

THE UNEQUAL GAINS FROM PRODUCT INNOVATIONS: EVIDENCE FROM THE U.S. RETAIL SECTOR*

XAVIER JARAVEL

This article examines how product innovations led to inflation inequality in the United States from 2004 to 2015. Using scanner data from the retail sector, I find that annual inflation for retail products was 0.661 (std. err. 0.0535) percentage points higher for the bottom income quintile relative to the top income quintile. When including changes in product variety over time, this difference increases to 0.8846 (std. err. 0.0739) percentage points a year. In CEX-CPI data covering the full consumption basket, the annual inflation difference is 0.368 (std. err. 0.0502) percentage points. I investigate the following hypothesis: (i) the relative demand for products consumed by high-income households increased because of growth and rising inequality; (ii) in response, firms introduced more new products catering to such households; (iii) as a result, the prices of continuing products in these market segments fell due to increased competitive pressure. Using a shift-share research design, I find causal evidence that increasing relative demand leads to increasing product variety and lower inflation for continuing products. A calibration indicates that the hypothesized channel accounts for a large fraction (over 50%) of observed inflation inequality. *JEL* Codes: E31, I31, I32, O30, O31, O33.

The capitalist achievement does not typically consist in providing more silk stockings for queens but in bringing them within the reach of factory girls in return for steadily decreasing amounts of effort...

*I am particularly indebted to Daron Acemoglu, Philippe Aghion, Raj Chetty, Edward Glaeser, Lawrence Katz, Pete Klenow, Josh Lerner, Marc Melitz, and Andrei Shleifer for extensive advice and comments. For thoughtful suggestions, I thank Katharine Abraham, Michael Boskin, Gabriel Chodorow-Reich, John Coglianesse, Anastassia Pedyk, Matt Gentzkow, Jessie Handbury, Erik Hurst, and Leslie McGranahan. Hua Jin provided excellent research assistance. I thank numerous seminar participants at Berkeley, Chicago Booth, the Collège de France, the Econometric Society, Harvard, the Kauffman Foundation, LSE, MIT, NBER, Princeton, SED, Stanford, and Yale for useful comments. I gratefully acknowledge financial support from the HBS doctoral office, the Kauffman Foundation, the Stanford Institute for Economic Policy Research, and the Washington Center for Equitable Growth. The analysis in this paper is based in part on data from The Nielsen Company (US), LLC and marketing databases provided through the Nielsen Datasets at the Kilts Center for Marketing Data Center at The University of Chicago Booth School of Business. The conclusions drawn from the Nielsen data are those of the author alone and do not reflect the views of Nielsen. Nielsen is not responsible for, had no role in, and was not involved in analyzing and preparing the results reported herein.

© The Author(s) 2018. Published by Oxford University Press on behalf of the President and Fellows of Harvard College. All rights reserved. For Permissions, please email: journals.permissions@oup.com

The Quarterly Journal of Economics (2019), 715–783. doi:10.1093/qje/qjy031.
Advance Access publication on December 4, 2018.

The capitalist process, not by coincidence but by virtue of its mechanism, progressively raises the standard of life of the masses.

—Joseph Schumpeter, *Capitalism, Socialism and Democracy*

I. INTRODUCTION

Who benefits from product innovations? A long tradition emphasizes the notion that everyone benefits from innovation thanks to the “product cycle” (e.g., Hayek 1931; Schumpeter 1942; Vernon 1966; Matsuyama 2002). Innovation may initially be aimed at the rich, who buy first; but soon enough a trickle-down process brings to the mass market the new products that were originally enjoyed by a select few. Historical examples abound, from automobiles to street lighting. But is the product cycle the only important force at play? Many product markets target different populations and are clearly segmented, such that there is no product cycle between them (e.g., scotch and tobacco). It could be the case that product innovations affect purchasing power inequality by increasing the variety and quality of goods available in specific consumer segments, as well as by driving down the price of existing products in these market segments via increased competitive pressure.

Schmookler (1966) pointed out that innovations are often inspired by specific demand. In the context of economic growth and rising income inequality, demand grows faster for premium products. As a result, the endogenous entry of new products (targeting product segments with increased demand) may lead to an increase in product variety and potentially reduced prices for existing products in these fast-growing premium categories, which are predominantly consumed by high-income households. In this article, I conduct several empirical tests showing the importance of this channel, primarily by using barcode-level scanner data from the U.S. retail sector in recent years.¹

The article starts by establishing that higher-income households experienced a faster increase in product variety and lower

1. The U.S. food industry illustrates particularly well some of the core ideas developed in this article. Organic food sales have grown at an average annualized rate of 11.2% between 2004 and 2015, compared with 2.8% for total food sales, in the context of increasing demand from higher-income households. The price premium for organic products shrank significantly: for instance, organic spinach cost 60% more than nonorganic spinach in 2004, compared with only 7% more in 2015. Low inflation for organic products brought down the food CPI, which reduced the rate of increase in food stamps through indexation, although most food stamp recipients do not purchase organic products.

inflation in the U.S. retail sector from 2004 to 2015 (consistent with [Argente and Lee 2016](#); [Kaplan and Schulhofer-Wohl 2017](#)). The magnitude of these effects is large: using a Törnqvist index, annual inflation for retail products was 0.661 (std. err. 0.0535) percentage points higher for the bottom income quintile relative to the top income quintile. When accounting for changes in product variety over time with a Constant Elasticity of Substitution (CES) price index, annual inflation inequality increases to 0.8846 (std. err. 0.0739) percentage points. These results are very stable for a wide variety of price indexes and hold before, during, and after the Great Recession, both across and within product categories.²

The scanner data used for this analysis, provided by the Nielsen Company, are representative of a large subset of products within the retail sector and account for a sizable fraction of households' expenditures ([Section II](#) provides a detailed analysis of expenditure coverage). Using a matched data set from the Consumer Expenditure Survey (CEX) and Consumer Price Index (CPI) covering the full consumption basket, I find that the inflation difference between the bottom and top income quintiles was 0.368 (std. err. 0.0502) percentage points a year (consistent with [McGranahan and Paulson 2006](#)). This result suggests that similar patterns hold within and outside of the sample covered by the Nielsen data.

Next, I examine whether the equilibrium response of supply to faster growth in demand from high-income consumers can explain the patterns of differential inflation and increase in product variety. It is well documented that in recent decades the share of U.S. national income accruing to high-income consumers has steadily increased, both because more and more households entered high-income brackets as the economy grew and because of rising income inequality (e.g., [Piketty and Saez 2001](#); [Autor, Katz, and Kearney 2008](#); [Kopczuk, Saez, and Song 2010](#)). Intuitively, firms can respond to changes in relative market size by skewing product introductions toward market segments that are growing faster. This process can lead to a decrease in the price of existing products in the fast-growing market segments because increased competitive pressure from new products pushes markups down. I investigate this hypothesis, which I view as the

2. To keep the analysis tractable, I posit the existence of separate homothetic price indexes for each income quintile; overcoming this limitation and fully accounting for nonhomotheticities is an important direction for future work.

main contribution of this article, by leveraging the rich scanner data available for the retail sector.

A variety of patterns in the data support the hypothesized channel. Product categories that grow faster indeed feature a greater increase in product variety, lower inflation, and disproportionately cater to higher-income households. To address potential reverse causality or omitted variable biases, I develop a shift-share research design (broadly similar to [Acemoglu and Linn 2004](#) and [DellaVigna and Pollet 2007](#)). This design relies on two components: predetermined spending shares across the product space for a large number of sociodemographic groups and heterogeneity in the population growth rates for these various groups during the sample period. Age groups, education group, racial groups, and regional populations all have different budget shares across the product space; variation in the size of these groups over time generates changes in demand. Spending profiles across the product space are measured in the initial period and kept constant, such that the variation in the shift-share instrument comes entirely from changes in the size of the sociodemographic groups over time.

The IV estimates from the shift-share design indicate that increases in demand lead to a substantial fall in prices and increases in product variety. When the growth rate of demand increases by 1 percentage point, the inflation rate for products available in consecutive years falls by 0.42 percentage points (std. err. 0.139). Accounting for changes in product variety, inflation falls by 0.62 percentage points (std. err. 0.258). To shed light on the mechanism, I amend the shift-share research design (exploiting additional variation across states) and test whether the supply response is driven by increasing demand or merely by a higher (but stagnating) level of demand. I find that increasing demand drives the effect. Finally, I show that the estimated market size effects on inflation for continuing products can largely be explained by a change in markups. Both reduced-form evidence on store markups and a structural approach following [Hottman, Redding, and Weinstein \(2016\)](#) indicate that changes in markups explain a substantial fraction (between 50% and 100%) of the price response to increasing market size.

The article concludes with a calibration that assesses the extent to which inflation inequality results from the endogenous response of supply to increasing demand. Because of growth and increasing inequality, changes in the income distribution

during the sample period induced changes in demand that were asymmetric across the product space. The calibration considers changes in demand resulting from changes in the income distribution across detailed cells of the product space (product modules by price deciles). Based on the IV estimates from the shift-share research design, the calibration shows that these changes in demand and their induced price effects are sufficiently large to explain much of the observed inflation inequality. The predicted Törnqvist inflation difference is 0.34 percentage points a year, which is 70.8% of the observed inflation difference at the same level of aggregation (product modules by price deciles) and 51.5% of the full observed inflation difference. The patterns are very similar for the CES inflation difference when also accounting for changes in product variety. The predicted CES inflation difference is 0.49 percentage points, which is 74.2% of the benchmark at the same level of aggregation and 55.6% of the full CES inflation difference.

This article relates and contributes to two main literatures. First, a vast literature has studied the implications of innovations for inequality. Although most of this literature has studied skill-biased factor-augmenting technical change (see [Violante 2008](#) for a survey), my article investigates the distributional effects of innovations in the product market. The idea that larger markets offer benefits to consumers through increased product variety and potentially lower prices for these varieties goes back to the seminal work of [Dixit and Stiglitz \(1977\)](#), [Krugman \(1979\)](#), [Shleifer \(1986\)](#), and [Romer \(1990\)](#). More recently, a few papers have examined how demand-side forces determine the direction of innovation across sectors. [Acemoglu et al. \(2012\)](#), [Boppart and Weiss \(2013\)](#), and [Comin, Lashkari, and Mestieri \(2017\)](#) study this question primarily theoretically. Empirically, [Acemoglu and Linn \(2004\)](#) provide estimates of the effect of changes in market size on the entry of new drugs, and [Broda and Weinstein \(2010\)](#) show that net product creation is procyclical. Relative to this literature, the contribution of my article is to study the implications of endogenous innovations across sectors for inequality.³ Two recent papers

3. Adapting the concepts from [Acemoglu \(2007\)](#) to sector-augmenting technical change, I test both the “weak bias” and “strong bias” hypotheses for technical change across sectors; namely, when demand for a sector becomes relatively more abundant, does product entry endogenously increase in this sector (weak bias)? And is this effect sufficiently strong such that the observed relative supply curves for goods are downward-sloping (strong bias)? The scanner data provide a positive answer to these questions.

document that consumers benefit unequally from innovations and firm dynamics: [Eizenberg \(2014\)](#) shows that the rapid innovations in CPUs mostly benefited the 20% least price-sensitive consumers; and [Faber and Fally \(2017\)](#) find that the most productive firms endogenously target wealthier households.

Second, a long-standing literature has investigated patterns of inflation inequality. [Amble and Stewart \(1994\)](#), [Garner, Johnson, and Kokoski \(1996\)](#), [Hobijn and Lagakos \(2005\)](#), and [McGranahan and Paulson \(2006\)](#) measure inflation across household groups in the United States using CEX-CPI data covering the full consumption basket. In more recent work, [Broda and Romalis \(2009\)](#), [Argente and Lee \(2016\)](#), and [Kaplan and Schulhofer-Wohl \(2017\)](#) document inflation inequality using scanner data. Relative to this literature, my article refines and extends existing estimates of inflation inequality, and it proposes a new mechanism to explain them: directed product innovations.

The remainder of this article is organized as follows. [Section II](#) describes the data sources and summary statistics. [Section III](#) presents measures of inflation inequality across quintiles of the income distribution. [Section IV](#) estimates the response of supply to changes in demand. [Section V](#) presents the calibration, establishing that much of the observed inflation inequality measured in [Section III](#) can be explained by the supply dynamics estimated in [Section IV](#).

II. DATA AND SUMMARY STATISTICS

This section describes the data sources, defines the samples and key variables used in the analysis, and presents summary statistics.

II.A. Data Sources, Samples, and Variable Definitions

1. Scanner Data. The analysis is primarily based on the Nielsen Homescan Consumer Panel and Nielsen Retail Scanner data sets. The Nielsen Homescan Consumer Panel (HMS) records consumption from 2004 to 2015 for a rotating panel of about 50,000 households, who are instructed to scan any product they purchase that has a barcode.⁴ These products are typically found in department stores, grocery stores, drug stores, convenience stores, and other similar retail outlets. The HMS

4. The sample includes about 55,000 households in the panel from 2007 onward and around 40,000 prior to 2007.

data have the key benefit of providing 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.

The Nielsen Retail Scanner data set (RMS) records consumption from 2006 to 2015 in more than 100 retail chains across all U.S. markets. The database includes more than 45,000 individual stores. A key advantage of this data set is that it provides a better measure of changes in product variety across the product space, as sampling error is much smaller than in the HMS data.

All products in the HMS and RMS data are classified into broad “departments” (dry grocery, general merchandise, health and beauty care, alcoholic beverages, deli, etc.), which are themselves subdivided into detailed “product groups” (grooming aids, soup, beer, pet care, kitchen gadgets, etc.) and very detailed “product modules” (ricotta cheese, pet litter liners, bathroom scales, tomato puree, women’s hair coloring, etc.). Over time, Nielsen expanded coverage of certain product modules (for instance, in-store baked goods). I only keep product modules that are available throughout the HMS/RMS samples, which leaves me with 10 departments, 112 product groups, and 1,042 product modules in the finalized HMS/RMS samples.

Finally, to measure manufacturer entry and competition, I follow [Hottman, Redding, and Weinstein \(2016\)](#) and match manufacturer identifiers from GS1, the company in charge of allocating bar codes in the United States, to the UPC codes in the HMS and RMS samples. The match rate is 95%.

2. CEX-CPI Data. The CEX is a widely used consumption survey tracking spending in all product categories, including goods, services, housing, and health. It consists of two parts, the interview and diary surveys, which I use in combination; both provide information on household characteristics including income. The strengths and weaknesses of the CEX are well understood (e.g., [Garner, McClelland, and Passero 2009](#)). Relative to the scanner data, the CEX has the benefit of much larger coverage; but it only provides information on spending, not on quantities and prices separately. To obtain information on price changes, I match by hand the 650 detailed product categories in the CEX to the most disaggregated Consumer Price Index (CPI) data series made available by the Bureau of Labor Statistics. Carrying out this exercise from 2004 to 2015 to cover the same period as the scanner

data, I obtain a matched data set (CEX-CPI) recording inflation and spending patterns for 256 detailed product categories that span the full consumption basket.

3. *Choice of Income Groups.* In the HMS data, households are asked to report how much they earn using discrete bins across the income distribution. In the main analysis, I focus on inflation patterns across five groups corresponding approximately to the quintiles of the household income distribution. Between 2004 and 2015, the thresholds separating the five quintiles of the household income distribution were approximately \$20,000, \$40,000, \$60,000, and \$100,000, on average. I use these thresholds to assign the panelists in the HMS and CEX-CPI data sets to income quintiles. For robustness, I repeat the analysis across income deciles.

4. *Price Indexes.* Following Broda and Romalis (2009), I assume the existence of groups of individuals with different homothetic preferences; in my setting, the groups are given by income quintiles. This approach provides a characterization of how price changes for continuing products and changes in product variety affect welfare across income quintiles. Conceptually, the change in the price index for income quintile i between t and $t + 1$, denoted $\Pi_{t,t+1}^i$, gives the compensating variation (as a percentage of nominal income) that one would need to give to consumers at $t + 1$ to bring them back to the same level of utility as at time t .⁵

The change in the price index (specific to each income quintile) from t to $t + 1$ can be expressed as the combination of two components: (i) a weighted average of the price changes for products that are available at both t and $t + 1$; and (ii) an adjustment reflecting the willingness to pay for the products available at $t + 1$ but not at t (and vice versa). I build these two components as follows.

5. Positing the existence of separate price indexes for each income quintile implies that preferences jump discontinuously at the quintile thresholds. This feature would be concerning if households were often crossing the income quintile threshold from one year to the next; but in practice it happens only for a small fraction of panelists. A related limitation is that the approach I take (following Broda and Romalis 2009) requires assuming homothetic preferences within each income quintile. This approach has the advantage of tractability, but it does not account for the fact that, as real income changes, people's preferences may also change, that is, it does not fully account for nonhomotheticities.

5. *Computing Inflation for Continuing Products.* In the main analysis, I compute the inflation rate for continuing products (available at both t and $t + 1$) using the Törnqvist index:

$$1 + \pi_{t,t+1}^i = \prod_{k=1}^n \left(\frac{p_{k,t+1}^i}{p_{k,t}^i} \right)^{\frac{s_{k,t}^i + s_{k,t+1}^i}{2}} = \exp \left(\sum_{k=1}^n \frac{s_{k,t}^i + s_{k,t+1}^i}{2} \cdot \log \left(\frac{p_{k,t+1}^i}{p_{k,t}^i} \right) \right),$$

where i indexes income quintiles, k products, and t years; $s_{k,t}^i$ is the spending share of income quintile i on product k in year t ; and $p_{k,t}^i$ is the average price (with quantity weights) paid by income quintile i for product k in year t . The Törnqvist index is a natural benchmark because it is a superlative price index, providing a second-order approximation to any twice continuously differentiable, homothetic expenditure function (e.g., [Diewert 1976](#)). For robustness, I consider other standard price indexes.

Note that the spending shares $s_{k,t}^i$ are updated each year, that is, the price index is “chained,” which provides a better approximation to the change in the expenditure function. In the HMS and RMS data, a product k is a barcode, while in the CEX-CPI data it refers to one of the 256 detailed product categories. In the HMS data, the price paid $p_{k,t}^i$ can freely vary across income quintiles; for instance, they may purchase the same barcode in different stores, and certain quintiles may use coupons more often. In contrast, in the RMS data the measured price is the average price (with quantity weights) for all consumers purchasing the barcode. In the CEX-CPI data, there is no price variation across income quintiles within a detailed product category; moreover, due to the methodology of the Bureau of Labor Statistics, inflation reflects changes in posted prices rather than in the effective prices actually paid by households.

6. *Changes in Product Variety.* As pointed out in [Handbury and Weinstein \(2014\)](#), the HMS data are not well suited to the measurement of changes in product variety across household groups. Because only products purchased by panelists appear in this data set, there may be a mechanical tendency to find larger increases in product variety among household groups with faster expenditure growth. Accordingly, I use the HMS data in combination with the RMS data to compute spending on “new” and “exiting” goods across income quintiles.

In the baseline analysis, a new good in year t is defined as a good that never appeared in the HMS or RMS data at any point before year t ; conversely, an exiting good in year t is defined as a

good that never appeared in the HMS or RMS data at any point after year t . Then, for each income quintile i in each year t , I compute the spending shares on new ($s_{N,t}^i$) and exiting products ($s_{E,t}^i$).

Structural assumptions on the demand system are required to adjust the price index for willingness to pay for new and exiting products. I introduce and discuss these assumptions in [Section III](#). For robustness, I also consider alternative definitions of new and exiting products, defining “new” relative to the previous year and “exiting” relative to the following year (i.e., without considering availability in other years in the sample).

7. Construction of Price Deciles. For several parts of the analysis, it will prove useful to segment the product space at a level of disaggregation finer than the 1,042 products modules provided by Nielsen. I segment each product module into 10 price deciles, which can be viewed as an intuitive proxy for the quality ladder within each product module. In each year t , each product k is assigned to a price decile as follows: compute the average (quantity-weighted) unit price for product k in year t ; standardize the average price by the weight of the item (in ounces) whenever it is provided by Nielsen (e.g., weight is available for most product modules within food); assign each product to a price decile (within the product module) in year t based on its standardized average unit price in that year.

Thus, within each price decile by product module I can compute inflation for continuing goods as well as spending on new and exiting products. In the baseline analysis, I characterize inflation for continuing goods from t to $t + 1$ across deciles based on the price deciles assigned to products in the base year t . Spending on new goods (available at $t + 1$ but not t) uses the price deciles assigned at $t + 1$ while spending on exiting goods (available at t but not $t + 1$) uses the price deciles assigned at t . For robustness, I consider an alternative definition of price deciles, assigning each product that exists for two consecutive years to a decile based on the weighted average unit price over the full two-year period instead of the base period only.

II.B. Summary Statistics

[Table I](#) provides summary statistics on expenditures in the HMS, RMS, and CEX data.

Panel A focuses on the HMS and RMS samples, reporting the share of expenditures and barcodes across the 10 broad

TABLE I
SUMMARY STATISTICS ON EXPENDITURES COVERAGE

Department	Examples of product groups	Expenditure Share		Barcode Share	
		HMS	RMS	HMS	RMS
Panel A: Distribution of expenditures and barcodes in Nielsen data					
Alcoholic Beverages	Beer, liquor, wine	3.76	5.42	3.14	4.98
Dairy	Butter and margarine, cheese, milk, yogurt	9.01	8.06	4.00	2.83
Deli	Deli dressings, deli salads	2.04	2.14	1.76	1.69
Dry grocery	Baby food, bread, breakfast food, carbonated beverages	39.68	38.41	33.87	28.00
Fresh produce	Fresh produce	2.83	2.92	1.44	1.52
Frozen food	Frozen baked goods, ice cream, frozen drinks, frozen pizza	8.92	7.14	5.57	4.21
General merchandise	Cookware, electronics, gardening, household supplies	6.66	6.26	18.09	25.46
Health and beauty aids	Deodorant, diet aids, hair care, shaving needs, skin care	10.30	14.21	17.32	19.21
Nonfood grocery	Detergents, disposable diapers, laundry supplies, tobacco	13.40	12.42	13.26	10.84
Packaged meat	Fresh meat, deli packaged meat	3.40	3.02	1.56	1.28

TABLE I
CONTINUED

Sample	All Households, \$/HH (% of CEX, All)	By income quintile, \$/HH (% of CEX, All)				
		(1)	(2)	(3)	(4)	(5)
Panel B: Comparison of expenditures in Nielsen data and CEX						
(A) CEX, all	52,187 (100%)	29,906 (100%)	37,206 (100%)	51,478 (100%)	91,156 (100%)	
(B) CEX, all except shelter	40,323 (77.27%)	23,079 (77.17%)	28,885 (77.64%)	40,048 (77.8%)	70,393 (77.22%)	
(C) CEX, all except shelter, healthcare, transportation, services, utilities	16,439 (31.5%)	9,735 (32.55%)	11,785 (31.67%)	16,109 (31.29%)	27,896 (30.6%)	
(D) CEX, food and drinks at home, housekeeping supplies, and household cleaning products	4,474 (8.57%)	2,502 (12.04%)	3,459 (9.3%)	4,274 (8.3%)	6,097 (6.69%)	
(E) CEX, Food and drinks at home, housekeeping supplies, cleaning products, personal care products, smoking products, tableware, tools, nonelectric cookware, and apparel	6,987 (13.39%)	4,604 (15.39%)	5,280 (14.19%)	6,745 (13.1%)	10,162 (11.15%)	
(F) Nielsen (HMS), all	3,503 (6.71%)	3,191 (10.67%)	3,563 (9.58%)	3,858 (7.49%)	4,126 (4.53%)	
(G) Nielsen (HMS), food and drinks at home	2,440 (4.67%)	2,217 (7.41%)	2,481 (6.67%)	2,686 (5.22%)	2,882 (3.16%)	

Notes. This table reports summary statistics in the HMS, RMS, and CEX samples. The HMS and CEX samples extend from 2004 to 2015, and RMS is available from 2006 to 2015.

departments covered in these data. The table shows that spending patterns are very similar in the HMS and RMS data. Food products account for a substantial fraction of overall spending. It is important to note that the Nielsen data cover a very selected set of products, such that the results of the analysis may not generalize to all product categories.

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 technological innovations in recent decades. However, these product categories are characterized by a sizable rate of increase in product variety, as documented in [Section III](#). The Nielsen data are therefore useful to study one particular manifestation of innovation, increasing product variety, and to assess its benefits for households across the income distribution.

[Table I](#), Panel B reports various statistics that help assess the extent of expenditures coverage in the Nielsen data, on average and across income quintiles.⁶ Row (A) shows that average spending per capita is about \$52,000 in the CEX data, ranging from \$20,785 for the bottom income quintile to \$91,156 for the top income quintile. Rows (B) and (C) report how expenditures coverage declines when major spending categories (not covered in Nielsen) are excluded. Coverage falls to 77.27% without shelter, and excluding healthcare, transportation, services, and utilities reduces it further to 31.5%.

Row (D) focuses on expenditure categories that are likely to be covered accurately in the Nielsen data: food and drinks at home, housekeeping supplies, and household cleaning products account for 8.57% of overall spending in CEX on average, ranging from 12.04% for the bottom income quintile to 6.69% for the top income quintile. Row (E) repeats this calculation for a more extensive set of product categories within CEX that also includes personal care products, smoking products, tableware, tools,

6. Although the CEX is a natural benchmark to assess expenditure coverage in Nielsen across income groups, CEX may well understate overall consumption spending. For instance, according to the Bureau of Economic Analysis (BEA), average annual consumption per household was approximately \$85,000 during the sample period. This discrepancy is in part due to the fact that the BEA consumption measure includes purchases by nonprofits, by the military and various institutions, as well as employer-paid insurance. [Garner, McClelland, and Passaro \(2009\)](#) provide a detailed discussion of these issues.

nonelectric cookware, and apparel. The Nielsen data offer some coverage of all of these categories but it may be imperfect. The statistics reported in row (E) are thus likely to constitute an upper bound for spending coverage in Nielsen. The table reports that these product categories account for 13.39% of overall spending on average, which corresponds to close to 40% of expenditures on goods given that goods accounted for about 34% of overall spending during this period.

Row (F) provides another instructive calculation of expenditures coverage in Nielsen data by directly computing spending per capita in the HMS sample. I use the panelists' projection weights provided by Nielsen, so that the panelists are representative of the full U.S. household population. This row reports that average spending per capita in HMS accounts for 6.71% of average spending in CEX, ranging from 13.19% for the bottom income quintile to 4.53% for the top income quintile. These numbers are likely to provide a lower bound for expenditure coverage in Nielsen, because it could be the case that the Homescan Consumer Panel is representative of all household spending in categories in which it has some coverage (bringing coverage closer to rows (D) or (E)). Finally, row (G) considers only food and drinks at home within HMS, which yields coverage equal to about two-thirds of row (F).

In sum, the Nielsen data account for a nontrivial share of spending, but the sample is far from providing comprehensive coverage of consumption. Although the analysis in the next sections primarily relies on the HMS and RMS samples, I also use the CEX-CPI data to provide suggestive evidence that some of the patterns may hold across a broader set of consumption categories.

III. INFLATION INEQUALITY ACROSS INCOME QUINTILES

This section takes a series of steps to estimate inflation rates across income quintiles between 2004 and 2015, first focusing on continuing products and then accounting for the welfare gains from increasing product variety.

III.A. Inflation across Income Groups for Products Available in Consecutive Years

For each income quintile, I compute the Törnqvist index for continuing products (introduced in [Section II](#)). These products are available at both t and $t + 1$ and account for about 90% of spending

on average. The key assumption underlying this computation is that the quality of a given product is constant over time.⁷

Figure I, Panel A shows a clear pattern of declining inflation across income quintiles. The sample is the HMS data. Between 2004 and 2015, for the fifth income quintile average annual Törnqvist inflation for continuing products was 1.21% (std. err. 0.031), while for the first income quintile it was 1.87% (std. err. 0.045). Note that measured inflation is subject to sampling uncertainty. Figure I, Panel A reports the 95% confidence intervals, obtained by bootstrapping, and shows that in the HMS sample inflation is estimated quite precisely for each income quintile. The OLS best-fit line shows that one cannot reject a pattern of linearly decreasing inflation rates across the five income quintiles.

Figure I, Panel B examines the extent to which the choice of price index affects the estimated inflation difference between the bottom and top income quintiles. The panel reports the inflation difference along with 95% confidence intervals for Törnqvist and four other standard price indexes: Laspeyres, Paasche, CES, and nested CES. These price indexes all handle substitution patterns (from t to $t + 1$) differently. By using base period (t) spending shares, Laspeyres makes substitution impossible. By using end period ($t + 1$) spending shares, Paasche allows for an extreme form of substitution. The CES and nested CES price indexes are widely used price indexes derived from the corresponding utility functions. I use three nests in nested CES: departments, product groups, and product modules. The mathematical formulas for each of these price indexes are standard and are reported in [Online Appendix A](#).

Figure I, Panel B shows that inflation inequality among continuing products is very similar across price indexes. The average annual inflation difference is 0.661 (std. err. 0.0535) percentage point for Törnqvist, compared with 0.669 (std. err. 0.0520) for Laspeyres, 0.657 (std. err. 0.0596) for Paasche, 0.651 (std. err. 0.0531) for CES, and 0.6305 (std. err. 0.0474) for nested CES. It is comforting that the results are very stable across price indexes

7. The assumption that quality is constant at the UPC level is supported by institutional details: GS1 recommends using the same barcode only for exactly the same products, and the inventory management systems used by retailers would be disrupted if different products had the same barcode. It is also in line with the fact that the set of available characteristics (flavor, label, and scent) are stable within barcode over time.

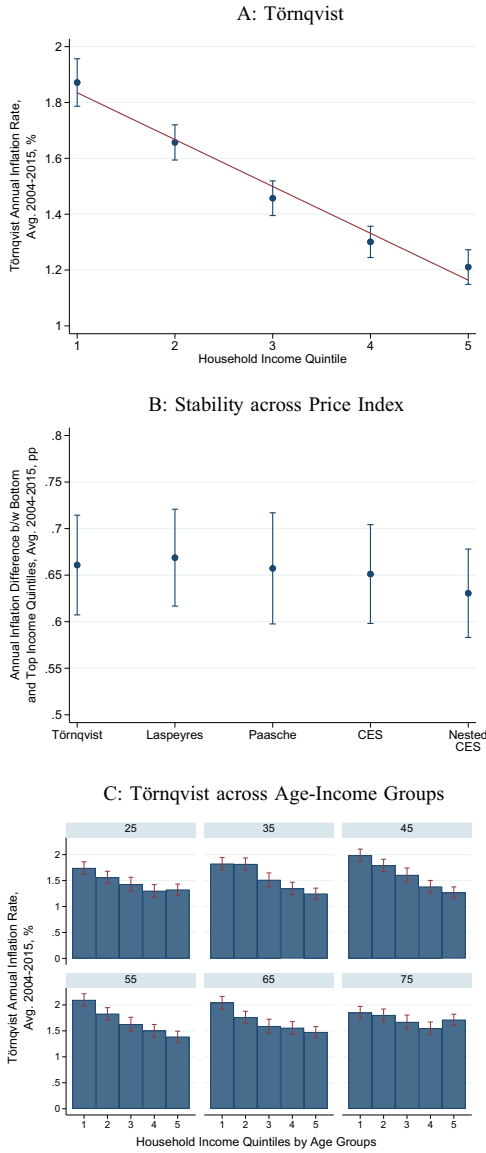


FIGURE I

Inflation for Continuing Products across Income Groups

Panels A, B, and C report the average annual inflation rate for various household groups in the HMS data from 2004 to 2015. Continuing products are observed across consecutive years. The 95% confidence intervals are obtained by bootstrap.

that assume different types of substitution effects. It indicates that both qualitatively and quantitatively, the results for continuing products do not depend on the way substitution effects are handled.

Finally, [Figure I](#), Panel C repeats the analysis of inflation inequality across income quintiles within six age groups. The panel shows that inflation for continuing products is lower for higher-income groups within each age group, which illustrates the robustness of this finding.

III.B. Changes in Product Variety across Income Groups

Do welfare effects from increasing or decreasing product varieties also differ across income groups? [Figure II](#), Panel A reports simple reduced-form evidence suggesting that changes in product variety benefit higher-income households more.

First, the left side of this panel shows that spending shares on new products are higher for richer households, ranging from about 8% for households in the bottom income quintile to close to 10% for those in the top income quintile. At the same time, spending shares on exiting products are very similar across income quintiles, around 2%. So product variety is increasing faster for higher-income households.

Moreover, the right side of the panel shows that for every \$10,000 increase in the mean income of the consumers buying from a product module (using spending weights to compute mean consumer income), the share of spending on new products in this product module increases by 2.6 percentage points, a large change equal to approximately a third of the average share of spending on new products. The binned scatter plot, where each dot represents 1% of the data, shows the strength of the relationship between consumer income and introduction of new goods across the product space: the relationship is not driven by a few outlier categories.⁸

To obtain a measure of consumers' willingness to pay for changes in product variety, structural assumptions are needed.

8. Note that plotting the data in this way, through the lens of the product space rather than by directly looking at the consumption baskets of consumers of different income levels, has the key advantage that the product cycle will not mechanically generate differences across income groups. In other words, the fact that new products may first be purchased by higher-income consumers will not generate an increasing relationship between income and share of spending on new products, given that we are looking at patterns across product modules while the product cycle operates within product modules.

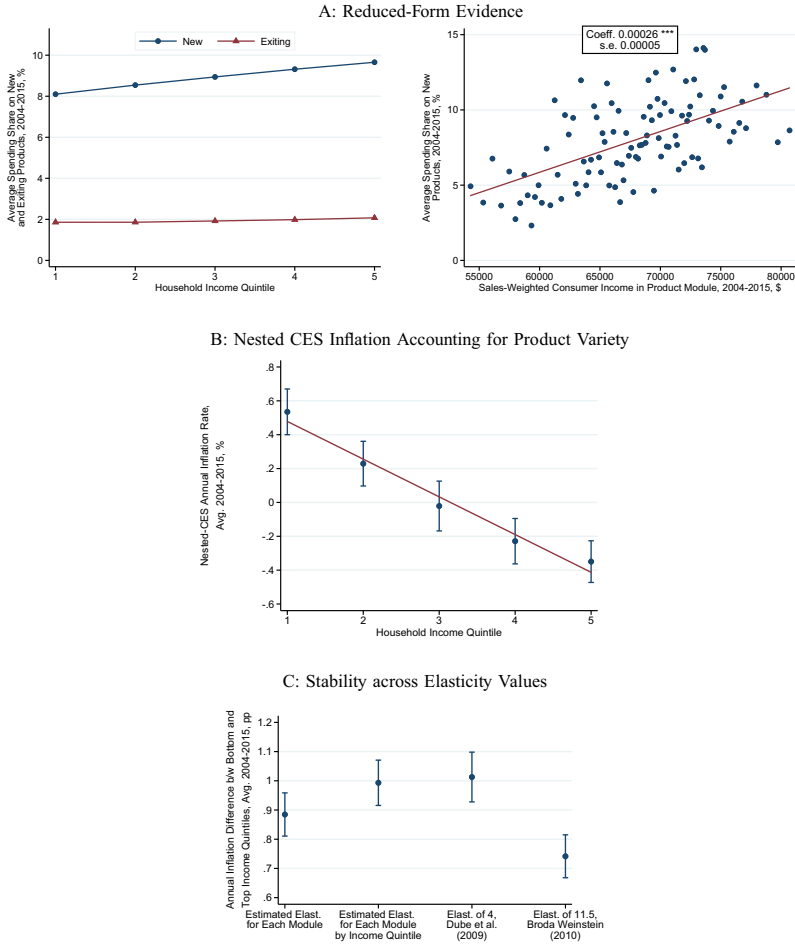


FIGURE II

Welfare Gains from Changes in Product Variety across Income Groups

Panel A reports spending on new and exiting products across household groups and across the product space; the regression uses spending weights and standard errors are clustered at the level of product modules. Panels B and C report inflation patterns accounting for changes in product variety using the nested CES price index. Spending patterns are measured in the HMS data, while product entry and exit is measured using the combined HMS and RMS data sets. The 95% confidence intervals are obtained by bootstrap.

Conceptually, by assuming a well-behaved utility function, if one knows the relevant demand elasticities one can infer the infra-marginal consumer surplus created or destroyed by changes in product variety from the observed spending shares on new and exiting products. A prominent choice in the literature is the CES utility function, following Feenstra (1994) and Broda and Weinstein (2006, 2010). With product entry and exit, the change in the exact CES price index from t to $t + 1$ for product module m and income

quintile i is given by: $\Pi_{t,t+1}^{CES,i,m} = \pi_{t,t+1}^{CES,i,m} \cdot \left(\frac{1-s_{N,t+1}^{i,m}}{1-s_{E,t}^{i,m}} \right)^{\frac{1}{\sigma_m-1}}$, where

$\pi_{t,t+1}^{CES,i,m}$ is the CES price index for continuing products; $s_{N,t+1}^{i,m}$ and $s_{E,t}^{i,m}$ are the spending shares on new and exiting products; and σ_m is the elasticity of substitution between products within module

m . In the remainder of the article, I refer to the term $\frac{1-s_{N,t+1}^{i,m}}{1-s_{E,t}^{i,m}}$ as the Feenstra ratio. This term shows that the higher the expenditure share of new products or the lower the expenditure share of exiting products, the smaller the exact price index ($\Pi_{t,t+1}^{CES,i,m}$) relative to the price index focusing on continuing products ($\pi_{t,t+1}^{CES,i,m}$).

The strength of the welfare effect from changes in product variety depends on the module-specific elasticity of substitution between varieties, σ_m . As σ_m grows, the term $\left(\frac{1-s_{N,t+1}^{i,m}}{1-s_{E,t}^{i,m}} \right)^{\frac{1}{\sigma_m-1}}$ converges to 1 and the inflation bias from product variety goes to 0. Intuitively, when existing varieties are close substitutes to new or disappearing varieties, a law of one price applies and price changes in the set of existing products perfectly reflect price changes for new and exiting varieties.

To compute the module-level price index with variety effects shown above, a high-dimensional set of elasticities of substitution $\{\sigma_m\}$ must be estimated. In the benchmark specification, all consumers are assumed to have the same elasticity of substitution. I follow the well-known approach of Feenstra (1994) and Broda and Weinstein (2006) to estimate these elasticities (for completeness, the estimation procedure is described in Online Appendix A). For robustness, I estimate elasticities $\{\sigma_m^i\}$ separately for each income quintile. Online Appendix Table A1 reports the distribution of the estimated elasticities (the 10th, 50th, and 90th percentiles are 3.5, 5.7, and 9.3) along with their standard errors. Online Appendix Figure A1 reports a histogram showing that

the elasticities vary substantially across product modules but are very similar across income quintiles.

In the baseline specification, I use common elasticities across income quintiles and I compute the nested CES inflation rate for each income quintile from the estimates of $\{\Pi_{t,t+1}^{CES,i,m}\}$. Ninety-five percent confidence intervals are obtained by bootstrapping the entire estimation procedure (in particular, accounting for the fact that the elasticities of substitution are estimated).

Figure II, Panel B presents the results. As for continuing products, inflation is clearly lower for higher-income groups and one cannot reject a linear decline in inflation across income quintiles. Between 2004 and 2015, for the fifth income quintile average annual nested CES inflation was -0.349% , compared with -0.228% for the fourth quintile, -0.021% for the third quintile, 0.229% for the second quintile, and 0.535% for the first quintile. Compared with inflation for continuing products alone, incorporating the adjustment for changes in product variety, inflation is lower for each income quintile. This finding is consistent with the results of Broda and Weinstein (2010): consumers value increasing product variety, which lowers overall inflation.

The comparison of Figures I and II show that the product variety adjustment benefits higher-income households relatively more. In the baseline specification with the product variety adjustment, the average annual nested CES inflation difference between the bottom and top income quintiles is 0.8846 (std. err. 0.0739) percentage points, while for continuing products it is only 0.6305 (std. err. 0.0474) percentage points. Thus, inflation inequality is about 40% stronger when changes in product variety are accounted for.

Figure II, Panel C examines the sensitivity of inflation inequality to the CES elasticities. In the baseline specification, the estimated elasticities vary across product modules but are common to all income quintiles. When the elasticities are also allowed to vary across income quintiles, the annual inflation difference increases to 0.9932 percentage points a year. I also consider the results when using a single elasticity across all product modules, taking the elasticities across the ranges of estimates from the literature. With an elasticity of 4 as in Dubé, Fox, and Su (2012), which is in the lower range of the literature, the annual inflation difference is 1.012 percentage points. With an elasticity of 11.5 as in Broda and Weinstein (2010), which is in the higher range of the literature, the difference falls to 0.7414 percentage points,

which is 17.5% larger than when changes in product variety are not accounted for.

These results show, not surprisingly, that the quantitative estimates of the impact of increasing product variety on inflation inequality vary substantially depending on the elasticity estimates. Intuitively, the lower the elasticities, the higher the amount of inframarginal consumer surplus created by increasing product variety. Despite this sensitivity to elasticity estimates, all specifications yield the robust conclusion that higher-income groups benefit more from increasing product variety. In that sense, the patterns of inflation inequality for continuing products can be viewed as a lower bound for overall inflation inequality.

The [Online Appendix](#) presents additional evidence on the robustness of these findings. [Online Appendix Figure A2](#) shows that there is more entry of new manufacturers (not just of new barcodes) in product modules or product modules by price deciles that sell to richer consumers. [Online Appendix Figure A3](#) uses another specification to address the two main limitations of the nested CES demand system. First, as pointed out by [Hausman \(2003\)](#), CES assumes infinite reservation prices and may overestimate the amount of inframarginal consumer surplus created by increasing product variety. [Hausman \(2003\)](#) suggests using a linear demand curve to estimate inframarginal consumer surplus instead. Second, CES assumes that the elasticity of utility to increasing product variety is constant, whereas it may in fact be the case that the product space gets crowded out as new varieties get introduced. Following [Feenstra and Weinstein \(2017\)](#), I compute the gains from increasing product variety across income quintiles using the translog demand system, which allows for such crowding-out effects. These alternative specifications do not alter the result that higher-income consumers benefit relatively more from changes in product variety, although the inflation difference becomes lower than with CES, decreasing to around 0.75 percentage points a year.

III.C. Evidence from the CEX

Are the findings from the previous subsections specific to consumer packaged goods? To assess whether inflation inequality is likely to hold more generally, I compute inflation across income quintiles using the CEX-CPI sample, which covers the full consumption basket.

Figure III, Panel A reports the results for the benchmark Törnqvist index. The patterns are similar to those documented within consumer packaged goods: inflation declines with income. In the CEX-CPI sample between 2004 and 2015, average annual Törnqvist inflation was 1.86% for the fifth income quintile, while for the first income quintile it was 2.20% .

Figure III, Panel B shows that inflation inequality in the CEX-CPI sample is similar across price indexes. The average annual inflation difference between the top and bottom income quintiles is 0.346 (std. err. 0.0476) percentage points for Törnqvist, compared with 0.368 (std. err. 0.0502) for Laspeyres, 0.349 (std. err. 0.0480) for Paasche, and 0.349 (std. err. 0.0435) for CES. These numbers are about 50% smaller than the amount of inflation inequality found in the HMS sample.

Finally, Figure III, Panel C repeats the exercise focusing on the subset of products within the CEX-CPI sample that belong to consumer packaged goods. Here as well, a robust pattern of lower inflation for the top income quintile relative to the bottom income quintile is found. The average annual inflation difference varies between 0.264 and 0.280 percentage points depending on the price index.

Thus, inflation inequality in the CEX-CPI sample is smaller than in the HMS sample, but it remains sizable.

III.D. Decompositions

Next I examine which product categories drive the patterns of inflation inequality. I find that inflation inequality largely arises within detailed product categories, which can explain why the degree of measured inflation inequality is smaller in the CEX-CPI sample than in the HMS sample.

Inflation differences across income groups reflect the combined effects of both price and quantity changes, as well as baseline differences in spending patterns across income groups. For instance, it could be that high-income households spend more on fresh produce and that inflation tends to be lower in this broad item category. Alternatively, it could be the case that high-income households experience different inflation rates compared with low-income households on the same barcodes, for instance because they shop at different stores or have different propensities to use coupons. Accordingly, the inflation difference between high-income and low-income households can be decomposed into a

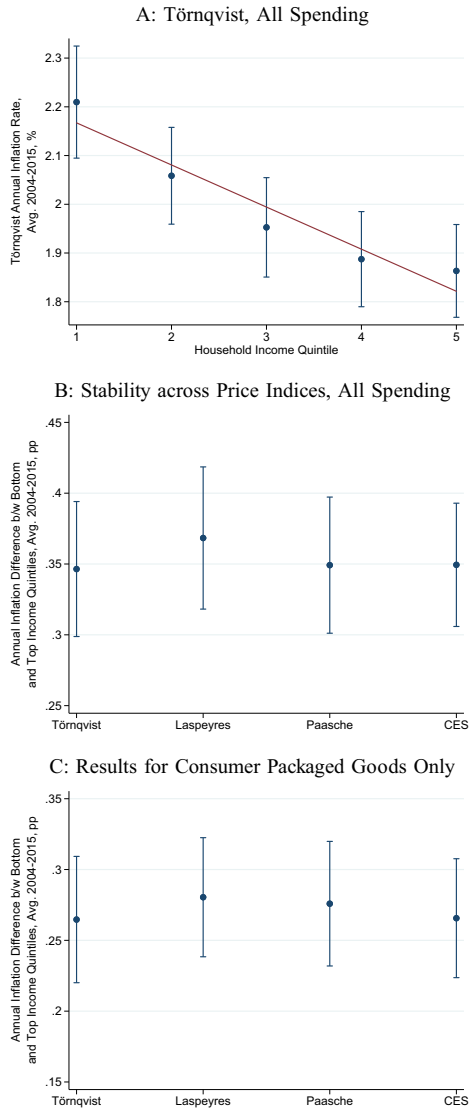


FIGURE III

Inflation across Income Groups in CEX-CPI

Panels A and B report the average annual inflation rate for various household groups in the CEX-CPI data from 2004 to 2015. In Panel C, the sample is restricted to the following product categories: food and alcohol at home, smoking products, stationery, personal care, and household cleaning products. The 95% confidence intervals are obtained by bootstrap.

“between” component and a “within” component. The “between” component corresponds to the inflation difference that would prevail if households differed only in terms of their expenditure shares across item categories and experienced the same inflation rate within each 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 , the inflation difference between high- and low-income households can be decomposed as follows (Diewert 1976):

$$(1) \quad \pi^R - \pi^P \equiv \sum_G s_G^R \pi_G^R - \sum_G s_G^P \pi_G^P =$$

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

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

Table II carries out this decomposition for the inflation difference between the top and bottom income quintiles, using both the HMS and RMS samples. The decomposition is reported separately for inflation for continuing products, using the Törnqvist index, and for the Feenstra ratio capturing changes in product variety.

Columns (1) and (2) report the decomposition for inflation for continuing products. Differences in inflation rates across barcodes induce a difference in average annual inflation rates of 0.541 percentage points, whereas the full inflation difference in Figure I was 0.661 percentage points. Of the inflation difference across barcodes, 88.6% arises between product modules by price deciles, 66.2% across product modules, 53.9% across product groups, and only 13% across departments. Thus, inflation inequality arises at a pretty fine level of disaggregation. About half of the difference is lost by going from the level of barcodes ($N = 2,240,278$) to the level of product groups ($N = 112$).

Columns (3) and (4) show that the patterns are very similar for the welfare effects from changes in product variety. Of the

TABLE II
 DECOMPOSITION OF DIFFERENCES IN INFLATION AND CHANGES IN PRODUCT VARIETY BETWEEN TOP AND BOTTOM INCOME QUINTILES IN
 NIELSEN DATA

Aggregation Level	HMS (2004–2015)				RMS (2006–2015)			
	Δ Inflation rates, continuing products		Δ Log Feenstra variety adj.		Δ Inflation rates, continuing products		Δ Log Feenstra variety adj.	
	pp (1)	% Explained (2)	pp (3)	% Explained (4)	pp (5)	% Explained (6)	pp (7)	% Explained (8)
Barcodes <i>N</i> = 2,240,278	0.541	100	1.487	100	0.527	100	1.277	100
Product Modules by price deciles <i>N</i> = 10,371	0.479	88.6	0.974	65.5	0.420	79.8	0.797	62.4
Product Modules <i>N</i> = 1,042	0.358	66.2	0.578	38.9	0.276	52.4	0.485	37.9
Product groups <i>N</i> = 112	0.291	53.9	0.493	33.2	0.237	44.9	0.47	36.8
Departments <i>N</i> = 10	0.071	13	-0.048	-3.3	0.077	14.7	0.038	3

Notes. This table decomposes the inflation difference between the top and bottom income quintiles in the HMS/RMS data, following equation (1).

TABLE III
 DECOMPOSITION OF DIFFERENCES IN INFLATION AND PRODUCT VARIETY BETWEEN TOP
 AND BOTTOM INCOME QUINTILES IN CEX-CPI DATA SET (2004–2015)

Aggregation level	Δ Inflation rates	
	pp	% Explained
Detailed categories $N = 256$	0.3464	100
Subcategories $N = 22$	0.0739	21.3
Main categories $N = 11$	0.0965	27.8

Notes. This table decomposes the inflation difference between the top and bottom income quintiles in the CEX-CPI data, following [equation \(1\)](#).

difference, 65.5% arises within product modules by price deciles, 38.9% across product modules, and 33.2% across product groups, while differences across department account for essentially nothing.

Columns (5), (6), (7), and (8) repeat the exercise using the RMS data, which is available from 2006 to 2015. The spending shares across income quintiles are taken from the HMS data, but prices are measured in the RMS data. The patterns are very similar to those found with the HMS data, indicating the robustness of this finding.

The fact that much of inflation inequality arises within detailed product categories may explain why measured inflation inequality in the CEX-CPI sample is smaller than with the HMS/RMS data. [Table III](#) investigates this hypothesis. I repeat the decomposition exercise following [equation \(1\)](#), but now using the product hierarchy available in the CEX-CPI data. The average annual inflation difference between the top and bottom income quintile is 0.3464 percentage points when I use spending patterns across the most detailed categories ($N = 256$). This difference falls by about 80% when I only consider spending differences across subcategories ($N = 22$). Likewise, the inflation difference across the main categories ($N = 11$) is only 0.0965 percentage points a year.

These decomposition results paint a consistent picture: inflation inequality tends to arise within detailed product categories. The CEX-CPI data may be missing a substantial fraction of

inflation inequality, although spending patterns across detailed categories already generate relatively high inflation inequality.

III.E. Robustness

The [Online Appendix](#) reports several robustness checks. [Online Appendix Table A2](#) shows that the findings are similar when excluding the Great Recession, when considering non-durables only (alcohol, dairy, deli, dry grocery, fresh produce, and packaged meat), or using an alternative definition of new and exiting products. [Online Appendix Figure A4](#) shows that inflation decreases approximately at a linear rate across deciles (rather than quintiles) of the income distribution, in both the HMS-RMS and CEX-CPI data sets. [Online Appendix Table A3](#) reports a decomposition showing that much of inflation inequality arises because of spending patterns within retail chains and within stores (rather than across). [Online Appendix Table A4](#) presents results using additional scanner data from 1998 to 2004; the results are similar to the main sample. Finally, [Online Appendix Figure A5](#) extends the CEX-CPI data to obtain much longer coverage, going back to 1953. Doing so requires using less detailed product categories (48 instead of 256). This extended sample shows that inflation inequality persists over the long run; however, consistent with the results presented above on the role of aggregation, measured inflation inequality is smaller in this sample with coarser categories.⁹

III.F. Discussion

The results presented here shed new light on a number of facts that were partly documented in previous work.

First and foremost, [Argente and Lee \(2016\)](#) construct income-group-specific price indexes from 2004 to 2010 using Nielsen scanner data and report that annual inflation for the highest income quartile was on average 0.59 percentage points lower than for

9. Unreported robustness checks indicate that the results are not driven by selection effects induced by the exit of certain products (as in [Erickson and Pakes 2011](#)) and remain stable when using quarterly data, nonchained price indexes, alternative measures of household income (adjusted for household size), and the alternative definition of price deciles (decile based on the weighted average unit price over a two-year period instead of the base period only).

the lowest income quartile.¹⁰ They interpret this inflation difference as being driven by the Great Recession.¹¹ [Online Appendix Figure A6](#) shows that the inflation difference between high- and low-income households in fact existed before, during, and after the Great Recession. The key difference between [Argente and Lee \(2016\)](#) and this article is the mechanism. They summarize their product quality substitution mechanism as follows: “The recent literature shows that households’ shopping behavior changed during the Great Recession. Households changed the quality of the items they bought [...] Because households have different margins within which to adjust their shopping behavior, they face heterogeneous inflation rates” (1). Although this channel is theoretically plausible, it can be checked that in practice it explains little of inflation inequality. [Figure I](#), Panel B shows that inflation inequality remains quantitatively very similar when using price indexes that do not allow for any substitution patterns, such as the Laspeyres and Paasche indexes. By definition, such indexes rule out the shopping behavior adjustment margin hypothesized by [Argente and Lee \(2016\)](#). The Laspeyres and Paasche price indexes do not allow consumers to substitute across barcodes, let alone across the quality ladder.

In contemporaneous work also using Nielsen scanner data, [Kaplan and Schulhofer-Wohl \(2017\)](#) compute inflation at the household level and find lower inflation for higher-income groups

10. Besides the time period, the main differences between the data and techniques used by [Argente and Lee \(2016\)](#) and my work are twofold. First, they focus on a nested CES price index accounting for changes in product variety with one nest (product modules). In contrast, I use three nests for nested CES (departments, groups, and modules) to account for changes in product variety and I show that the inflation patterns for continuing products are very stable across price indexes ([I](#)). Second, [Argente and Lee \(2016\)](#) obtain estimates of CES elasticities of substitution that are very large, with a median of 16.5, a 25th percentile of 9.1, and a 75th percentile of 44.8. My elasticity estimates are significantly smaller ([Online Appendix Table A1](#)) and fall in between other estimates in the literature (e.g., [Broda and Weinstein 2010](#); [Hottman, Redding, and Weinstein 2016](#)). Because of these smaller estimated elasticities, I find a larger impact of changes in product variety on the nested CES inflation rate than do [Argente and Lee \(2016\)](#). While they find that changes in product variety increased inflation inequality by about 15%, with my benchmark elasticities the increase is about 40%.

11. They write: “we find substantial differences across income groups that arise during the Great Recession” (abstract); “[Figure 1] shows that the indexes for all income groups track each other closely but drastically vary during the Great Recession” (16). According to the NBER, the Great Recession lasted from December 2007 to June 2009.

conditional on various sociodemographic controls. The main difference is that their methodology does not allow them to take into account the role of product variety.

Finally, the findings I obtain with the CEX-CPI data set are consistent with prior work by [McGranahan and Paulson \(2006\)](#). They also find lower inflation for higher income groups. However, since they used more aggregated product categories, the degree of inflation inequality they found was significantly smaller.

In sum, the contribution of the first part of this article is to show that inflation inequality is likely to be more common and more quantitatively significant than previously thought—more common because it exists outside of recessions, and more quantitatively significant because aggregation bias led to smaller estimates in prior work. Going beyond past contributions, this article explores a new channel to explain this fact and evaluates it quantitatively.

IV. THE EQUILIBRIUM RESPONSE OF SUPPLY TO CHANGES IN DEMAND

This section presents estimates of the response of supply to changes in demand. After presenting a series of stylized facts, I use a shift-share research design to obtain causal estimates. I then present complementary evidence to shed light on the mechanism at play, investigating the role of markups in particular.¹²

IV.A. *Stylized Facts*

I first present stylized facts on product variety, inflation, market size growth, and consumer income.

The analysis relies primarily on the RMS data to avoid potential measurement issues with the HMS data. The sampling methodology of the HMS data is such that the HMS panelists are representative of the U.S. population. As more and more households enter high-income brackets over time, there are more and more high-income households in the HMS data over time. This could result in a mechanical increase in measured market size

12. In [Section III](#), the Törnqvist inflation was a natural benchmark because of its flexibility. Having established that Törnqvist and CES produce similar measures for inflation for continuing products in practice, [Section IV](#) proceeds with the CES price index as its benchmark. Doing so allows for a direct comparison for the patterns with and without changes and product variety, and the CES demand system can be used to infer markups in a tractable way.

in the HMS data for product categories selling to these high-income households. Using the RMS data alleviates this concern because spending is measured in stores and doesn't depend on the sampling of panelists. The RMS data are available from 2006 to 2015.

Figure IV presents the stylized facts. I run simple OLS regressions of the form $Y_\ell = X_\ell + \epsilon_\ell$, where ℓ indexes product modules by price deciles. Segmenting the product space by product modules by price decile is a natural choice given that Section III has shown the importance of segmented consumption patterns even within product modules. Standard errors are clustered at the level of product modules. In Panels A and B, the independent variable is average annual "real" spending growth in ℓ between 2006 and 2015. From t to $t + 1$ nominal spending growth is $S_{\ell,t,t+1} = \frac{\sum_{k \in \ell} D_{k,t+1} Q_{k,t+1}}{\sum_{k \in \ell} D_{k,t} Q_{k,t}}$ and real spending growth is defined as $Q_{\ell,t,t+1} \equiv \frac{S_{\ell,t,t+1}}{\Pi_{t,t+1}^{CES,\ell}}$, where $\Pi_{t,t+1}^{CES,\ell}$ is the CES price index from t to $t + 1$ in ℓ , accounting for changes in product variety.¹³

In Panel A, the average Feenstra ratio between 2006 and 2015 is regressed on average annual real spending growth in ℓ during the same period. This panel shows a clear negative linear relationship between the average Feenstra ratio, which proxies for increasing product variety, and real spending growth. Each dot on the graphs represents 1% of the data when weighted by spending. This relationship is consistent with the notion that increasing market size induces entry and hence increasing product variety.

In Panel B, the average annual Törnqvist inflation for continuing products between 2006 and 2015 is regressed on average annual real spending growth in ℓ . There is a clear negative linear relationship: inflation for continuing products is lower in parts of the product space that grow faster.

Panel C regresses annualized average real spending growth on average (spending-weighted) consumer income. There is a clear upward relationship: product modules by price deciles that sell to richer consumers have grown faster during the sample period. Taken together, Panels A–C are consistent with the

13. Formally, $\Pi_{t,t+1}^{CES,\ell} = \pi_{t,t+1}^{CES,\ell} \cdot \left(\frac{1-s_{N,t+1}^\ell}{1-s_{E,t}^\ell} \right)^{\frac{1}{\sigma_\ell-1}}$, where $\pi_{t,t+1}^{CES,\ell}$ is the CES inflation rate for continuing products in ℓ . As a benchmark, I use the same elasticities of substitution $\sigma_\ell = \sigma_m$ for all product module by price decile cells ℓ within product module m .

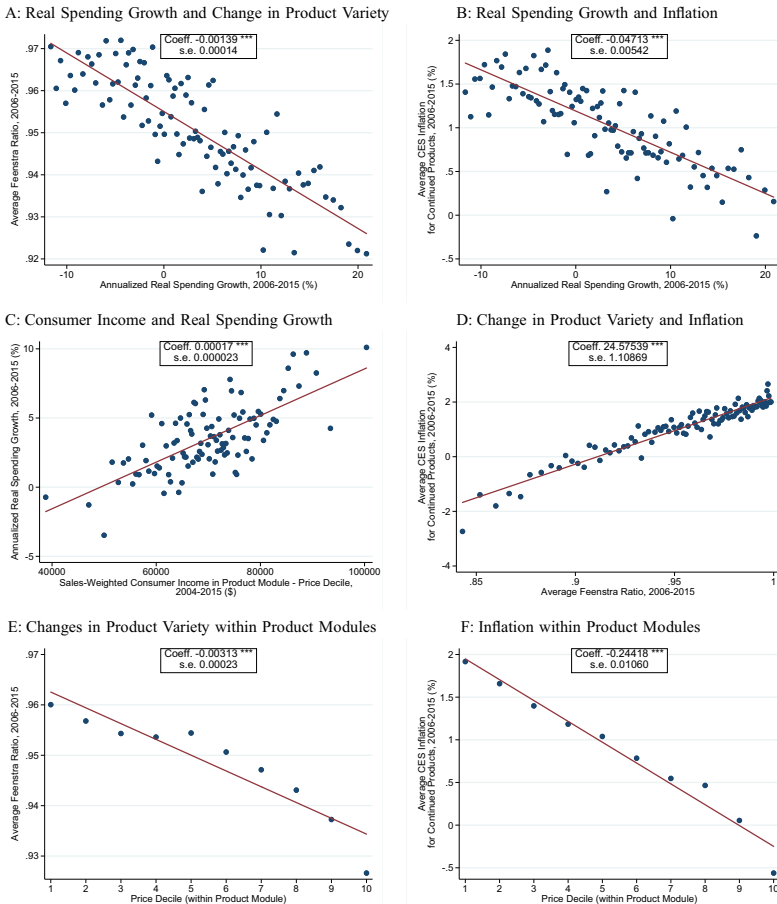


FIGURE IV
Stylized Facts

In all panels, the level of observation is a product module by price decile. In Panels A–D each dot represents 1% of observations. Panels E and F report patterns across price deciles, controlling for product module fixed effects. All regressions use spending weights and standard errors are clustered at the level of product modules.

following hypothesis: (i) demand from high-income consumers increased during the sample period, which led to increasing market size for premium products relative to entry-level products; (ii) in response, suppliers directed their product innovations towards premium market segments; (iii) in turn, this led to increased

competitive pressure and lower inflation for continuing products in these market segments.

Panel D provides another illustration of the hypothesized mechanism by showing the strong positive relationship between changes in product variety and inflation for continuing products across the product space. Product modules by prices deciles where product variety increases faster (with a lower Feenstra ratio) also have substantially lower inflation for continuing products (as measured with the Törnqvist index). The relationship appears to be linear.

Finally, Panels E and F report the patterns of changes in product variety and inflation for continuing products across price deciles within product modules. Describing these patterns within modules is instructive given that the decompositions in Section III showed that patterns across the quality ladder are a key determinant of inflation inequality. These panels show that product variety increases faster in higher price deciles, while inflation for continuing products is lower in higher price deciles.

A potential concern with Panel F is that the pattern of decreasing inflation for continuing products across price deciles may partly be mechanical due to mean reversion. Namely, a product that has a high price in a given period will be assigned to a high price decile, and it may have a lower price in the next period due to mean reversion. To address this concern, [Online Appendix Figure A7](#) repeats the exercise by building a measure of “leave-one-out brand unit price” for each barcode. For each barcode, I compute the average unit price for all other barcodes belonging to the same brand; I then assign barcodes to price deciles based on that brand-level measure. The results are similar.¹⁴

IV.B. A Shift-Share Research Design

I develop a research design to assess the causal effects of changes in demand on the price index, via changes in product variety and changes in inflation for continuing products.

1. Intuition. The equilibrium relationships between product innovations, price changes, and quantities across the product space, documented in the previous subsection, do not identify the

14. [Online Appendix Figure A8](#) shows that all stylized facts are very similar when using the HMS data (from 2004 to 2015) instead of the RMS data. [Online Appendix Figure A9](#) shows the patterns with nominal expenditure growth.

causal effect of demand because of (i) reverse causality: better products will have larger markets, that is, causality might run from supply to demand; and (ii) omitted variable bias: there might be unobserved heterogeneity in the difficulty of innovating across the product space, which could happen to coincide with spending patterns.

To address this issue, I use a shift-share research design. This design relies on two components: predetermined spending shares across the product space for a large number of sociodemographic groups, and heterogeneity in the population growth rates for these various groups during the sample period. For instance, age groups, education groups, racial groups, and regional populations all have different budget shares across the product space—variation in the size of these groups over time generates changes in demand. Spending profiles across the product space are measured in the initial period and kept constant, such that the variation in the shift-share instrument comes entirely from changes in the size of the sociodemographic groups over time. The strategy of using time-invariant spending profiles and changes in the size of household groups to address reverse causality follows a well-established literature, including [Acemoglu and Linn \(2004\)](#) and [DellaVigna and Pollet \(2007\)](#).

In the remainder of this section, I describe formally how the shift-share instrument helps address identification concerns, I discuss the assumptions underlying its validity, and I discuss the specific features of the data I use to build the instrument.

2. IV Framework. The goal is to understand how the price index P_ℓ responds to changes in the quantity index Q_ℓ induced by changes in demand. Conceptually, we want to find a demand shifter to vary Q_ℓ and observe the impact on P_ℓ across the cells of the product space indexed by ℓ . So we wish to estimate β in the following specification: $\Delta \log(P_\ell) = \beta \Delta \log(Q_\ell) + \gamma X_\ell + \varepsilon_\ell$, where X_ℓ is a vector of characteristics of product category ℓ including a constant, ε_ℓ is the (unobserved) potential outcome that would prevail in ℓ absent changes in demand, and $\mathbb{E}[\varepsilon_\ell] = \mathbb{E}[X_\ell \cdot \varepsilon_\ell] = 0$. Consistent estimation by OLS would require $\mathbb{E}[\Delta \log(Q_\ell) \cdot \varepsilon_\ell] = 0$, which is not a plausible assumption because quantities are endogenous to prices.

The shift-share design uses variation in Q_ℓ that comes from the variation in the size of household groups only, not from changes in prices (or preferences). Consider H household groups

indexed by h of size L^h with average income Y^h , facing similar price indexes P_ℓ across the product space. Total quantity demanded (Q_ℓ) is the sum of quantities demanded by each group (Q_ℓ^h): $Q_\ell = \sum_{h=1}^H Q_\ell^h = \sum_{h=1}^H (L^h Y^h \omega_\ell^h)$, where $\omega_\ell^h \equiv \frac{Q_\ell^h}{L^h Y^h}$ depends on prices and households' preferences.¹⁵ By log differentiation,

(2)

$$\Delta \log(Q_\ell) = \sum_{h=1}^H s_{h\ell} \Delta \log(L^h) + \sum_{h=1}^H s_{h\ell} \Delta \log(Y^h) + \sum_{h=1}^H s_{h\ell} \Delta \log(\omega_\ell^h),$$

with $s_{h\ell}$ the share of total sales to group h in cell ℓ . The first term shows that one can obtain variation in demand from changes in the composition of households over time. The third term depends on changes in prices, which shows the endogeneity problem.¹⁶

Motivated by equation (2), I build the following shift-share instrument to obtain variation in demand from the change in household size: $Z_\ell = \sum_{h=1}^H s_{h\ell 0} \cdot g^h$, where $g^h \equiv \Delta \log(L^h)$ and $s_{h\ell 0}$ denotes the share of sales in ℓ to households of type h in a base period $t = 0$. It is important to use predetermined spending shares at $t = 0$ because it follows from equation (2) that shares at $t > 0$ are endogenous to price changes occurring at any $t > 0$ during the sample. The instrument Z_ℓ uses only the part of the variation in demand in equation (2) that comes from changes in the size of various household groups. This instrument addresses the concern that changes in demand might be driven by reverse causality (price changes).

15. With $U_h = \left(\sum_{\ell=1}^L (\alpha_\ell^h)^{\frac{1}{\sigma}} (Q_\ell^h)^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}}$, that is, household groups have CES preferences with similar elasticity of substitution σ but different taste parameters α_ℓ^h , standard CES algebra yields Marshallian demand for ℓ by an agent of type h : $Q_\ell^h = \alpha_\ell^h \left(\frac{P_\ell}{P^h} \right)^{-\sigma} \frac{Y^h}{P^h}$ with $P^h = \left(\sum_{j=1}^L \alpha_j^h (P_j)^{1-\sigma} \right)^{\frac{1}{1-\sigma}}$. So total Marshallian demand for ℓ from h is $Q_\ell^h L^h = \omega_\ell^h Y^h L^h$ with $\omega_\ell^h \equiv \frac{\alpha_\ell^h}{(P_\ell)^\sigma} \frac{1}{\sum_{j=1}^L \alpha_j^h (P_j)^{1-\sigma}}$.

16. The second term of equation (2) shows that variation in household income $\Delta \log(Y^h)$ can also provide variation in demand. But in practice it is difficult to accurately measure income changes within detailed sociodemographic groups. Moreover, general equilibrium effects that affect income could at the same time affect prices, therefore it is preferable to use only changes in the size of the various household groups.

I use this instrument in a standard IV framework:

$$(3) \quad \begin{cases} \Delta \log(P_\ell) = \alpha Z_\ell + \gamma X_\ell + \varepsilon_\ell \\ \Delta \log(Q_\ell) = \lambda Z_\ell + \varphi X_\ell + \eta_\ell, \end{cases}$$

where $\frac{\alpha}{\lambda} \rightarrow \beta$ under suitable identification conditions. Before discussing the identification assumptions underlying this approach, I describe the practical implementation of the instrument.

First, I use product modules by price deciles as the cells ℓ of the product space. Second, I define household groups based on five dimensions: age, education, race, presence of children, and state of residence. Specifically, using the information provided by Nielsen for the household head, I build 12 age groups (from 25 to 75 in five-year age bins, below 25, and above 75), 5 education groups (below high school, high school graduates, some college, college graduates, and post-college graduate), 3 racial groups (White, African-American, Other), and 2 groups depending on whether the household has children. Finally, I keep track of households' place of residence across 49 states. I then construct household groups corresponding to all the observed stratas resulting from the interactions of these five household characteristics. I obtain 7,326 household groups, indexed by h in what follows.¹⁷ I measure baseline spending share ($s_{h\ell 0}$) for each of these households from 2004 to 2006 in the HMS data.

Second, the change in the size of each household group (g^h) is measured as the annualized log change in the population in group h between two four-year periods, at the beginning and end of the sample (2000–2004 versus 2012–2016). I obtain this information from IPUMS USA (see [Ruggles et al. 2017](#) and www.ipums.org), which provides consolidated data from the American Community Surveys (ACSs) from 2000 to the present. The ACS is a survey administered by the U.S. Census Bureau to track shifting demographics, making it an ideal data source for my purposes. All household variables described above in the Nielsen data are also available in the ACS. For each household group, I compute the annualized log change in population size over the sample period. For instance, on average older groups grow faster, as the Baby Boomers enter retirement.

17. The full interaction over all dimensions yields a potential number of $12 \times 5 \times 3 \times 2 \times 49 = 8,820$ household groups, but in practice only 7,326 are observed in the Nielsen data.

3. *Summary Statistics on the Baseline Shift-Share Instrument.* Table IV, Panel A, column (1) describes the variation in g^h , the annual log change in the size of household group h . The mean is 0.021, with substantial heterogeneity across groups: the standard deviation is 0.041 and the interquartile range is 0.052. This is the underlying source of variation which, combined with the initial spending shares $s_{h\ell 0}$, generates variation in demand across the product space.

Summary statistics on the shift-share instrument Z_ℓ are reported in Table IV, Panel B. Column (1) shows that the average value of the instrument is 0.0129, with substantial variation relative to the mean. The standard deviation is 0.0051 and the interquartile range is 0.0056. I only keep product modules for which enough data are available to construct price deciles in each year in the sample, leaving me with a total of 8,600 product module by price decile cells (accounting for over 90% of total spending).

4. *Identification Conditions.* IV consistency requires two conditions to be satisfied: instrument relevance, that is, Z_ℓ and $\Delta \log(Q_\ell)$ are correlated; and instrument validity, that is, Z_ℓ is uncorrelated with ε_ℓ . Relevance can be directly checked in the data (the first stage). To understand the exclusion restriction underlying instrument validity, I follow [Borusyak, Hull, and Jaravel \(2018\)](#). Applied to my setting, their results show that the exclusion restriction can be expressed as follows:

(4)

$$\text{Cov}[Z_\ell, \varepsilon_\ell] = \mathbb{E} \left[\left(\sum_{h=1}^H s_{h\ell 0} g^h \right) \cdot \varepsilon_\ell \right] = \sum_{h=1}^H s_h g^h \mathbb{E} \left[\frac{s_{h\ell 0}}{s_h} \cdot \varepsilon_\ell \right] \rightarrow 0,$$

where the covariance and expectation are taken over the product space indexed by ℓ . The key identification condition shown in [equation \(4\)](#) is a weighted covariance (in household space indexed by h , with spending weights s_h) between the shocks g^h and the unobservable term $\mathbb{E} \left[\frac{s_{h\ell 0}}{s_h} \cdot \varepsilon_\ell \right]$. The term $\mathbb{E} \left[\frac{s_{h\ell 0}}{s_h} \cdot \varepsilon_\ell \right]$ is a weighted average of product space unobservable potential outcomes ε_ℓ . If the quasi-experiment is valid, shocks to household populations are “as good as random” and the exclusion restriction is satisfied.¹⁸

18. In addition to the identification condition in [equation \(4\)](#), the other key assumption underlying my shift-share research design is that manufacturers are

TABLE IV
SUMMARY STATISTICS ON SHIFT-SHARE INSTRUMENT FOR CHANGES IN DEMAND

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Annual log change in group population								
Panel A: Changes in population of household Groups (2004–2016, yearly avg.)								
Mean	0.021	0.021	0.021	0.021	0.021	0.021	0.021	0.021
Standard deviation	0.041	0.038	0.036	0.037	0.041	0.039	0.041	0.026
Interquartile range	0.052	0.045	0.040	0.049	0.052	0.050	0.052	0.028
Residual change after controlling for								
age (linear)		✓						
age f.e.			✓					✓
education f.e.				✓				✓
state f.e.					✓			✓
race f.e.						✓		✓
children f.e.							✓	✓
Samples sizes								
<i>N</i> total	7,326							
<i>N</i> age groups	12							
<i>N</i> education group	5							
<i>N</i> states	49							
<i>N</i> race groups	3							
<i>N</i> presence of children	2							

TABLE IV
CONTINUED

	Instrument for annual log change in demand					
	(1)	(2)	(3)	(4)	(5)	(6)
Panel B: Implied changes in demand across product modules by price deciles						
Mean	0.0129	0.0122	0.0122	0.0122	0.0122	0.0185
Standard deviation	0.0051	0.0040	0.0039	0.0039	0.0036	0.0018
Interquartile range	0.0056	0.0041	0.0040	0.0039	0.0035	0.0013
Residual change after controlling for						
age (linear)		✓	✓	✓	✓	✓
age, educ., state, race, and children f.e.						✓
price decile f.e.			✓			✓
department f.e.				✓		
product group f.e.					✓	✓
<i>N</i> product modules — price deciles	8,600					

Notes: Panel A, the level of observation is a household group. The average annual log change in household population sizes is computed between 2000–2004 and 2012–2016 in the American Community Survey. Panel B reports statistics across the product space (product module by price decile cells) on the shift-share instrument between the same periods.

However, in practice certain household shocks observed in the historical data may happen to violate the exclusion restriction. For instance, as previously mentioned, during the sample period older household groups tend to grow faster (i.e., they have a high g^h). It may be intrinsically more difficult to innovate and reduce prices in product categories that sell more to older household groups, for instance because these households are likely to have defined their tastes earlier in life and to be less likely to adopt new products (i.e., their $\mathbb{E} \left[\frac{s_{h0}}{s_h} \cdot \varepsilon_\ell \right]$ is large). If this is so, then age groups with a high g^h also have a high $\mathbb{E} \left[\frac{s_{h0}}{s_h} \cdot \varepsilon_\ell \right]$, invalidating the exclusion restriction across age groups. Because of such potential concerns, it is important to assess whether the results are stable when using more “idiosyncratic” household population shocks (e.g., using only variation within age groups, not across). I now discuss how this can be achieved through the use of fixed effects.

5. *Alternative Specifications: Fixed Effects and Residualized Shift-Share Instrument.* To assess the robustness of the estimates, it is useful to consider the following statistical decomposition of the shocks g^h :

$$(5) \quad g^h = \mu + g_{age} + g_{education} + g_{state} + g_{race} + g_{children} + v_h.$$

This expression says that the observed shocks g^h ($N = 7,326$) can be expressed as the combination of the average shocks along the five dimensions that segment the household space—age, education, state, race, and presence of children—as well as a residual component v_h . One can compute a residualized household population shock \tilde{g}^h after controlling for age, education, state, race, and presence of children, either simultaneously or separately. Then, one can build the residualized shift-share instrument $\tilde{Z}_\ell = \sum_{h=1}^H s_{h\ell 0} \cdot \tilde{g}^h$. For instance, controlling for age fixed effects means that the instrument only relies on variation in household shocks that occur within each of the 12 age groups, which addresses the aforementioned potential concern about the validity of the exclusion restriction across age groups. Thus, one can build

able to anticipate trends in market demand coming from changes in the population sizes for the various sociodemographic groups described above. Under this assumption, the IV estimates reflect the supply response to properly anticipated changes in demand.

the residualized shift-share instrument in two steps: first regress g^h on household group fixed effects as in equation (5) to obtain the residualized household population shocks \tilde{g}^h ; then build the residualized shift-share instrument \tilde{Z}_ℓ .¹⁹

Table IV, Panel A, columns (2)–(8) present summary statistics on the residualized household population shocks, introducing different controls in turn. The mean is mechanically the same across all columns. Reassuringly, the amount of variation in household shocks remains very similar across specifications with different types of controls, regardless of whether the standard deviation or the interquartile range are considered. For instance, the standard deviation of household population shocks falls from 0.041 with the raw shocks to 0.038 with linear age controls and to 0.036 with age fixed effects. Across specifications (2) to (7), the standard deviation varies between 0.041 and 0.036 and the interquartile range from 0.04 to 0.052. The standard deviation remains sizable in column (8) when including all fixed effects at once. This result indicates that the household shocks are not driven by a single dimension of the data and are rather idiosyncratic, lending support to the notion that they can be used in a quasi-experimental framework. The choice of the specification (i.e., how to residualize) can be thought of as a bias-variance trade-off. As more fixed effects are included, the quasi-experimental interpretation potentially becomes more plausible (reducing bias), but the instrument loses some of its power (increasing variance).

In addition to investigating the robustness of the patterns by residualizing the household shocks (and thus the shift-share instrument), it is also instructive to control for characteristics of product modules by price deciles l . If the quasi-experiment is valid, that is to say if the household shocks driving the shift-share instrument are truly idiosyncratic, then the results should remain stable as one varies the set of product space controls X_ℓ in equation (3).

Table IV, Panel B, columns (2)–(6) report summary statistics on the residualized shift-share instrument. The overall message is very similar to that of Panel A: substantial variation in the instrument remains as fixed effects are introduced. When introducing linear age controls in column (2), the variance of the

19. Borusyak, Hull, and Jaravel (2018) show that residualizing the instrument in this way is equivalent to running a one-step IV specification with household characteristics projected onto the product space using initial spending shares.

instrument falls to 0.004, from 0.0051 for the raw instrument in column (1). Adding price decile fixed effects, department fixed effects, and product groups reduces the variance of the instrument only slightly further, down to 0.0036 in column (5). Column (6) reports that the standard deviation remains sizable when including all household fixed effects and the most detailed product space fixed effects (product groups and price decile); this will be my preferred specification.

6. Inference. As discussed in [Adao, Kolesar, and Morales \(2018\)](#) and [Borusyak, Hull, and Jaravel \(2018\)](#), it is important to consider the presence of correlated errors in shift-share research designs. Intuitively, there might be a correlation between residuals ε_ℓ in parts of the product space ℓ that sell to similar consumers, because these cells may be exposed to similar combinations of unobserved demand or supply shifters. Accordingly, in addition to clustering by product modules in the main specifications, for robustness I also report standard errors accounting for the shift-share correlation structure.

IV.C. Baseline Estimates

I use the IV framework summarized in [equation \(3\)](#) to provide the main estimates of the impact of changes in demand on price indexes. Specifically, I consider the impact on the CES index for continuing products and on the CES price index accounting for changes in product variety. As in [Section IV.A](#), real spending growth $\Delta \log(Q_\ell)$ is defined as the growth of nominal spending deflated by the CES price index. The price indexes and real spending growth are all computed using the RMS data from 2006 to 2015.

1. Main Results. [Figure V](#) provides a graphical depiction of the main results using the baseline shift-share instrument. The baseline specification includes linear controls for age in household space (using [equation \(5\)](#) to address potential concerns about inflation systematically varying with age). Panel A reports the first stage: there is a clear positive relationship between the shift-share instrument on the x -axis and the actual change in market size on the y -axis, which shows the power of the instrument. Note that the actual change in market size is measured at the level of stores in the RMS data. This is a helpful feature of RMS relative to HMS. With HMS, one would be worried about finding a




FIGURE V

The Effect of Changes in Demand on Inflation and Product Variety

This figure presents the first-stage and reduced-form relationships for the shift-share research design, using the RMS data from 2006 to 2015 to build the dependent variables. Each dot represents 1% of the data. All regressions use spending weights and standard errors are clustered at the level of product modules.

mechanical first stage due to the sampling scheme, because to ensure that panelists are representative, Nielsen increasingly samples more households in groups whose populations increase over time.

Figure V, Panels B and C report the reduced-form relationships. The patterns are very clear: both Törnqvist inflation for continuing products and CES inflation (with the adjustment for changes in product variety) fall for larger values of the shift-share instrument. The fall is steeper in Panel C, indicating that increasing product variety plays a role above and beyond declining inflation for continuing products. In the panels of Figure V, each dot represents 1% of the data and the relationships described above appear to be approximately linear. These findings lend support to the hypothesis that supply endogenously responds to changes in market size.

Table V reports regression results for the first-stage and reduced-form relationships, as well as IV estimates. Column (1) confirms the strength of the first stage, column (2) shows the negative response of the CES index for continuing products, and column (7) reports a stronger negative response for CES inflation with changes in product variety. The other columns present the IV estimates.

The IV estimates for the CES index for continuing products are reported in columns (3) to (6) of Table V. The first stage F -statistic and Cragg-Donald F -statistic confirm that the instrument is not weak. Column (3) presents the IV estimates controlling for linear age effects only. Column (4) reports the results with the full set of fixed effects in household space (age, education, state, race, and children), and column (5) adds the full set of fixed effects in the product space (price deciles and product groups). The estimates are similar across columns, statistically significant at the 1% level, and statistically indistinguishable from each other. Because it includes the full set of fixed effects, column (5) serves as my preferred specification: a 1 percentage point increase in demand leads to a fall in the CES inflation rate for continuing products of 0.418 percentage points. Adjusting the standard errors for shift-share clustering in column (6) increases standard errors only slightly.

The IV estimates for CES inflation with changes in product variety paint a similar picture and are reported in columns (8) to (11). The IV coefficients are stable across specifications and larger than for continuing products. Column (10) gives our

TABLE V
THE EFFECT OF CHANGES IN DEMAND ON INFLATION AND PRODUCT VARIETY

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
	$\Delta \text{Log market size}$ (annual avg., 2006–2015)		$\text{CES inflation for continuing products}$ (annual avg., 2006–2015)				$\text{CES inflation with changes in product variety}$ (annual avg., 2006–2015)				
Instrument for annual log change in demand	2.760*	-1.042*					-1.721*				
	(0.556)	(0.201)					(0.308)				
$\Delta \text{Log market size}$ (annual avg., 2006–2015)			-0.378*	-0.458*	-0.418*	-0.418†		-0.624*	-0.673*	-0.618†	-0.618†
			(0.081)	(0.14)	(0.139)	(0.169)		(0.128)	(0.241)	(0.258)	(0.288)
Estimator	OLS	OLS	IV	IV	IV	IV	OLS	IV	IV	IV	IV
Instrument relevance											
First-stage F			24.638	12.964	13.981	13.981		24.638	12.964	13.981	13.981
Cragg-Donald F			43.216	24.249	15.114	15.114		43.216	24.249	15.114	15.114
Controls			✓	✓	✓	✓	✓	✓	✓	✓	✓
age (linear)	✓										
age/educ/state/race/children f.e.				✓	✓	✓			✓	✓	✓
price decile/product group f.e.					✓	✓				✓	✓
Standard errors											
Clustered by modules	✓		✓	✓	✓	✓		✓	✓	✓	✓
Shift-share clustering						✓					✓
$N = 8,600$											

Notes: This table reports the main results of the shift-share research design presented in Section IV.B. The outcome variables are measured in the RMS data from 2006 to 2015. * denotes significance at the 1% level and † at 5%.

preferred specification: a 1 percentage point increase in demand leads to a fall in the CES inflation rate of 0.618 percentage points, about 50% larger than the effect for continuing products alone.

2. *Robustness.* Table VI presents a series of robustness checks. Panel A examines the stability of the IV estimates as various fixed effects for the product space are added iteratively to the IV specification. Columns (1) to (3) focus on continuing products, and columns (4) to (6) report the results accounting for product variety. Iteratively adding price decile fixed effects, department fixed effects, and product group fixed effects in the other columns leads to a fall in the magnitude of the effect, down to coefficients of -0.23 for continuing products in column (4) and -0.38 with changes in product variety in column (8). Although the coefficients are smaller, they remain large and statistically significant at the 1% level. Moreover, they are statistically indistinguishable from the estimates reported in Table V.

Table VI, Panel B presents the IV estimates across fixed effects in household space. With age fixed effects or education fixed effects the IV estimates remain similar for continuing products (columns (1) and (2)) and when considering changes in product variety (columns (8) and (9)). Adding price decile fixed effects and department fixed effects leads to a relatively small fall in the magnitude of the estimates, as shown in columns (3) and (10). The remaining columns add household fixed effects on top of the price decile and department fixed effects. But the point estimates always remain sizable and statistically significant at the 1% level: the smallest effects are -0.27 for continuing products and -0.53 with changes in product variety. Moreover, there is no specification in which we can reject the preferred IV estimates reported in Table V.

The Online Appendix reports additional robustness checks. Online Appendix Figure A10 shows that the results remain similar when using other price indexes for continuing products and other elasticities of substitution in the CES framework with changes in product variety. Online Appendix Figure A11 shows that similar results are obtained with the HMS data instead of the RMS data. Online Appendix Table A5 shows that the results are similar for Törnqvist inflation for continued products. Online Appendix Table A6 reports the results when instrumenting for nominal expenditure growth.

TABLE VI
ROBUSTNESS OF IV ESTIMATES

	CES inflation for continuing products (annual avg., 2006–2015)		CES inflation with changes in product variety (annual avg., 2006–2015)			
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Stability across fixed effects in product space						
$\Delta \text{Log market size}$	-0.32*	-0.31*	-0.23*	-0.57*	-0.48*	-0.38*
(annual avg., 2006–2015)	(0.075)	(0.089)	(0.065)	(0.116)	(0.122)	(0.089)
Instrument relevance						
First-stage F	31.27	24.72	21.77	31.27	24.72	21.77
Cragg-Donald F	52.11	52.43	31.02	52.11	52.43	31.02
Controls						
age (linear)	✓	✓	✓	✓	✓	✓
price decile f.e.	✓			✓		
department f.e.		✓			✓	
product group f.e.			✓			✓
$N =$	6098					

TABLE VI
CONTINUED

	CES inflation for continuing products (annual avg., 2006–2015)				CES inflation with changes in product variety (annual avg., 2006–2015)									
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)
Panel B: Stability across fixed effects in household space														
Δ Log market size (annual avg., 2006–2015)		-0.36* (0.12)	-0.42* (0.073)	-0.32* (0.092)	-0.33* (0.051)	-0.29* (0.078)	-0.27* (0.066)	-0.28* (0.077)	-0.64* (0.167)	-0.70* (0.107)	-0.60* (0.146)	-0.60* (0.075)	-0.54* (0.123)	-0.53* (0.121)
Instrument relevance														
First-stage F	15.38	43.91	15.11	61.11	17.01	21.70	16.11	15.38	43.91	15.11	61.11	17.01	21.70	16.11
Cragg-Donald F	22.17	68.18	22.82	92.22	27.97	31.95	27.29	22.17	68.18	22.82	92.22	27.97	31.95	27.29
Controls														
age fe.	✓							✓						
education fe.		✓		✓					✓			✓		
state fe.					✓								✓	
race fe.						✓								✓
children fe.							✓							✓
price decile fe.			✓	✓	✓	✓	✓			✓	✓	✓	✓	✓
department fe.			✓	✓	✓	✓	✓			✓	✓	✓	✓	✓
$N = 8,600$														

Notes. The outcome variables are measured in the RMS data from 2006 to 2015. Standard errors are clustered by product modules. * denotes significance at the 1% level.

IV.D. Market Size Effects in Levels versus Changes

To understand which properties of the income distribution matter for induced innovations, it is particularly important to understand whether changes in demand or the level of demand induce the supply responses previously documented.²⁰ If changes in demand are the driver then the main beneficiaries from induced innovations should be in the ranges of the income distribution in which purchasing power is increasing, such as at the top of the distribution. In contrast, if the level of demand drives the supply response, then the middle class should always be the main beneficiary of endogenous innovations, because it accounts for most of the market size.

1. *Research Design.* The IV estimates from Tables V and VI cannot help determine whether demand in levels or in changes leads to lower price indexes. Instead, one needs a quasi-experiment in which both the initial level of demand and the growth of demand are as good as randomly assigned, making it possible to estimate the impact of both channels on product innovations and price changes. To approximate this ideal experiment, I rely on another shift-share design. Conceptually, one would like to estimate the following equation: $\Delta \log(P_\ell) = \beta_{Change} \Delta \log(Q_\ell) + \beta_{Level} \log(Q_{\ell 0}) + \gamma X_\ell + \varepsilon_\ell$, where $\log(Q_{\ell 0})$ is total market size in the baseline period. If it were possible to build a shift-share instrument for the level and the change in demand, one could run a horse race between the two channels. However, the shift-share design from Section IV.B makes this impossible because the initial spending shares are endogenous to prices. Instead, I use the

20. The results from the baseline shift-share research design indicate that higher demand leads to a fall in the price index; but this result could be interpreted in two ways. On one hand, it could be the case that a higher level of demand always leads to a higher flow of innovations and lower prices. For instance, in Romer (1990) and Aghion and Howitt (1992) a higher level of market size always leads to more innovations, because it is more profitable to pay the fixed cost of innovation to try to capture a large market. Intuitively, a higher level of demand creates higher returns to R&D and in turn increasing returns to scale. On the other hand, it could be the case that increasing demand is the driver of innovations and declining prices. For example, Jones (1995) developed a model in which it is the growth of market size, not its level, which drives innovation. Intuitively, if it becomes harder to innovate as the market becomes larger, then a higher level of demand creates higher returns to R&D but also higher costs of R&D—only additional growth of demand can induce more innovations.

fact that for consumer packaged goods, markets are quite differentiated across states. For instance, certain manufacturers have strong local brand capital, such that what matters for them is changes in the level of local (in particular, state) demand rather than changes in nationwide demand (Bronnenberg, Dubé, and Gentzkow 2012). One can rewrite the desired specification across cells of the product space which are now “product module by price decile by state” cells, indexed by ℓs : $\Delta \log(P_{\ell s}) = \beta^{Change} \Delta \log(Q_{\ell s}) + \beta^{Level} \log(Q_{\ell s0}) + \gamma X_{\ell s} + \varepsilon_{\ell s}$.

I then consider household groups which differ along four dimensions: age, education, race, and presence of children. As previously, there are 12 age groups, 5 education groups, 3 race groups, and an indicator for the presence of children. I fully interact these characteristics to define household groups ($N = 341$). In each state for each household group, I keep track of the initial population in 2000–2016, denoted $g_{Level}^{h,s}$, where h indexes the household group and s the state. I also keep track of the change in population size between 2000–2004 and 2012–2016, denoted $g_{Change}^{h,s}$. As before, I obtain the data from IPUMS USA. Using spending shares in the HMS data from 2004 to 2006, I build two shift-share instruments: $Z_{\ell s}^{Change} = \sum_{h=1}^H s_{h\ell 0(-s)} \cdot g_{Change}^{h,s}$ and $Z_{\ell s}^{Level} = \sum_{h=1}^H s_{h\ell 0(-s)} \cdot g_{Level}^{h,s}$, where $s_{h\ell 0(-s)}$ is the initial spending share (in 2004–2006) for household group h on product category ℓ , measured by leaving out all households in state s . By using “leave-one-out” spending shares (leaving out the state of reference s) interacted with differences in population size across states, one can obtain an instrument for the level of market size. This was not possible in the baseline research design because nationwide spending shares are endogenous to prices; using segmented markets across states addresses this limitation.

Intuitively, this research design exploits the fact that certain states have a larger baseline population with tastes skewed toward particular goods, which provides an instrument for initial market size. The variation comes entirely from differences in the sociodemographic composition of the population in different states, and how it changes over time. As previously, one can control for potential confounding factors by adding fixed effects. State fixed effects ensure that the identifying variation does not capture generic scale effects (for instance, states with a larger population may have more efficient distribution networks for reasons other than the endogenous response of supply to demand).

I use these shift-share instruments in standard IV frameworks. The main IV specification focuses on the level of market size only:

$$(6) \quad \begin{cases} \Delta \log(P_{\ell s}) = \alpha Z_{\ell s}^{Level} + \gamma X_{\ell s} + \varepsilon_{\ell s}, \\ \log(Q_{\ell s 0}) = \lambda Z_{\ell s}^{Level} + \gamma X_{\ell s} + \eta_{\ell s}. \end{cases}$$

For robustness, I repeat the IV exercise with two endogenous variables, the change in market size and the level of market size:

$$(7) \quad \begin{cases} \Delta \log(P_{\ell s}) = \alpha_1 Z_{\ell s}^{Change} + \alpha_2 Z_{\ell s}^{Level} + \gamma X_{\ell s} + \varepsilon_{\ell s}, \\ \Delta \log(Q_{\ell s}) = \lambda_1 Z_{\ell s}^{Change} + \lambda_2 Z_{\ell s}^{Level} + \gamma X_{\ell s} + \eta_{\ell s}, \\ \log(Q_{\ell s 0}) = \tilde{\lambda}_1 Z_{\ell s}^{Change} + \tilde{\lambda}_2 Z_{\ell s}^{Level} + \tilde{\gamma} X_{\ell s} + \tilde{\eta}_{\ell s}. \end{cases}$$

The identification assumption is analogous to the baseline shift-share design: conditional on the covariates $X_{\ell s}$, the initial sociodemographic composition of the state population and its changes over time must be as good as randomly assigned. As previously, the change in price indexes $\Delta \log(P_{\ell s})$ and real spending growth $\Delta \log(Q_{\ell s})$ are all computed using the RMS data, from 2006 to 2015. $\log(Q_{\ell s 0})$ is measured as the log of total spending in product module by price decile ℓ in state s in 2006.

2. *Results.* Figure VI provides a graphical depiction of the first-stage and reduced-form relationships in the IV design focusing on the level of market size (equation (6)). As can be seen in Panel A, the state-level shift-share instrument is a strong predictor of the level of market size. Panel B shows that there is no relationship between the instrument for the level of market size and the Törnqvist index for continuing products; likewise, Panel C shows no relationship between the instrument and CES inflation with changes in product variety. These results provide a clear indication that inflation and product variety do not respond to the level of market size.

Table VII, Panel A reports the results for specification (6). Column (1) confirms that the first stage is strong. The shift-share instrument is not correlated with either CES inflation for continuing products (column (2)) or CES inflation with adjustment for product variety (column (5)). The IV estimates in columns (3) and (4) report a precisely estimated 0 effect of large market size on CES inflation, regardless of whether standard errors are clustered by product modules or by state. Columns (6) and (7) show the

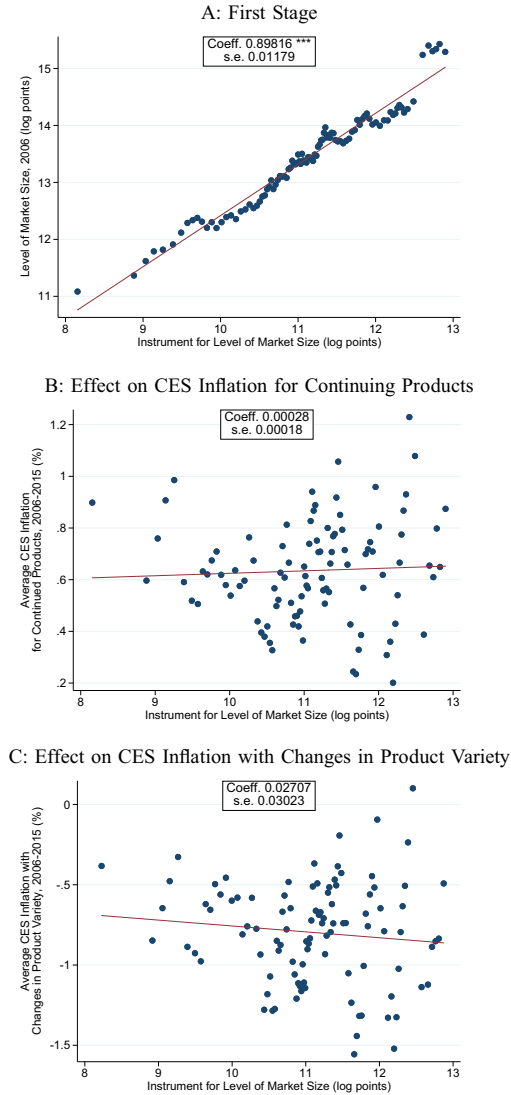


FIGURE VI

The Effect of the Level of Market Size on Inflation and Product Variety

This figure describes the relationship between the state-level shift-share instrument for the level of market size and the actual level of market size (Panel A) as well as price indexes (Panels B and C). The dependent variables are measured in the RMS data. Each dot represents 1% of the data. All regressions use spending weights and standard errors are clustered at the level of product modules.

TABLE VII
THE EFFECT OF THE LEVEL OF MARKET SIZE ON INFLATION AND PRODUCT VARIETY

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Log market size (2006)	CES inflation for continuing products (annual avg., 2006–2015)	CES inflation for continuing products (annual avg., 2006–2015)	CES inflation with changes in product variety (annual avg., 2006–2015)			
Panel A: Main specifications							
Instrument for log market size	0.898* (0.012)	0.00028 (0.00018)			0.0003 (0.0003)		
Log market size (2006)			0.00031 (0.00019)	0.00031 (0.00034)		0.0003 (0.0003)	0.0003 (0.0005)
Estimator	OLS	OLS	IV	IV	OLS	IV	IV
Instrument relevance							
First-stage <i>F</i>			5,799.80	161.87		5,799.80	161.87
Cragg-Donald <i>F</i>			55,426.66	55,426.66		55,426.66	55,426.66
Standard errors							
Clustered by modules	✓	✓	✓	✓	✓	✓	✓
Clustered by states							
<i>N</i>	201,247						

TABLE VII
CONTINUED

	CES inflation for continuing products (annual avg., 2006–2015)		CES inflation with changes in product variety (annual avg., 2006–2015)			
	(1)	(2)	(3)	(4)	(5)	(6)
Panel B: Robustness						
Log market size	0.01152 [†] (0.00448)		-0.01925 (0.01422)	0.0198* (0.00727)		-0.01665 (0.01499)
Δ Log market size (annual avg., 2006–2015)		-0.406* (0.0787)	-0.3364* (0.1190)		-0.633* (0.125)	-0.3575* (0.1300)
Estimator	IV	IV	IV	IV	IV	IV
Instrument relevance						
First-stage F	13.47	26.88	11.76	13.47	26.88	11.76
Cragg-Donald F	2,635.88	1,274.83	23.71	2,635.88	1,274.83	23.71
Controls						
location f.e.	✓	✓	✓	✓	✓	✓
age, educ., race, & children f.e.			✓			✓
price decile & product group f.e.			✓			✓
$N = 201,247$						

Notes: Panel A reports the results of the shift-share research design for the level of market size, using specification (6). Panel B reports robustness checks. Columns (4) and (8) of Panel B report the results from IV specifications with two endogenous variables, using specification (7). * denotes statistical significance at the 1% level and [†] at the 5% level.

same result for CES inflation with changes in product variety. In all IV specifications, the F -statistics confirm that the instrument is strong.

Table VII, Panel B reports various robustness specifications. In column (1), state fixed effects are included. In column (1), the effect of a larger market size on CES inflation for continuing products becomes positive and statistically significant at the 5% level, but the magnitude is small. CES inflation is 1.2 basis points smaller in a market that's 1% larger. Column (2) reports the results for the state-level shift-share research design focusing on changes in market size. In this column, the endogenous variable is the change in market size and the instrument is Z_{ts}^{Change} . The IV estimate is large, negative, and similar in magnitude to the baseline IV estimates reported in Table V. Finally, column (3) shows the result when including both the level and change in market size in the IV, as in equation (7), as well as the full set of fixed effects. The results indicate that a 1 percentage point increase in demand leads to a 33.6 basis point fall in inflation (std. err. 0.119), while a 1% larger market (in the base period) is characterized by a level of inflation that is merely 1.9 basis points lower (not statistically significant). Columns (4) to (6) report very similar results for CES inflation with the adjustment for product variety.

Across all specifications, a clear pattern emerges: while increasing market size leads to lower inflation and increasing product variety, an initially larger market is not on a different trajectory in terms of either inflation or changes in product variety. Put another way, supply expands in response to increasing demand, but not in response to larger initial demand. Of course, larger markets may still benefit from lower price levels and larger product variety (in a static sense), but the results show that the dynamics of supply are driven by changes in demand rather than by its level.

IV.E. *The Role of Markups*

The results so far indicate that inflation falls substantially when demand increases and that much of this effect comes from inflation for continuing products. How much of the estimated market size effects on price growth can be explained by a change in markups? I address these questions in two steps, first by presenting reduced-form evidence suggesting that falling markups are an important force in the data; and second by using the structural approach of Hottman, Redding, and Weinstein (2016) to quantify

the role of markups in explaining the price response to increasing market size.

1. *Reduced-Form Evidence on Store Gross Markups.* To provide direct evidence on the role of markups, I use data for a subset of products indexed by k for which retailer price p_{kst} and wholesale cost c_{kst}^w are observed from January 2004 to December 2006. The data set covers 250 grocery stores indexed by s , operated by a single grocery chain in 19 U.S. states (see [Gopinath et al. 2011](#)). [Online Appendix Table A7](#) shows that the main expenditure categories from the Nielsen data are properly covered in this data set. Because the retailer price and the wholesale cost are observed, the “gross” store markup over the wholesale cost can be recovered as $m_{kst} \equiv \frac{p_{kst}}{c_{kst}^w}$, which includes the “net” store markup and store costs such as store payroll and rent. To the extent that these store costs are similar for all products in the same store, one can account for them in regressions by including fixed effects. To assess whether falling markups play a role in inflation inequality, I examine the correlation between consumer income and changes in wholesale cost and store markups, in regressions that include store fixed effects. For each product k , using the HMS data I measure the average income of consumers who buy from the product module m_k that product k belongs to; using consumer income at the module level rather than at the product level is preferable to address measurement error. I then run the following specifications: $\Delta \log(p_{kst}) = \beta^p I_{m_k} + \lambda_{st} + \varepsilon_{st}$, $\Delta \log(c_{kst}^w) = \beta^{c^w} I_{m_k} + \lambda_{st} + \tilde{\varepsilon}_{st}$, and $\Delta \log(m_{kst}) = \beta^m I_{m_k} + \hat{\lambda}_{st} + \hat{\varepsilon}_{st}$, where λ_{st} denotes store-by-year fixed effects. Because $\Delta \log(p_{kst}) = \Delta \log(c_{kst}^w) + \Delta \log(m_{kst})$, we have $\beta^p = \beta^{c^w} + \beta^m$, which provides a convenient decomposition of overall inflation inequality (β^p) into wholesale cost effects (β^{c^w}) and store markup effects (β^m).

A potential concern with this approach is that the measure for the change in store markups may in fact include changes in store costs over time (i.e., stores that sell to richer consumers may be located in areas where rents decrease). To alleviate this concern, in the main specification I include store-by-year fixed effects, such that the variation is entirely within store-year, that is, the regression does not attribute changes in store-level cost to store markups. For robustness, I include store-by-department-by-year fixed effects, effectively allowing store costs to vary across departments.

The results are presented in [Table VIII](#), Panel A. Columns (1) to (3) report the results for the overall price change, columns (4) to (6) for the change in the gross retailer markup, and columns (7) to (9) for the change in wholesale cost. Various specifications are reported for robustness. Given the relationship between the outcome variables, the point estimates in columns (4) and (7) mechanically sum up to the point estimate in column (1)—likewise for columns (5) and (8) summing up to column (2), and columns (6) and (9) summing up to column (3). In the baseline specification including store-by-year fixed effects and standard errors clustered by stores (columns (1), (4), and (7)), both store markup effects and wholesale cost effects are significant at the 1% level. Markup effects account for 57.9% ($= \frac{0.545}{0.941}$) of the relationship between inflation and consumer income. Standard errors increase when clustering by product modules but statistical significance is retained (columns (2), (5), and (8)). To account for potential differential changes in store costs across the product space, columns (3), (6), and (9) include store-by-department-by-year fixed effects. The role of markup effects increases, now accounting for 68.4% ($= \frac{0.608}{0.889}$) of the relationship between inflation and consumer income. These correlations suggest that a sizable fraction of the relationship between inflation and consumer income can be accounted for by changes in store markups (which may be a lower bound on the overall importance of changes in markups for inflation inequality, because wholesalers, and in turn manufacturers, also have a markup).

2. *Structural Evidence Following [Hottman, Redding, and Weinstein \(2016\)](#).* A limitation of the analysis reported in [Table VIII](#), Panel A is that the patterns do not necessarily reflect causal relationships related to market size; moreover, the sample covers only three years for a subset of goods. To make additional progress on this front, I make assumptions following [Hottman, Redding, and Weinstein \(2016\)](#) to measure markups for all products and quantify the importance of the markup response in explaining price changes following changes in market size.

Assume that preferences are Cobb-Douglas across cells ℓ of the product space. Within each cell ℓ , assume two CES nests for firms (upper tier) and barcodes (lower tier). The elasticity of substitution between firms is denoted σ^F . Under these assumptions on the preference structure, the firm's profit maximization problem is separable across cells ℓ . Within each ℓ , the firm sets prices to

TABLE VIII
CONTINUED

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Δ Log markup (annual avg., 2006–2015)				Δ Log marginal cost (annual avg., 2006–2015)			
Panel B: Structural evidence on overall markups								
Δ Log market size (annual avg., 2006–2015)	-0.22 [†] (0.089)	-0.43 [†] (0.212)	-0.28 [†] (0.135)	-0.44* (0.166)	-0.45 [†] (0.212)	-0.35 [†] (0.155)	-0.15 (0.102)	0.07 (0.164)
Estimator	IV	IV	IV	IV	IV	IV	IV	IV
Instrument relevance								
First-stage F	24.63	12.96	13.98	24.63	12.96	13.98	24.63	24.63
Cragg-Donald F	43.21	24.24	15.11	43.21	24.24	15.11	43.21	43.21
Markup equilibrium assumption								
Bertrand	✓	✓	✓	✓	✓	✓	✓	✓
Cournot								
Controls								
age (linear)	✓	✓	✓	✓	✓	✓	✓	✓
age, educ., state, race, & children f.e.		✓	✓	✓	✓	✓	✓	✓
price decile & product group f.e.			✓	✓	✓	✓	✓	✓
N	8,600							

Notes: In Panel A, consumer income is measured in the HMS data and the outcome variables are measured in the grocery store sample. Panel B uses the RMS data from 2006 and 2015 and the shift-share instrument from Section IV.B. Spending weights are used in all specifications and standard errors are clustered by product modules unless noted otherwise. * denotes statistical significance at the 1% level and [†] at the 5% level.

maximize profits, taking into account how its decisions affect the price index for ℓ . [Hottman, Redding, and Weinstein \(2016\)](#) show that multiproduct firms that are large relative to the market internalize the effects of their pricing decisions on the sales of other varieties. They derive the firm markup for each product under Bertrand and Cournot competition: $\mu_{Xkf} = \frac{\varepsilon_{Xkf}}{\varepsilon_{Xkf} - 1}$, where f indexes firms, k products, $X = B/C$ refers to Bertrand/Cournot competition, and the perceived firm elasticities of demand are given by $\varepsilon_{Bkf} = \sigma^F(1 - S_{kf}^\ell) + S_{kf}^\ell$ under Bertrand and $\varepsilon_{Ckf} = \frac{1}{\frac{1}{\sigma^F} - (\frac{1}{\sigma^F} - 1)S_{kf}^\ell}$ under Cournot. S_{kf}^ℓ is the expenditure share of firm f supplying k in cell ℓ , and σ^F is the elasticity of substitution between firms.²¹

To assess the importance of variable markups, I repeat the shift-share research design while decomposing the response of inflation for continuing products between a markup effect and a residual that can be viewed as the change in cost. Measuring inflation for continuing products with CES preferences, a convenient decomposition follows:

$$(8) \quad \Delta \log(P_\ell) = \underbrace{\sum_{k=1}^n \omega_{k,t,t+1} \log\left(\frac{\mu_{Xkf,t+1}}{\mu_{Xkf,t}}\right)}_{\Delta \log(P_\ell^m)} + \underbrace{\sum_{k=1}^n \omega_{k,t,t+1} \log\left(\frac{p_{k,t+1}}{p_{k,t}} \frac{\mu_{Xkf,t}}{\mu_{Xkf,t+1}}\right)}_{\Delta \log(P_\ell^c)},$$

where $\omega_{k,t,t+1}$ is the Sato-Vartia weight (see [Online Appendix A](#)). [Equation \(8\)](#) shows that the CES inflation rate for continuing products is the sum of two CES inflation rates, one for the change in markups ($\Delta \log(P_\ell^m)$) and the other for the change in cost ($\Delta \log(P_\ell^c)$). I measure the change in markups $\frac{\mu_{Xkf,t+1}}{\mu_{Xkf,t}}$ in the RMS data using $\mu_{Xkf} = \frac{\varepsilon_{Xkf}}{\varepsilon_{Xkf} - 1}$ and then infer the change in costs by $\frac{p_{k,t+1}}{p_{k,t}} \frac{\mu_{Xkf,t}}{\mu_{Xkf,t+1}}$.

I rerun the IV estimator with the baseline shift-share instrument from [Section IV.B](#). I run IV specifications similar to [equation \(3\)](#), where the dependent variable is in turn the full CES inflation

21. Intuitively, consumers have CES preferences but each firm internalizes the effect of its pricing decisions on market price indexes and hence perceives a variable elasticity of demand. The firm's variable markup μ_{Xkf} is increasing in its expenditure share within cell ℓ , S_{kf}^ℓ . If increasing demand induces entry of new manufacturers and falling market shares for incumbents, then incumbent markups will fall.

$\Delta \log(P_\ell)$ or the component corresponding to markups $\Delta \log(P_\ell^m)$, or costs $\Delta \log(P_\ell^c)$.

Table VIII, Panel B reports the results. Columns (1) to (6) report the results for the markup component, and columns (7) and (8) show the estimates for the cost component. The IV estimates in columns (1) and (7) consider the case of Bertrand competition; since $\Delta \log(P_\ell) = \Delta \log(P_\ell^m) + \Delta \log(P_\ell^c)$, they mechanically sum up to the IV estimate with the same specification in Table V (column (3)). According to these estimates, the markup channel accounts for 57.9% of the overall price response to demand ($= \frac{-0.22}{-0.38}$). The estimate for the Bertrand markup response is significant at the 5% level, while the estimate for the cost response is not significant. The IV estimates in columns (5) and (9) repeat the exercise under the assumption of Cournot competition. In that case, the entirety of the price response is accounted for by markups. The IV estimate for the change in cost is close to 0 and insignificant. The fact that the estimates vary depending on the equilibrium assumption is not surprising: as pointed out by Hottman, Redding, and Weinstein (2016), Cournot competition yields much higher markups for larger firms than does Bertrand competition (and hence a larger fall in markups when their market shares decrease). Columns (3), (4), (6), and (7) show that the IV estimates for markups under Bertrand or Cournot remain large as the full set of fixed effects is introduced. Online Appendix Table A8 finds similar patterns when computing markups across product groups.

The results in Table VIII, Panel B indicate that changes in manufacturers' market shares induced by changes in demand (through the entry of new manufacturers across the product space) can yield substantial markup effects. Although the share of the price response explained by markups differs depending on whether Bertrand or Cournot competition is assumed, the effects are substantial. Falling markups induced by entry explain between 50% and 100% of the price response for continuing goods. These results illustrate the importance of the "procompetitive effects" of growing demand and endogenous entry.²²

22. Endogenous markups may explain why the price response is driven by changes in demand rather than by the level of demand. When demand increases in a market, there are excess returns for firms to enter this market and introduce new goods because the level of competition is low relative to the level of demand. However, as more firms enter the market, markups endogenously decline, reducing profits up to the point where there are no longer excess returns in this market. At

V. CALIBRATION: WHAT ARE THE IMPLICATIONS OF CHANGES IN THE INCOME DISTRIBUTION FOR INFLATION INEQUALITY?

To what extent can market size effects (documented in [Section IV](#)) account for the observed inflation differences between income groups (documented in [Section III](#))? I present a simple calibration focusing on the role of market size effects resulting from changes in the income distribution during the sample period. Conceptually, the calibration ignores all supply effects that could result from changes in the income distribution other than the supply response induced by changes in demand across the product space. The goal is to assess whether this type of supply response alone generates substantial inflation inequality.

1. *Framework.* As in [Section III](#), assume the existence of five groups of households with different preferences, which correspond approximately to the average income quintiles between 2004 and 2015, with L^i the number of households in i and Y^i their average income. Following [equation \(2\)](#), the change in demand induced by changes in the income distribution is: $\Delta \widehat{\log(Q_\ell)} = \sum_{i=1}^5 s_{i\ell 0} (\Delta \log(L^i) + \Delta \log(Y^i))$, where $s_{i\ell 0}$ is the share of sales to i in ℓ in the base period. From this equation, we can derive implications for price indexes across the product space, using the IV estimates from the research design in [Section IV](#) to predict Törnqvist inflation for continuing products and CES inflation with changes in product variety. $\beta_\Pi = -0.618$ is the response of the CES index to a change in demand, accounting for product variety (column (10) of [Table V](#)), while $\beta_\pi = -0.431$ gives the response of the Törnqvist index for continuing products (column (2) of [Online Appendix Table V](#)). Using these point estimates gives:

$$(9) \quad \begin{cases} \widehat{\pi}_\ell = \beta_\pi \Delta \widehat{\log(Q_\ell)}, \\ \widehat{\Pi}_\ell = \beta_\Pi \Delta \widehat{\log(Q_\ell)}. \end{cases}$$

this point, the process of targeted entry stops; the fact that the market is larger doesn't lead to a constantly higher flow of entrants because there are both more consumers and more firms in this market, such that profits per firm are equalized across markets in equilibrium. [Online Appendix B](#) presents a model describing this mechanism formally.

One can then obtain a first-order approximation to the implied inflation rates for each income quintile:

$$(10) \quad \begin{cases} \widehat{\pi}_i = \sum_{\ell=1}^L s_{\ell 0}^i \widehat{\pi}_\ell, \\ \widehat{\Pi}_i = \sum_{\ell=1}^L s_{\ell 0}^i \widehat{\Pi}_\ell, \end{cases}$$

where $s_{\ell 0}^i$ is the spending share of i on ℓ in the base period, $\widehat{\pi}_i$ is implied inflation for continuing products for income quintile i , and $\widehat{\Pi}_i$ also takes into account changes in product variety.

I use IPUMS USA to measure the average annual log change in market size for group i , $\Delta \log(L^i) + \Delta \log(Y^i)$, between two four-year periods at the beginning and end of the sample (2000–2004 versus 2012–2016). I measure $s_{i\ell 0}$ and $s_{\ell 0}^i$ in 2004. As before, the cells ℓ of the product space are product modules by price deciles. Given that the approach described above is not informative about the average inflation rate, I add a constant (common to all cells ℓ) such that the predictions $\widehat{\pi}_\ell$ and $\widehat{\Pi}_\ell$ match the average inflation rates observed in the data between 2004 and 2016.

2. Results. Figure VII and Table IX report the results. Figure VII, Panel A plots the patterns of predicted inflation across the product space for the Törnqvist index and CES index (from equation (9)) against the 2004 spending share of the top income quintile relative to the bottom income quintile. There is a strong negative relationship: cells of the product space that sell relatively more to the top income quintile have much lower predicted inflation. This is because demand from the top income quintile grows faster in IPUMS USA and because the estimates from Table V show that increasing demand leads to lower inflation.

Figure VII, Panel A shows that the magnitude of the effect is substantial. For instance, when a product module by price sells 80% to the top income quintile, predicted annual Törnqvist inflation is under 1%, whereas when the sales share of the top income quintile is only 20% predicted annual Törnqvist inflation is close to 2%. Depicting the patterns across the product space illustrates the strength and stability of the relationship. But what

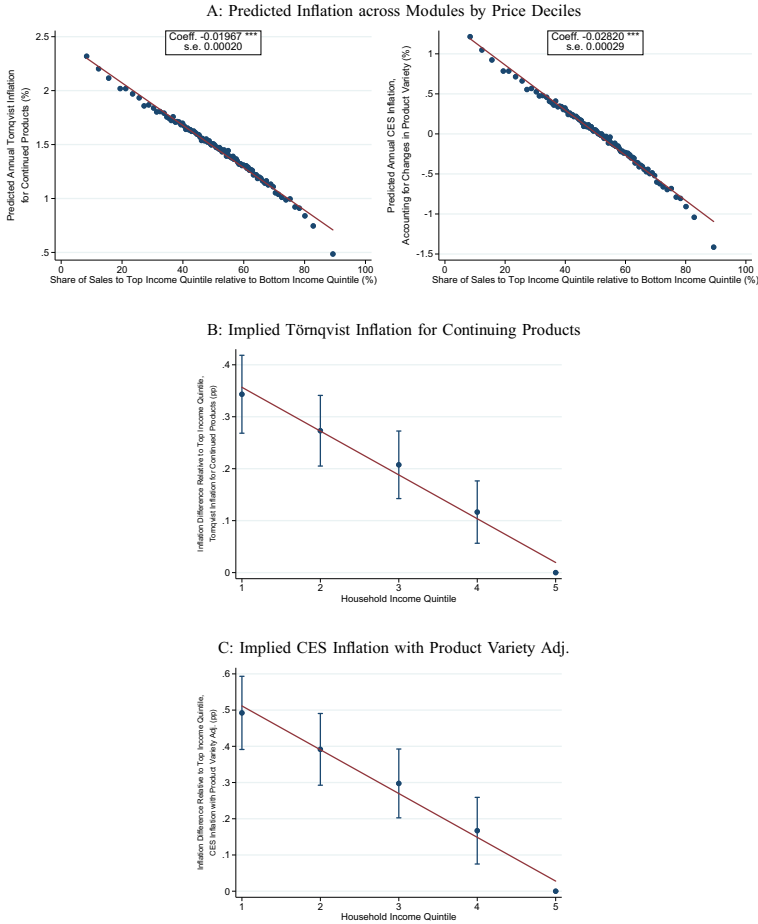


FIGURE VII

Inflation Inequality Implied by Changes in the Income Distribution

Panels A, B, and C report the results of the calibration presented in Section V.A. Continuing products are observed across consecutive years. In Panel A, the regressions use spending weights and standard errors are clustered by product modules. The 95% confidence intervals are obtained by bootstrap.

are the implications for differences in inflation between income groups?²³

23. From equation (10), the inflation difference between the top and bottom income quintiles can be written $\Delta \hat{\pi}^{Q5-Q1} = \sum_{\ell=1}^L (s_{\ell 0}^{Q5} - s_{\ell 0}^{Q1}) \hat{\pi}_{\ell}$. Denoting by μ_{ℓ}^{Q5} the share of sales going to the top income quintile in ℓ relative to the bottom

Figure VII, Panel B reports the implied Törnqvist inflation difference relative to the top income quintile. The figure shows that lower-income groups have larger predicted inflation. The magnitude is substantial: the difference between the bottom and top income quintiles is 34 basis points a year. On Figure VII, panel C the patterns for CES inflation accounting for product variety are similar. CES inflation implied by market size effects is significantly larger for lower income groups. The difference between the top and bottom income quintiles is about 50 basis points a year.

Table IX presents an explicit comparison between the inflation patterns implied by market size effects and the observed inflation differences reported in Section III. Columns (1) to (4) report the results for Törnqvist and columns (5) to (8) for CES. For Törnqvist, the full inflation difference between the top and bottom income quintiles is 0.66 percentage points a year. This number falls to 0.48 at the level of product modules by price deciles. This level of aggregation is the relevant benchmark for the patterns of predicted inflation: the calibration is conducted at the level of product modules by price deciles because the IV estimates are available across these cells of the product space. The predicted Törnqvist inflation difference is 0.34 percentage points a year, which is 70.8% of the observed difference at the same level of aggregation ($= \frac{0.34}{0.48}$) and 51.5% of the full observed inflation difference ($= \frac{0.34}{0.66}$).

The patterns are very similar for the CES inflation difference accounting for changes in product variety. While the full CES inflation difference between the bottom and top income quintiles is 0.88 percentage points a year, it falls to 0.66 when only patterns across product modules by price deciles are taken into

income quintile (as on the x -axis of Panel A), it can be checked that $\Delta \hat{\pi}^{Q5-Q1} = \frac{\text{Var}(\mu_l^{Q5})}{\bar{\mu}^{Q5}(1-\bar{\mu}^{Q5})} \beta_{\hat{\pi}_l, \mu_l^{Q5}}$ (see Borusyak and Jaravel 2018). $\beta_{\hat{\pi}_l, \mu_l^{Q5}}$ is the OLS slope of predicted inflation $\hat{\pi}_l$ regressed on μ_l^{Q5} with spending weights, which is shown in Panel A. To obtain the overall implied inflation difference between the top and bottom income quintiles, this slope is scaled by $\frac{\text{Var}(\mu_l^{Q5})}{\bar{\mu}^{Q5}(1-\bar{\mu}^{Q5})}$, where $\bar{\mu}^{Q5}$ is the share of total sales to the top income quintile. This ratio is increasing in the degree of heterogeneity in spending patterns between the top and bottom income quintiles. Intuitively, if income groups tend to purchase similar goods, even though the slope on Figure VIII, Panel A is large, income groups will devote most of their expenditures to products that are located very close to each other on the x -axis, hence the overall inflation difference between income groups will be small.

TABLE IX
INFLATION INEQUALITY IMPLIED BY CHANGES IN THE INCOME DISTRIBUTION

Income quintile	Törnqvist inflation for continuing products (avg., 2004–2015)			CES inflation with product variety adj. (avg., 2004–2015)			
	Level (%) (1)	Inflation diff. rel. to top quintile (pp) (2)	Explained (3)	Level (%) (5)	Inflation diff. rel. to top quintile (pp) (6)	Explained (8)	
		Full (2)	Across modules by price deciles (3)		Full (6)	Across modules by price deciles (7)	
1st (< \$20k)	1.87 (0.043)	0.66 (0.062)	0.48 (0.063)	0.54 (0.067)	0.88 (0.103)	0.66 (0.077)	0.49 (0.051)
2nd (\$20k — \$40k)	1.66 (0.032)	0.45 (0.065)	0.31 (0.043)	0.23 (0.067)	0.58 (0.089)	0.42 (0.066)	0.39 (0.051)
3rd (\$40k — \$60k)	1.46 (0.032)	0.25 (0.037)	0.21 (0.032)	-0.02 (0.075)	0.33 (0.068)	0.28 (0.045)	0.30 (0.048)
4th (\$60k — \$100k)	1.30 (0.028)	0.09 (0.035)	0.10 (0.032)	-0.23 (0.068)	0.12 (0.054)	0.13 (0.044)	0.17 (0.046)
5th (> \$100k)	1.21 (0.031)	—	—	-0.35 (0.062)	—	—	—

Notes. This table reports the results of the calibration presented in Section VA. Standard errors are obtained by bootstrap.

account. At this level of aggregation, the predicted CES inflation difference is 0.49 percentage points, which is 74.2% of the relevant benchmark ($= \frac{0.49}{0.66}$) and 55.6% of the full CES inflation difference ($= \frac{0.49}{0.88}$).

VI. CONCLUSION

This article has shown that product variety and inflation respond to changes in market size. In a period of economic growth or rising income inequality, these dynamics disproportionately benefit high-income households. These findings stand in contrast with the product cycle view, according to which innovation tends to benefit everyone equally. This article showed that the product cycle view is incomplete by studying consumer packaged goods, but further research is needed to characterize comprehensively the relative importance of the product cycle and increasing product variety across all consumption categories.²⁴

These findings have two potential implications for public policy. First, accurate measurement of inflation across income groups appears important. Large and sustained inflation differences were found across income groups in the retail sector using the HMS/RMS data but also more broadly using the CEX-CPI data. To appropriately index government transfers, income poverty thresholds, and tax brackets, it appears essential for statistical agencies to improve on their ability to measure income-group-specific spending patterns across detailed consumption categories.

Second, for cost-benefit analysis it may be important to account for the supply response to market size changes induced by policy. Many policy instruments, such as food stamps, income taxes, minimum wage laws, or Social Security benefits, affect the relative market size of different groups of households. These changes in market size can induce a targeted response of supply, whose price effects will determine the equilibrium real effects of the policy change.

LONDON SCHOOL OF ECONOMICS

24. [Online Appendix Figure A12](#) presents evidence that the product cycle is not a central feature of consumer packaged goods, but it could be very different in other sectors, such as cars and consumer electronics.

SUPPLEMENTARY MATERIAL

An [Online Appendix](#) for this article can be found at *The Quarterly Journal of Economics* online. Code replicating tables and figures in this article can be found in [Jaravel \(2018\)](#), in the Harvard Dataverse, doi:[10.7910/DVN/JXGYOB](https://doi.org/10.7910/DVN/JXGYOB).

REFERENCES

- Acemoglu, Daron, "Equilibrium Bias of Technology," *Econometrica*, 75 (2007), 1371–1409.
- Acemoglu, Daron, Philippe Aghion, Leonardo Bursztyn, and David Hemous, "The Environment and Directed Technical Change," *American Economic Review*, 102 (2012), 131–166.
- Acemoglu, Daron, and Joshua Linn, "Market Size in Innovation: Theory and Evidence from the Pharmaceutical Industry," *Quarterly Journal of Economics*, 119 (2004), 1049–1090.
- Adao, Rodrigo, Michal Kolesar, and Eduardo Morales, "Shift-Share Designs: Theory and Inference," Working Paper, Princeton University, 2018.
- Aghion, Philippe, and Peter Howitt, "A Model of Growth through Creative Destruction," *Econometrica*, 60 (1992), 323–351.
- Amble, Nathan, and Ken Stewart, "Experimental Price Index for Elderly Consumers," *Monthly Labor Review*, 117 (1994), 11.
- Argente, David, and Munseob Lee, "Cost of Living Inequality during the Great Recession," Working Paper, University of Chicago, 2016.
- Autor, David H., Lawrence F. Katz, and Melissa S. Kearney, "Trends in US Wage Inequality: Revising the Revisionists," *Review of Economics and Statistics*, 90 (2008), 300–323.
- Boppart, Timo, and Franziska Weiss, "Non-Homothetic Preferences and Industry Directed Technical Change," Working Paper, Institute for International Economic Studies, 2013.
- Borusyak, Kirill, Peter Hull, and Xavier Jaravel, "Quasi-Experimental Shift Share Designs," Working Paper, London School of Economics, 2018.
- Borusyak, Kirill, and Xavier Jaravel, "The Distributional Effects of Trade: Theory and Evidence from the United States," Working Paper, London School of Economics, 2018.
- Broda, Christian, and John Romalis, "The Welfare Implications of Rising Price Dispersion," Mimeo, University of Chicago, 3, 2009.
- Broda, Christian, and David Weinstein, "Globalization and the Gains From Variety," *Quarterly Journal of Economics*, 121 (2006), 541–585.
- , "Product Creation and Destruction: Evidence and Price Implications," *American Economic Review*, 100 (2010), 691–723.
- Bronnenberg, Bart J., Jean-Pierre H. Dubé, and Matthew Gentzkow, "The Evolution of Brand Preferences: Evidence from Consumer Migration," *American Economic Review*, 102 (2012), 2472–2508.
- Comin, Diego, Danial Lashkari, and Martí Mestieri, "Demand-Pull, Technology-Push, and the Sectoral Direction of Innovation," Working Paper, Dartmouth University, 2017.
- DellaVigna, Stefano, and Joshua M. Pollet, "Demographics and Industry Returns," *American Economic Review*, 97 (2007), 1667–1702.
- Diewert, Erwin, "Exact and Superlative Index Numbers," *Journal of Econometrics*, 4 (1976), 115–145.
- Dixit, Avinash K., and Joseph E. Stiglitz, "Monopolistic Competition and Optimum Product Diversity," *American Economic Review*, 67 (1977), 297–308.

- Dubé, Jean-Pierre, Jeremy T. Fox, and Che-Lin Su, "Improving the Numerical Performance of Static and Dynamic Aggregate Discrete Choice Random Coefficients Demand Estimation," *Econometrica*, 80 (2012), 2231–2267.
- Eizenberg, Alon, "Upstream Innovation and Product Variety in the US Home Pc Market," *Review of Economic Studies*, 81 (2014), 1003–1045.
- Erickson, Tim, and Ariel Pakes, "An Experimental Component Index for the CPI: From Annual Computer Data to Monthly Data on Other Goods," *American Economic Review*, 101 (2011), 1707–1738.
- Faber, Benjamin, and Thibault Fally, "Firm Heterogeneity in Consumption Baskets: Evidence from Home and Store Scanner Data," Working Paper, Berkeley University, 2017.
- Feenstra, Robert, "New Product Varieties and the Measurement of International Prices," *American Economic Review*, (1994), 157–177.
- Feenstra, Robert C., and David E. Weinstein, "Globalization, Markups, and US Welfare," *Journal of Political Economy*, 125 (2017), 1040–1074.
- Garner, Thesia I., David S. Johnson, and Mary F. Kokoski, "An Experimental Consumer Price Index for the Poor," *Monthly Labor Review*, 119 (1996), 32.
- Garner, Thesia I., Robert McClelland, and William Passero, "Strengths and Weaknesses of the Consumer Expenditure Survey from a BLS Perspective," Draft Presented at National Bureau of Economic Research Summer Institute, 2009.
- Gopinath, Gita, Pierre-Olivier Gourinchas, Chang-Tai Hsieh, and Nicholas Li, "International Prices, Costs, and Markup Differences," *American Economic Review*, 101 (2011), 2450–2486.
- Handbury, Jessie, and David E. Weinstein, "Goods Prices and Availability in Cities," *Review of Economic Studies*, 82 (2014), 258–296.
- Hausman, Jerry, "Sources of Bias and Solutions to Bias in the Consumer Price Index," *Journal of Economic Perspectives*, 17 (2003), 23–44.
- Hayek, Friedrich A., "Prices and Production," 1931.
- Hobijn, Bart, and David Lagakos, "Inflation Inequality in the United States," *Review of Income and Wealth*, 51 (2005), 581–606.
- Hottman, Colin J., Stephen J. Redding, and David E. Weinstein, "Quantifying the Sources of firm Heterogeneity," *Quarterly Journal of Economics*, 131 (2016), 1291–1364.
- Jaravel, Xavier, "Replication Data for: 'The Unequal Gains from Product Innovations: Evidence from the U.S. Retail Sector'," (2018). Harvard Dataverse, doi: 10.7910/DVN/JXGYOB.
- Jones, Charles, "R & D-Based Models of Economic Growth," *Journal of Political Economy*, (1995), 759–784.
- Kaplan, Greg, and Sam Schulhofer-Wohl, "Inflation at the Household Level," *Journal of Monetary Economics*, (2017).
- Kopczuk, Wojciech, Emmanuel Saez, and Jae Song, "Earnings Inequality and Mobility in the United States: Evidence from Social Security Data Since 1937," *Quarterly Journal of Economics*, 125 (2010), 91–128.
- Krugman, Paul R., "Increasing Returns, Monopolistic Competition, and International Trade," *Journal of International Economics*, 9 (1979), 469–479.
- Matsuyama, Kiminori, "The Rise of Mass Consumption Societies," *Journal of Political Economy*, 110 (2002), 1035–1070.
- McGranahan, Leslie, and Anna Paulson, "Constructing the Chicago Fed Income Based Economic Index–Consumer Price Index: Inflation Experiences by Demographic Group: 1983–2005," Federal Reserve Bank of Chicago Working Paper, 2006.
- Piketty, Thomas, and Emmanuel Saez, "Income Inequality in the United States, 1913–1998," *Quarterly Journal of Economics*, 118 (2003), 1–41.
- Romer, Paul, "Endogenous Technological Change," *Journal of Political Economy*, 98 (1990), S71–S102.
- Ruggles, Steven, Katie Genadek, Ronald Goeken, Josiah Grover, and Matthew Sobek, "Integrated Public Use Microdata Series: Version 7.0," University of Minnesota, 2017.

- Sato, Kazuo, "The Ideal Log-Change Index Number," *Review of Economics and Statistics* (1976), 223–228.
- Schmookler, Jacob, "Invention and Economic Growth," 1966.
- Schumpeter, Joseph, "*Capitalism, Socialism and Democracy*." (New York: Harper and Bros, 1942).
- Shleifer, Andrei, "Implementation Cycles," *Journal of Political Economy*, 94 (1986), 1163–1190.
- Vartia, Yrjö O., "Ideal Log-Change Index Numbers," *Scandinavian Journal of Statistics*, (1976), 121–126.
- Vernon, Raymond. "International Investment and International Trade in the Product Cycle," *Quarterly Journal of Economics*, (1966), 190–207.
- Violante, Giovanni, "Skilled-Biased Technical Change," In *New Palgrave Dictionary of Economics*, (2008).