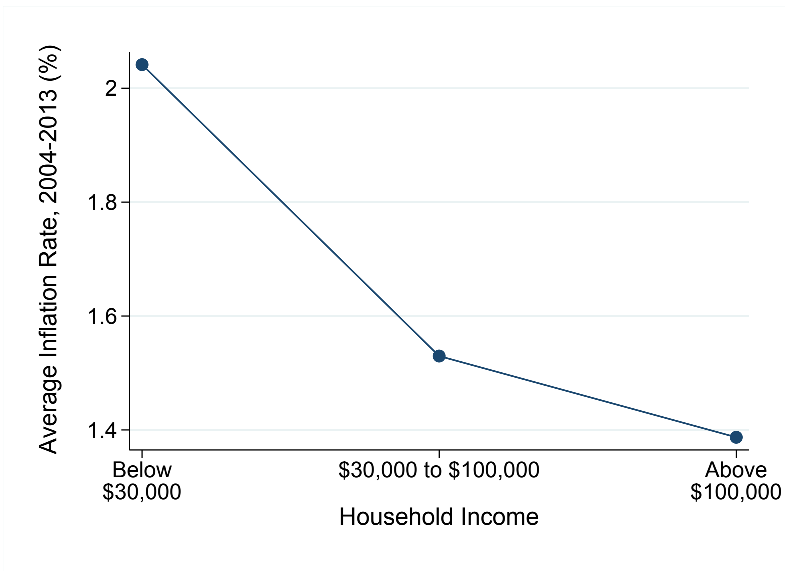

Figure 24: Elasticities of Substitution Differ across Income Groups

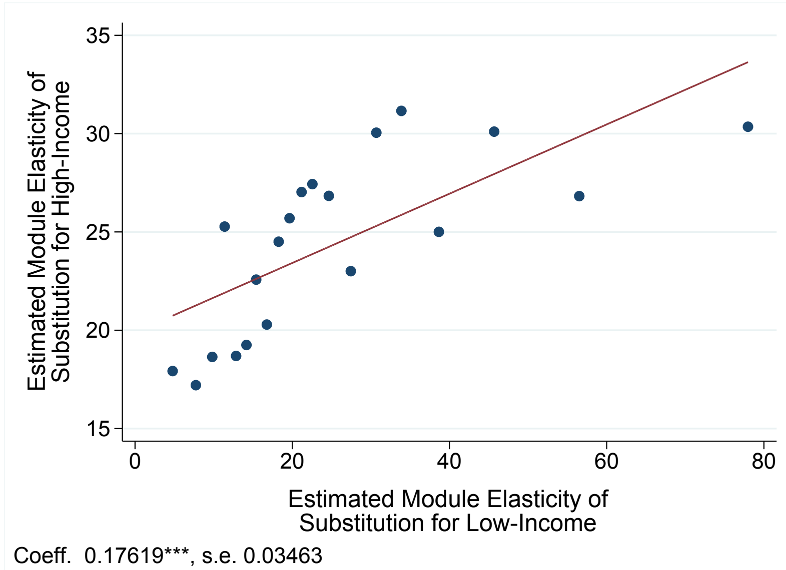

The remainder of Appendix C is available from the author upon request.

Appendix D

Robustness Checks on the Relationship Between Changes in Market Size and Quality-Adjusted Inflation

Additional Stylized Facts

Figure 25: Inflation across Brand Price Deciles, within Product Modules

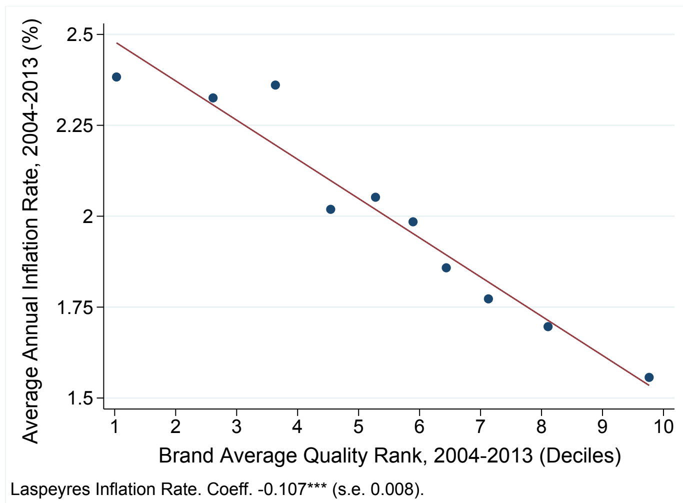

Additional Results on National Research Design

Figure 26: Predicted and Actual Increase in Market Size

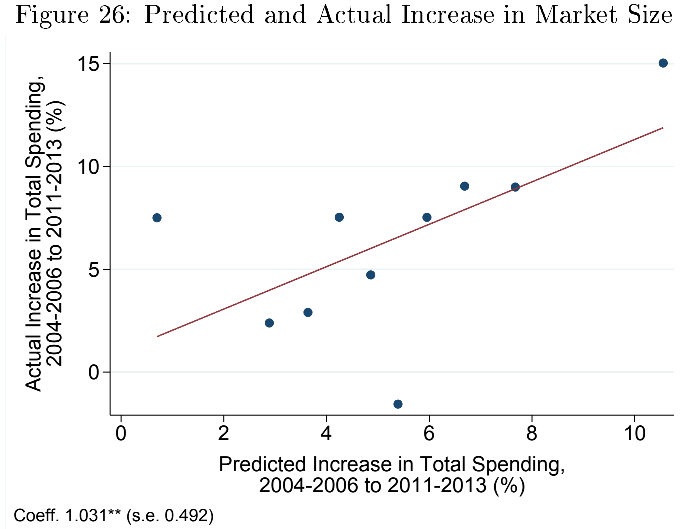

Table 24: Further Robustness Checks on Causal Effects of Changes in Market Size

Panel A: Controlling for 2004-2006 Age and Income Distributions and Price Decile Fixed Effects

	Share of Spending on New Products (pp)			Overlapping Goods Inflation Rate (pp)		
Predicted Increase in Spending (%)	0.527***	0.380***	0.419***	-0.159***	-0.137***	-0.125***
	(0.072)	(0.144)	(0.179)	(0.022)	(0.037)	(0.038)
Age Distribution Controls	Yes	Yes	Yes	Yes	Yes	Yes
Income Distribution Controls	No	Yes	Yes	No	Yes	Yes
Price Decile Fixed Effects	No	No	Yes	No	No	Yes
Product Module Fixed Effects	No	No	No	No	No	No
Spending Weights	Yes	Yes	Yes	Yes	Yes	Yes
Sample Restricted to Positive Spending Growth	Yes	Yes	Yes	Yes	Yes	Yes
Number of Observations	9,089	9,089	9,089	9,089	9,089	9,089
Number of Clusters	1,006	1,006	1,006	1,006	1,006	1,006

Standard errors clustered by product modules.

Panel B: Robustness to Other Weights

	Share of Spending on New Products (pp)	Overlapping Goods Inflation Rate (pp)
Predicted Increase in Spending (%)	0.296***	-0.086***
	(0.0585)	(0.0155)
Product Module Fixed Effects	Yes	Yes
Spending Weights	Yes	Yes
Sample Restricted to Positive Spending Growth	Yes	Yes
Number of Observations	8,545	8,545
Number of Clusters	1,000	1,000

Sample restricted to product modules below 95th percentile of total spending.
Standard errors clustered by product modules.

Geography Research Design

As a robustness check, I use time variation in the age and income distribution of households in 76 local markets tracked by Nielsen within the US (see Appendix B for details). A local market is a county group defined by Nielsen, which I match to local covariates from the American Community Survey. For each product module in each local market, I predict change in market size based on local change in age and income distributions. Some cities like San Francisco have experienced an increasing share of high-income and young households, while other cities like New Orleans have become poorer, with a decline in overall population. I then compare the change in fixed basket inflation across product module$\times$ local market cells with increasing or decreasing predicted market size. To control for supply factors, I include fixed effects: local market fixed effects control for local scale effects, while product group fixed effects control for national trends in inflation. In this setting, the identification assumption is that, conditional on the fixed effects, the direct effects of local changes in the age and income distributions on the equilibrium is only through demand.

I use 18 covariates X_{it} (all expressed in logs): total number of households, total population, total female population, total male population, total number of households in age $\times$ income groups (considering four age groups - below 25, 25 to 44, 45 to 64, above 65 -, and dividing each group into 3 income groups - below \$30k, \$30k to \$100k, above \$100k), median household income and mean household income.

Formally, I consider two periods, 2004-2006 and 2011-2013, and I predict (log) local total expenditures Q_{MIT} with local market covariates and fixed effects:

$$Q_{MIT} = \beta^M \cdot 1_M \cdot X_{IT} + \gamma_{IT} + \delta_{GT} + \epsilon_{IMT}$$

where M denotes the product module, G the product group, I the local market, and T the period. Note that the β^M coefficients are allowed to freely vary across modules (i.e. some modules will be very responsive to the number of low-income households, others more responsive to the number of high-income and old households, etc.). I estimate this specification in $2004-2006$ and predict market size out-of-sample in $2011-2013$ (the R^2 is very high: see Appendix D).

The predictor of residual market size growth between two periods is therefore $\beta^m \cdot 1_M \cdot (X_{IT_2} - X_{IT_1})$. Finally, I run specifications of the form:

$$Y_{MI} = \alpha \left[\beta^m \cdot 1_M (X_{IT_2} - X_{IT_1})\right] + \widetilde{\gamma_I} + \widetilde{\delta_G} + \widetilde{\epsilon_{IM}}$$

Figure 27 and the table below show that the relationship is very stable and very similar to the results found from the variation at the national level.

Table 25: Causal Effects of Changes in Market Size (Local Level)

	Difference in Fixed Basket Inflation Rate (pp)			
Predicted Increase in Spending (%)	-0.1471*** (0.0162)	-0.1276*** (0.0172)	-0.1271*** (0.0188)	-0.1276*** (0.0259)
F.E.	Department	Product Group	Product Group	Product Group
Weights	Yes	Yes	No	Yes
Cluster	Local Market	Local Market	Local Market	Product Module

Figure 27: Causal Effects of Changes in Market Size (Local Level)

A. Predicted and Actual Increase in Market Size

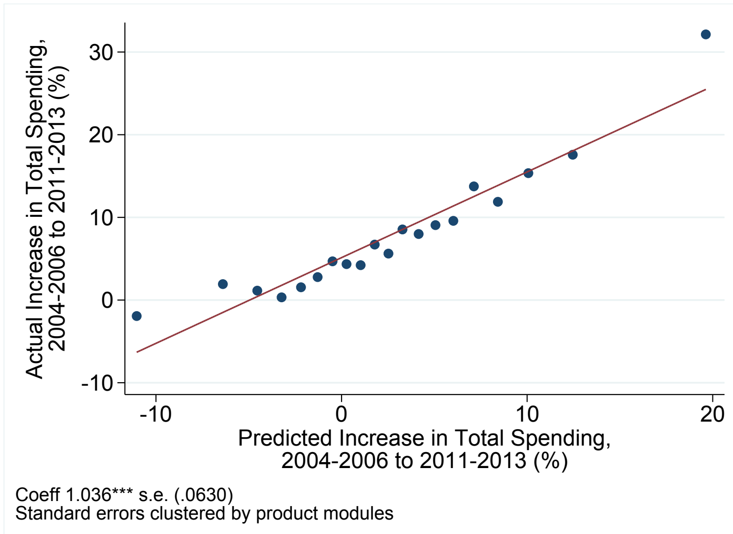

B. Higher Market Size Leads to Lower Inflation (Fixed Basket)

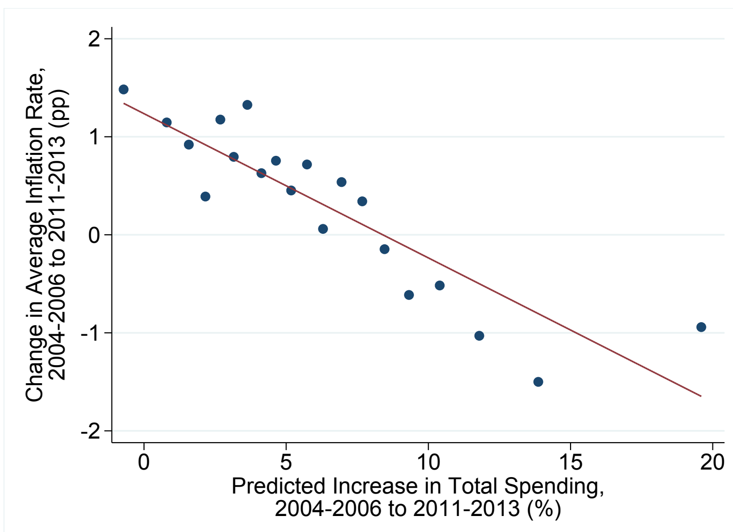

Figure 28: Geography Design: Predicting Market Size

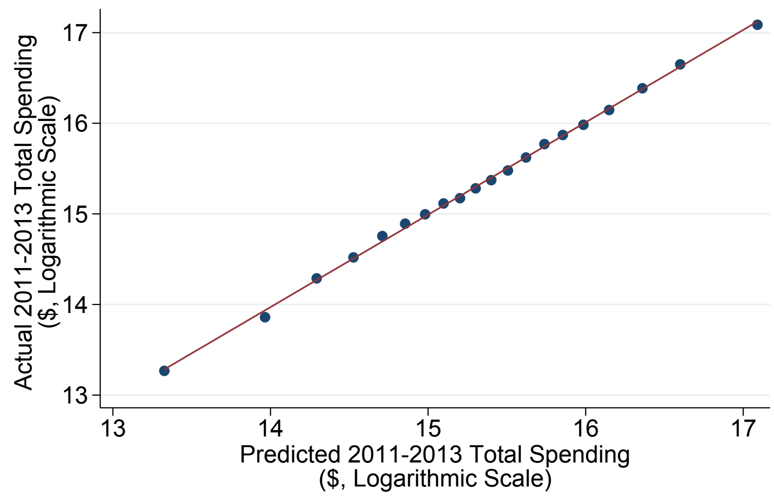

Note: Out-of-sample prediction using the prediction function estimated on 2004-2006 data and partialling out scantrack market and product group fixed effects

Figure 29: Geography Design: Predicted Market Size Growth and Fixed-Basket Inflation

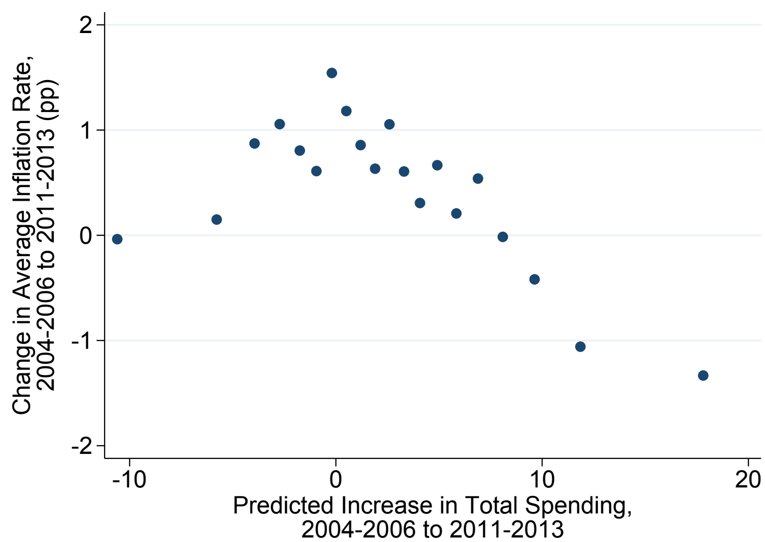

Additional Evidence on Changes vs. Level

Figure 30: Do Product Innovations Follow Market Size or Change in Market Size?

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

SNAP Research Design

Figure 31: Changes in SNAP Take-up Rate and Total Enrollment over Time

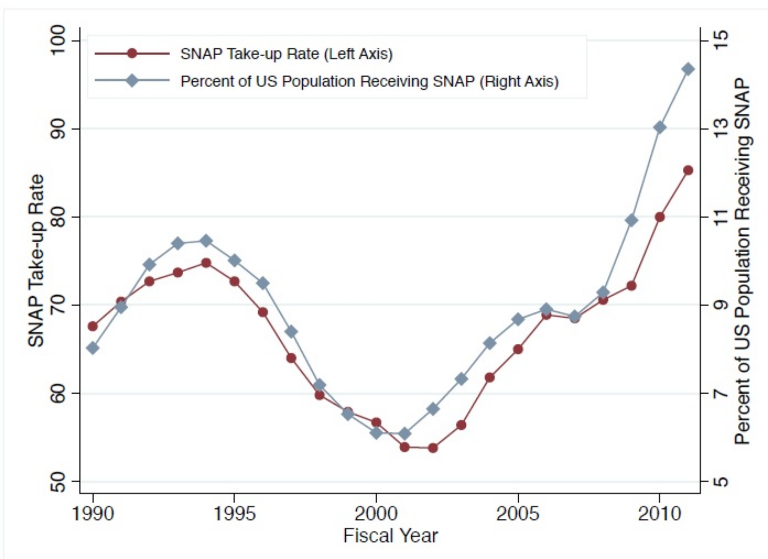

Source: Ganong and Liebman (2015)

The remainder of Appendix D is available from the author upon request.