Mobile Broadband and Consumer Price Response

John Reaves*†‡

October 23, 2026

See the most current version here.

Abstract

Households increasingly concentrate their spending on a smaller set of familiar products, which may give firms more power over loyal consumers. Theory suggests that Internet access could hasten this trend: consumers can more easily compare prices for their favorite products, and advertisers can use greater data availability to target ads and coupons. In this paper, I study how smartphone access affects consumers at the grocery store. Using state-level random forests, I construct the first public, national dataset on 3G rollout, a prerequisite for smartphone use. I find that smartphone access shifts purchases toward smaller brands and weakly reduces household spending concentration, as consumers buy more distinct brands and products. Counter to general trends, I find that consumers become more likely to purchase products where they differ from the "typical" consumer. I also find no significant change in posted prices or in-store offerings. Demand estimation shows that smartphones shrink price elasticities but do not generally affect brand relevance. These results run counter to the information mechanisms usually proposed for smartphones' effect on shopping. Instead, I propose in-store distraction as the most likely explanation.

JEL Codes: L19, R13

Keywords: 3G, smartphones, price elasticity, demand estimation, Internet, mobile broadband

1 Introduction

Consumers often find it easier to purchase something familiar rather than searching every time they go to the store. Inertia, the tendency to buy repeatedly from the same brand, is widespread—even in undifferentiated markets like gasoline (MacKay and Remer, 2026). Over the last twenty years, households have reached more for familiar options, and these habitual choices concentrated their spending. Grocery consumers now devote larger portions of their total spending to fewer products (Neiman and Vavra, 2023). They also shop at fewer distinct retailers (Leung and Li, 2026). Though market-level concentration (HHI) fell over the same period (Benkard, Yurukoglu, and A. L. Zhang, 2026), lower consumer experimentation could give firms more

*Department of Economics, Michigan State University — 486 W Circle Dr, East Lansing, MI 48824.

†I am grateful to Oren Ziv, Stacy Dickert-Conlin, Paul Kim, Steven Miller, and many others for their support. This work benefited from comments received at the 2026 Midwest Economics Association conference and the Michigan State Graduate Student Seminar.

‡Researcher(s)' own analyses calculated (or derived) based in part on data from Nielsen Consumer LLC and marketing databases provided through the NielsenIQ Datasets at the Kilts Center for Marketing Data Center at The University of Chicago Booth School of Business. The conclusions drawn from the NielsenIQ data are those of the researcher(s) and do not reflect the views of NielsenIQ. NielsenIQ is not responsible for, had no role in, and was not involved in analyzing and preparing the results reported herein.

power over loyal consumers. Consistent with this, markups rose by 30%, lowering consumer surplus for those earning below the median income (Döpper et al., 2025).

These trends coincide with large technological shifts. The Internet transformed search, and smartphones brought the Internet into physical stores. Smartphones let consumers compare prices while standing in the grocery aisle, potentially paying less by switching retailers instead of products (Morton, 2006). At the same time, smartphone adoption intensifies advertiser data collection (Clifford, 2009) and gives firms a platform to target consumers with ads and coupons. Both smartphone mechanisms could raise within-household concentration. Consumers can save on their favorites without trying substitutes, and firms can digitally encourage repeat purchases. Limited evidence from home Internet adoption is consistent with this narrative: Tran (2026) finds that households with Internet at home buy from fewer brands. But evidence evaluating the underlying mechanisms is largely absent, along with potential supply-side effects on prices or in-store offerings. Mobile Internet is accessible at the point of purchase in a unique way, and its lower barrier to entry may influence consumers—and firms—differently than desktop browsing.

In this paper, I explore how smartphone access affects consumers at the grocery store. Using the NielsenIQ Consumer Panel, I start with three facts: top brands' total market share is falling, household spending is becoming more concentrated, and brands' consumer bases are homogenizing over time. To see if smartphone access contributes to these three trends as theory would suggest, I first perform a staggered difference-in-differences analysis. I conclude that smartphone access decreases top brands' market share, but also weakly *decreases* household spending concentration, and leads consumers to buy more products where they are not in the consumer base's demographic majority. These results are inconsistent with smartphones hastening household spending concentration. Smartphone access instead changes what people buy, increasing the number of distinct brands and products that consumers purchase. While I find no significant supply-side effects in equilibrium, shifts in demand could mask potential changes. I therefore estimate demand for beer, cereal, and toilet paper to better account for possible changes. Consumer price elasticities tend to shrink with smartphone access, with larger movements in higher-elasticity markets. The effects of smartphone access are largely inconsistent with an information story for either firms or consumers—rather, as in Grewal et al. (2018), I propose in-store distraction as the likeliest mechanism for these behavioral changes.

To understand the effects of smartphone adoption on shoppers, I combine shopping data from NielsenIQ with a novel dataset on mobile Internet rollout. For consumers to effectively use their phones in the grocery aisle, I assume that consumers need both a smartphone (such as the iPhone) and 3G or faster mobile coverage, as older mobile standards are too slow to effectively use. Building on Hersh, B. J. Lang, and M. Lang (2022), I combine records from the NTIA's State Broadband Initiative (SBI), the FCC's Form 477, and the Antenna Structure Registration (ASR) program to create the first national, public dataset on 3G rollout in the United States from 2005-2017. I use machine learning to extend existing data on 3G coverage into the past, developing a model that is highly congruent with out-of-sample data and that allows me to track spatial and temporal variation in coverage. By combining zip-code-level predictions with the consumer and scanner data panels from NielsenIQ, I determine when consumers gained access to smartphones and track how access affects shopping behavior.

I next adopt a staggered difference-in-differences approach, aggregating zip-code-level differences in the timing of 3G/smartphone access to the city-quarter level. I first test whether smartphones contribute to observed trends. While the share purchased from top brands falls with smartphone access, I find that smartphones generally move *against* the second and third broader trends. Household spending concentration does not significantly grow, with increases of over 0.36% ruled out by my estimates. I find stronger evidence

against smartphones contributing to consumer homogenization. Because only some portions of brands' consumers have smartphones at a given time, I instead study consumers' propensity to buy products where most consumers demographically resemble them. By most metrics, consumers buy these products at weakly *lower* rates after gaining access to smartphones. I then explore other aspects of consumer behavior that smartphones may affect. I find no evidence of increasing coupon or sale usage. On the retailer side, market-level characteristics like the number of in-store options, average price, and standard deviation of prices are not consistently affected. Consumers are instead more likely to shop at their "favorite" (modal) retailer, suggesting they are less likely to "shop around" at other stores. Further evidence points to smartphones primarily changing consumer purchase decisions, as consumers purchase from more brands and product module codes.

While reduced form estimations are useful, they also reflect equilibrium outcomes. My findings thus far are the result of shifts in both supply and demand. To move beyond these equilibrium objects, I estimate demand using the approach outlined in Conlon and Gortmaker (2025) with additional micro moments and optimal price instruments. I examine several product markets, but focus on three in this paper: beer, cereal, and toilet paper. Demand estimations for all three products indicate that smartphones generally shrink own-price elasticities, with larger movements for cities/retailers starting out with larger elasticities. This pattern of movement indicates a small level of cross-market convergence.

Taking all these results together challenges the traditional narratives of how smartphones might change consumer shopping behavior. Given that they shrink price elasticities, it is unlikely that consumers are using them to compare prices across retailers. Without evidence of increased coupon use or repeat purchasing behavior, it also seems unlikely that firms are successfully leveraging this new technology to their advantage. I instead propose that smartphones may not be useful during the shopping process at all. Rather, distracted consumers may pay less attention to prices and find themselves examining shelves in unfamiliar parts of the store—consistent with my findings.

This paper contributes to a growing literature studying long-run trends in the retail industry. A dramatic rise in product markups from 2006-2019 seems to be driven by declining marginal production costs and less elastic consumer demand (Döpper et al., 2025). Households also concentrate more of their shopping at fewer distinct retailers (Leung and Li, 2026; Smith and Ocampo, 2025; Tran, 2026) and spend more on fewer products (Neiman and Vavra, 2023). In this paper, I am the first to document changes in consumer-base diversity, showing that the average brand is attracting more homogeneous customers. In addition, while Tran (2026) explores the effects of home Internet on shopping patterns, my paper is the first to study mobile (possibly in-store) Internet use in this setting.

I also build on the literature on the effect of Internet access on consumer behavior, exploring changes in what consumers buy and how they buy it. Prior studies are largely theoretical or experimental. Showrooming sees consumers use physical stores only as "showrooms" to browse before making a final purchase online (Bar-Isaac and Shelegia, 2023; Chan et al., 2021; Gensler, Neslin, and Verhoef, 2017). Experimental evidence from Grewal et al. (2018) suggests that smartphone use can also distract shoppers, leading them to purchase more distinct products and experiment more through deviation from "autopilot" shopping paths. Here, I extend Grewal et al. (2018) by exploring *initial* access, rather than the effects of using a familiar technology while shopping. I also examine what customers buy and how they buy it in greater detail—for instance, whether a product was on sale, if the consumer used a coupon, or if they bought the cheapest options.

I additionally supplement the literature on shifts in consumer price elasticities. Advertising is the most common outside intervention. Ads are typically thought to act on marginal consumers, shifting the demand

curve to the right. The incorporation of these price-elastic, marginal consumers leads to a higher average price elasticity. However, ads based on brand differentiation (instead of low prices) can have the opposite impact, making consumers care more about non-price attributes (Ataman et al., 2023; Erdem, Keane, and Sun, 2008; Kaul and Wittink, 1995; Krishnamurthi and Raj, 1985). The same is true when ads are targeted at a particular consumer, depending on whether the advertisement is intended for new or existing consumers (Bart, Stephen, and Sárváry, 2012; Book and Wallach, 2015; Lau, 2020). But ads are not the only exogenous shifter. Common determinants of consumer price elasticity include demographics (Hoch et al., 1995) and prior purchase behavior (Kim, Srinivasan, and Wilcox, 1999). Most similar to this paper, however, is a study of e-commerce shifting offline price responsiveness. Biswas, Yoganarasimhan, and H. Zhang (2025) find that online shopping options increase offline consumer price sensitivity in product markets with low switching costs. In my paper, I expand the limited literature on how online access affects offline behavior by being the first to study mobile Internet as a price-elasticity intervention.

Finally, by training random forest models and extending mobile Internet coverage backwards in time, I also improve researchers' knowledge of 3G's national rollout. Existing data on the spread of mobile broadband are proprietary, limited, or both. As an example, Collins Bartholomew (a major proprietary source of mobile Internet coverage across time) includes worldwide coverage measures based on broadband provider reports. However, because these reports are voluntary, these data are necessarily biased downwards.[1] Figure 1 in Dagorret and Guo (2023) shows that Collins Bartholomew's United States coverage in 2010 effectively includes only urban centers. However, this would both attenuate treatment effect estimates and contaminate them with the urban-rural divide. Contemporaneous press coverage of the AT&T-Verizon "color war" in Hansell (2009) suggests actual coverage was much higher, with maps that closely resemble government records[2]. To create my dataset, I use high-quality information from the federal government, as in Myers and Hooper (2026) and Hersh, B. J. Lang, and M. Lang (2022), to get more accurate estimates of 3G coverage across time. The code and data are freely available on my website.[3]

The remainder of the paper is organized as follows. Section 2 describes the NielsenIQ data. Section 3 details three major stylized facts about the changing retail environment. Section 4 provides more background on historical mobile broadband rollout, along with details on the data and machine learning I use to determine 3G/smartphone access. Section 5 presents my reduced-form findings, and Section 6 summarizes the results of my demand estimations in the beer, cereal, and toilet paper markets. Section 7 concludes by further discussing potential mechanisms and broader implications.

2 Data

The primary data sources for this project come from NielsenIQ—specifically, the Consumer Panel and Retail Scanner datasets. The Consumer Panel consists of a stratified, proportionate, randomly-recruited sample of households across the United States, weighted and balanced according to their demographic characteristics, with the goal of collecting a nationally-representative sample.[4] Consumer data include demographic characteristics of the household and household head(s), as well as detailed shopping records across the year. These

[1] This is mechanical, and not intended as a criticism. Note that Collins Bartholomew inherited their older "Legacy" data from the GSMA, and had no hand in compiling data from the 2000s. Data quality will also vary by country; where more of the "big players" reported their data, the result will be higher-quality.

[2] See Panel (a) of Figure 5, which visualizes the December 2010 data from the NTIA.

[3] Interested researchers should cite this paper, and can reach the resources here on my website.

[4] These characteristics are: household size, household income, head of household age, race, Hispanic origin, male head education, female head education, head of household occupation, presence of children, and county size.

include where the consumer purchased from (with better records for retailers participating in NielsenIQ data collection), what they purchased, and when they did so. The NielsenIQ panelist data have a projection factor associated with each household, to scale their stratified sample up to a nationally representative one. I therefore use the projection factor to weight averages where appropriate.

The NielsenIQ Retail Scanner data are collected from participating retailers. Participants represent a large portion of the grocery market, with 30,000-50,000 individual stores and roughly 90 retail chains participating. The scanner data feature sales at the UPC-store-week level from 2006-2019,[5] providing a rich sample of retailers across the United States. Market shares and prices within markets are derived from this scanner data.

When examining particular product markets, I focus on three primary goods: beer, cereal, and toilet paper. I chose these goods because they are purchased frequently enough that data on purchases of such products are rich, but are not necessarily "prime" grocery purchases (like milk or eggs). Consumers can also purchase them in bulk as they either do not expire or do not do so quickly, so the resulting product-level price can exhibit more substantial price variation than lower-priced goods. Table 1 includes summary information about these products. Ready-to-eat cereal (RTE cereal, or just "cereal") is the most commonly purchased good of the three, with 10.4% of recorded trips including a purchase of cereal; toilet paper and beer are purchased on 5.3 and 1.5% of trips, respectively. However, cereal purchases tend to be inexpensive, with an average price around $4.41. All markets have fairly substantial product variety available to consumers, though beer is the most diverse both in terms of what consumers can choose from (62 brands on average) and the broader market (almost 3,400 brands attested). Note that the number of markets recorded for beer is lower, though, due to differing state-level restrictions on alcohol sales.[6]

Table 1: Summary statistics by product

	Retailers	Cities	Quarters	Markets	Unique Prods.	Unique Brands	Price	Products*	Brands*	Inside Sh.*
Beer	145	205	44	40246	6653	3385	13.074	125.079	62.208	0.015
							(9.147)	(105.945)	(52.764)	(0.005)
RTE Cereal	169	205	44	58603	5748	1381	4.412	192.678	110.706	0.104
							(1.54)	(120.267)	(63.404)	(0.03)
Toilet paper	168	205	44	58763	2595	146	10.875	59.1	10.506	0.053
							(5.856)	(19.707)	(2.073)	(0.014)

Notes: Asterisks indicate a market-level measure; otherwise, measures are cross-market aggregates. Averages are calculated across all product size categories, reflecting overall rather than per-unit prices. Standard deviations are below averages in parentheses.

Markets are generally defined at the Nielsen DMA-retailer-year-quarter level, with the 206 Nielsen designated market areas (Nielsen DMAs) serving a similar role to a metro-/micropolitan area. As such, I will generally refer to Nielsen DMAs as "cities" from this point forward. Due to changes in data collection in 2007 and 2018 (and limited interest for this setting), I use data spanning the years 2007-2017, for a total of eleven observed years and 44 year-quarters, across 205 cities (as there are no recorded retailers in the Palm Springs area). The total number of markets for most products is between 40,000 and 58,000, though not all retailers are observed in every year/city. NielsenIQ collects sales information from many different types of store, including grocery stores, mass-market retailers,[7] convenience stores, and drugstores. I only consider

[5] A change to the way that the scanner data are organized in 2020 on, along with the effects of COVID-19, make 2019 a logical stopping point.

[6] One extreme involves state-owned liquor stores, some of which may be excluded due to my focus on grocery shopping.

[7] This classification also includes pet stores.

grocery stores and mass-market retailers for the purposes of this study.[8]

Following Döpper et al. (2025), who deflate prices to 2010 Q1 values when pooling estimation across time, I scale prices using the consumer price index (CPI) for 2024. This allows for more logical cross-time comparisons within and between markets.

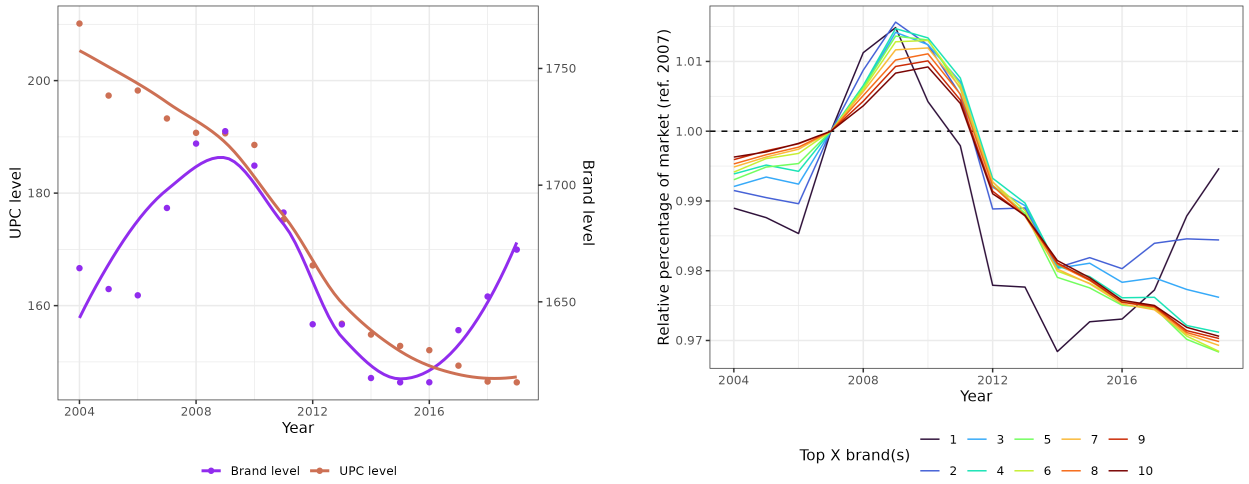
3 Stylized facts

In this section, I use the NielsenIQ Consumer Panel to outline broad changes in consumer behavior from the mid-2000s to 2019. Using these data, I find that brand concentration[9] as measured by the Herfindahl-Hirschman Index (HHI) is generally decreasing over this time period, seemingly due to increasing influence of small brands; the relative share of the top brands (whether the top-selling single brand or the top ten together) decreases from 2009-2019. This fall in concentration seems to be explained by differences between households, not within them, as household-level HHI measures indicate a *higher* concentration of spending. The average consumer base of a product/brand is also becoming more homogeneous over time, showing that some inter-household correlations in spending persist.

In other words, consumers are increasingly specializing into distinct spending silos, in ways that distinguish themselves from other, dissimilar consumers.

3.1 Retail grocers

Figure 1: Changes in grocery concentration



(a) Aggregate HHI, NielsenIQ Consumer Panel

(b) Relative share of top brands

Notes: Averages are calculated using the NielsenIQ Consumer Panel, averaging across all non-Magnet purchases in the panel and using the NielsenIQ projection factor for national representativeness. "UPC" (universal product code) indicates a unique product.

The grocery industry has seen falling brand concentration over the last twenty years. Panel (a) of Figure 1 shows that the average product market is increasingly competitive, as measured by the following HHI

[8]This helps to explain "missing" data in the beer market for states such as Pennsylvania, which have state-owned liquor stores. These stores would show up differently in the NielsenIQ data.

[9]I measure concentration throughout this paper at the brand or product level, not the firm level. Firms may have multiple brands active in the same market, though consumers may not necessarily know which brands are owned by which companies.

formula:

$$HHI_{pmc} = \sum_{j=1}^{N} \left(\frac{100 \cdot sales_{j,pmc}}{\sum_{j'=1}^{N} sales_{j',pmc}} \right)^2$$
$$HHI_{agg} = \sum_{pmc} HHI_{pmc} \cdot \frac{sales_{pmc}}{\sum_{pmc'} sales_{pmc'}}$$

where N is the total number of products/brands across all markets and *pmc* indicates the product module code[10] that product/brand j occupies.

Product-level HHI falls from a high of over 200 points in the early 2000s to around 150 fifteen years later. While product-level concentration is extremely low (in part because firms can offer many distinct products), the same phenomenon occurs from 2008-2016 at the brand level; HHI measures fall from a value of over 1700 in 2008 to around 1620 by 2016, though values climb back up in the next few years. Panel (b) indicates that among several possible definitions of product market-level "top" brands, smaller players are benefiting more. Relative to their 2007 share of the market, the highest-selling brand down to the top ten are all losing ground to smaller competitors.

It is important to note that these products are primarily substitutes on store shelves, not online. Online shopping does grow substantially over the 2010s, but this growth does not occur evenly across all goods, or across time. By 2019, with a mature online retail industry, the NielsenIQ consumer panel records that just 4% of recorded spending was spent while shopping online. For the products I study here, I therefore doubt that online shopping posed a realistic alternative for most consumers.

The market-level movements toward greater diversity could be driven by changes in household-level spending, as households "experiment" more by purchasing more products, or from changes in behavior between different households. To understand what is behind this increased diversity, I turn to the consumer side.

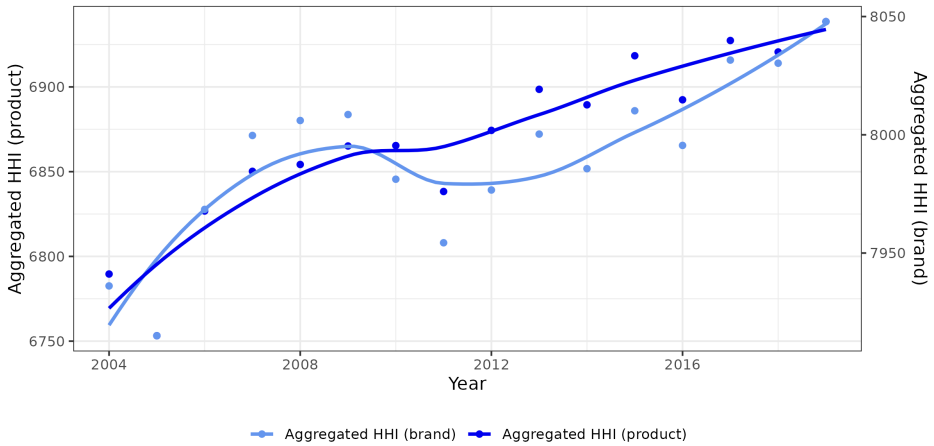
3.2 Consumers

To understand how households spend their money, I treat their spending in a given product market as that market's size for the household. While authors like Leung and Li (2026) employ measures of "household HHI", those metrics are focused on choice of retailer; I instead focus on choice of brand or product within a product market. This means that my measure of household HHI measures concentration of purchases, more akin to Neiman and Vavra (2023). I use the formula

$$household\ HHI_{h,pmc} = \sum_{j=1}^{N} \left(\frac{100 \cdot spending_{j,h,pmc}}{\sum_{j'=1}^{N} spending_{j',h,pmc}} \right)^2$$
$$household\ HHI_{h} = \sum_{pmc} \left(household\ HHI_{h,pmc} \cdot \frac{spending_{h,pmc}}{\sum_{pmc'} spending_{h,pmc'}} \right)$$
$$household\ HHI_{agg} = \sum_{h=1}^{H} household\ HHI_{h} \cdot \frac{w_h}{\sum_h w_h}$$

[10]Product module codes are effectively product markets, such as toilet paper, diet soft drinks, or frozen pizza.

Figure 2: Changes in household spending concentration



Notes: Values are author's own calculations of HHI within a household, using NielsenIQ Consumer Panel data. Within-household HHI is a function of that household's product-market spending habits, which are averaged according to spending in that category and the household's survey weight.

to calculate the household HHI of household h, based on product/brand designations j, in product market with product module code pmc. Product-market household HHI values are averaged according to households' spending in all categories to get a household-level value, then averaged using NielsenIQ's survey weights to yield an aggregate measure of within-household concentration.

Figure 2 shows how this measure has increased over time. Within the average product market, household spending is concentrated and growing more so. From 2004 to 2019, household HHI at the product level increased by 200 points (about 3.5%), and brand-level household HHI increased by 150 (about 1.8%). When combined with general increases in market-level diversity, this increased concentration in household spending implies an increase in siloing. Consumers are allocating a larger proportion of their spending to fewer brands and products, though these are not only the "big" brands. Households must either be trying other brands out (due to their preferred choice being out of stock, sales, or a desire for novelty) at lower rates, or be spending less money on those experimentation efforts.

Consumer sorting is a matter of preference, and this increase in siloing suggests a change in how consumers relate to brands. I therefore turn my attention away from within-household behavior and toward the result of that behavior: the consumer bases of distinct products and brands.

3.3 Consumer homogeneity

While all consumers are unique, their commonalities can lead to similar decision-making. For instance, the intentional gendering of products, such as men's/women's deodorant, leads such products to have relatively more homogeneous consumer bases. I measure a brand or product's level of consumer homogeneity with a metric very similar to HHI, using demographic groups or bins to calculate shares instead of firm identity. I call this metric the consumer homogeneity index (CHI), calculated as follows:

$$popshare_b = \frac{\sum_{h \in \mathcal{H}_b} w_h}{\sum_{h=1}^{H} w_h}$$

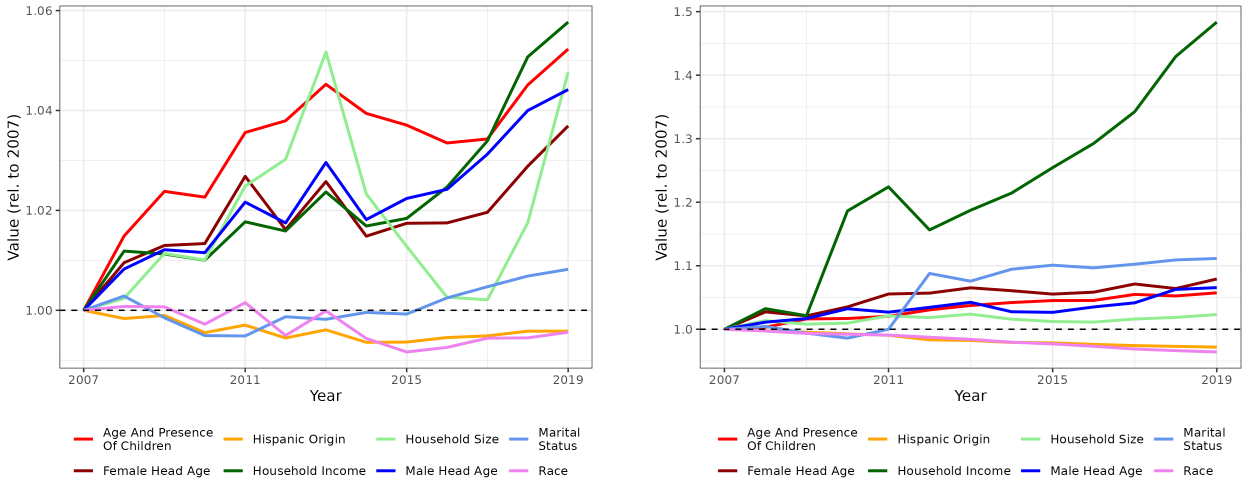
$$consshare_{b,j} = \frac{\sum_{h \in \mathcal{H}_b} spending_{h,j} \cdot w_h}{\sum_{h=1}^{H} spending_{h,j} \cdot w_h}$$

$$CHI_{j,unadj} = \sum_{b=1}^{B} \left(\frac{100 \cdot consshare_{b,j}}{\sum_{b'=1}^{B} consshare_{b',j}} \right)^2$$

$$CHI_{j,adj} = \sum_{b=1}^{B} \left(\frac{100 \cdot consshare_{b,j} / popshare_b}{\sum_{b'=1}^{B} (consshare_{b',j} / popshare_{b'})} \right)^2$$

for household h with NielsenIQ sample weight w_h, demographic bin b, and brand/product j. Household membership in bin b is denoted by that household belonging to set $\mathcal{H}_b$. In the "adjusted" CHI metric, each demographic's population share $popshare_b$ is used to correct for potential changes in demographic composition in the wider population; scaling consumer shares to the general population share prevents differences from being driven purely by demographic shifts. Aggregate measures of demographic sorting are weighted by the sales volume in USD of each good/product.

Figure 3: Consumer sorting, homogeneity



(a) Brand-level CHI

(b) Consumer in product's demographic plurality

Notes: The consumer homogeneity index (CHI) in Panel (a) is calculated identically to the HHI, but employs the demographic bins calculated by NielsenIQ rather than firm/brand identity. For instance, a brand that received 85% of its sales from White customers and 15% from Black customers would have an unadjusted CHI of 7,450. Adjusted values above account for changing population shares of these demographic categories, using that adjusted share as the input to the CHI rather than a raw proportion. Raw values are displayed in Figure A.4. Cross-industry values are weighted across product/brand by total sales of that product/brand. The plurality groups used in Panel (b) are determined annually, using NielsenIQ projection factors to determine the largest consumer group that purchased a given product.

Panel (a) of Figure 3 shows how brand-level CHI values have changed over time, relative to 2007. Consumers who buy the same brand/product are becoming less diverse by a number of metrics: age and presence of children, household income, household head ages (male and female), household size, and marital status.[11] The exceptions are race and ethnicity, which tend to slightly diversify. These adjusted estimates control

[11] See Appendix Table A.3 for details of these demographic bins.

for the population-level share of these bins, reflecting a greater increase in sorting behavior net of broader demographic shifts.

Another way to measure this trend is to examine the same phenomenon from the consumer's perspective. Panel (b) of Figure 3 shows that consumers are increasingly likely to purchase products where a plurality of consumers resemble them—in other words, where they look like the "typical" consumer of that good. As an example, if the largest segment of households who purchase frozen pizzas are single-member households, then a household with four members would not be in that demographic plurality. Apart from larger sorting by household income, most of these increases are on the order of 5-10% of 2007 levels.

To sum up, in an environment where general brand-share concentration is decreasing, within-household spending is actually concentrating. Consumers do not sort randomly, causing the consumer base of the average product/brand to become less diverse. As households become more distinct from one another in general, their purchases increasingly resemble what others of the same age/income/marital status buy.

3.4 Information

During this period, the Internet gives both firms and consumers substantially more information. The ability to target advertising has always been valuable, but the ability to track users' behavior, interests, and known or assumed demographics makes it increasingly effective (Farahat and Bailey, 2012; Yan et al., 2009). The results of this "data mining" are very rich, written about as early as Stein (2011). This approach achieved wider notoriety with the case of Target (correctly) determining that a teenage girl was pregnant, providing pregnancy-related coupons before her family was aware (Hill, 2012).

This means that smartphones, by putting the Internet at consumers' fingertips, would logically heighten this phenomenon. Smartphones let firms observe even more information about consumers, such as location, while also allowing advertisers to show them ads wherever they are. But this is not the only possibility. Contemporary news articles demonstrate a real concern about smartphones in the retail industry, with a Wall Street Journal article going so far as to say that smartphones "strike fear into retailers" (Bustillo and Zimmerman, 2010). Even when discussed in less alarming ways, smartphones were broadly acknowledged as changing the retail industry only a few years after their 2007 debut (Reisinger, 2010; Wood, 2010). These articles largely focus on consumer price shopping, giving consumers the ability to see whether the product they want is cheaper at a competing retailer.

Both the targeted marketing and price shopping narratives suggest that Internet access via smartphones should contribute to the broad increases in consumer sorting we observe. Firms can market to the most susceptible consumers, likely narrowing their consumer base but encouraging greater loyalty with ads or coupons. Consumers, after choosing the product they want, can find that product for a lower price more easily. I discuss the background of this technological change in greater detail in the next section.

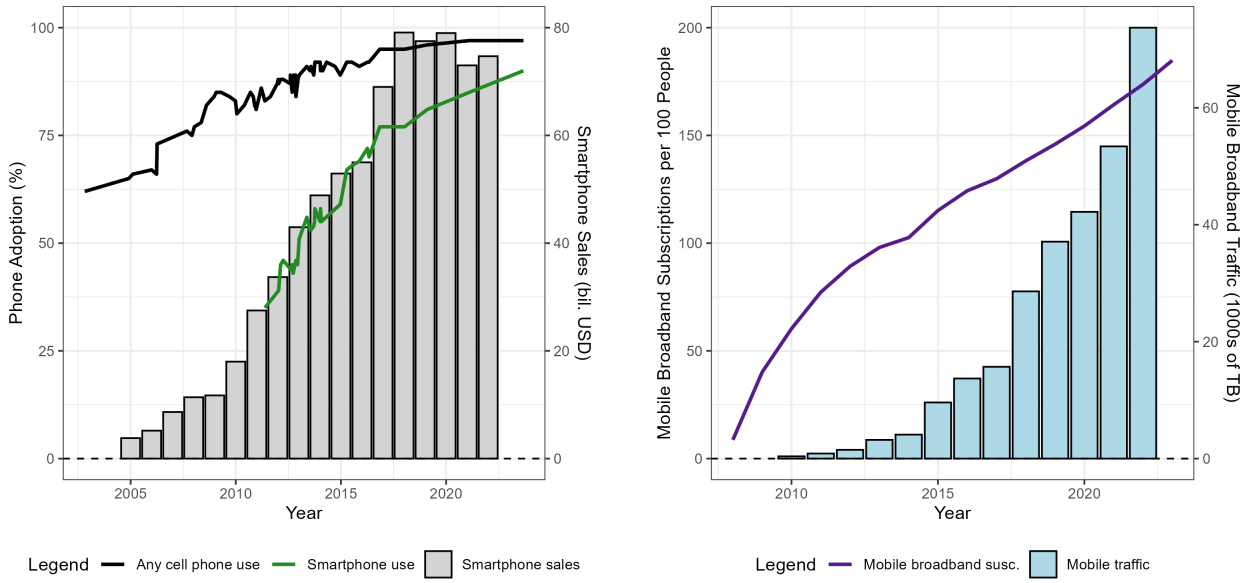
4 Mobile broadband and smartphones

In 2000, the majority of U.S. cell service consisted of 2G, which is sufficient for encrypted voice and text communication; the basic form of 2G has a maximum speed of 50 kilobits per second (Kbps).[12] This speed is too slow for consumers to comfortably go online, as a page of Google results is about 250KB, requiring approximately forty seconds to load at top speed.[13] 3G, a standard finalized in 1998 but not implemented in

[12]Source: `https://www.lifewire.com/1g-vs-2g-vs-2-5g-vs-3g-vs-4g-578681`

[13]Examined in 2025. Note that a kilobyte consists of eight kilobits, so this page of Google results is 2000 Kb large.

Figure 4: Smartphone and mobile broadband adoption



(a) Smartphone adoption

(b) Mobile broadband adoption

Notes: Panel (a) uses data on cell phone use from Pew Research (2024), and data on smartphone sales from Statista and the Consumer Technology Association. Panel (b) uses data on mobile broadband from the FCC (accessed via ITU) and mobile traffic from the CTIA. All data accessed in March 2025.

the United States until the early to mid 2000s, was significantly faster, with a maximum speed of 2 megabits per second (Mbps), forty times faster than the previous standard.[14]

Thus, prior to the rollout of 3G, consumers were effectively unable to access the Internet via their mobile devices, as even the fastest connection speeds were simply too slow for a usable browsing experience. While still slower than a contemporaneous wired connection, 3G was substantially faster than previous mobile standards. As it became more common, manufacturers began to release devices capable of taking advantage of 3G's faster speeds, though these devices' interfaces did not necessarily keep up. Early 3G-capable phones were generally traditional "dumb phones" despite their upgraded connection capabilities. While devices like BlackBerries had the capacity to go online and send emails, these were *not* the typical 3G-capable devices—and in fact, the company's 2009 Curve 8900 model was only 2G-compatible.[15] The first modern, touchscreen smartphone[16]—the Apple iPhone—debuted in 2007, with the Q3 2008 release of the iPhone 3G bringing to market the first 3G-compatible phone with an interface truly suitable for mobile Internet use. In the absence of more granular adoption data, these factors together suggest that 2009 is the earliest reasonable year in which smartphone access could be a relevant factor for consumers.

By 2008, many American consumers already had some kind of cell phone, and smartphone adoption was rapid—by 2013, a majority of American adults reported having a smartphone,[17] just six years after the invention of the category. Figure 4 depicts this near-exponential increase in smartphone adoption and mobile broadband use across time. Panel (a) shows how sales and adoption closely followed each other, with

[14] ibid.

[15] Source: `https://www.slashgear.com/1092246/the-most-iconic-blackberry-phones-of-all-time/`

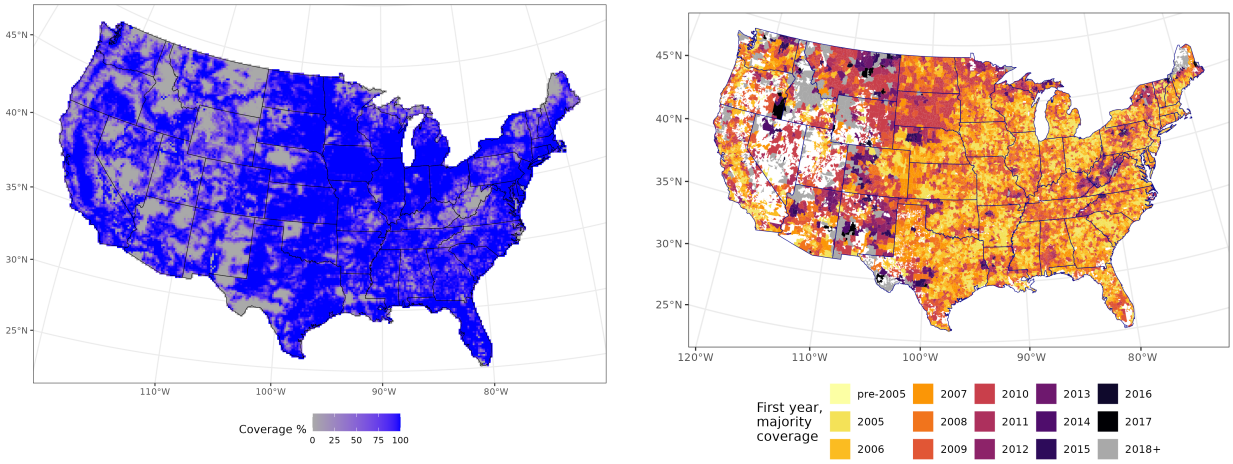
[16] Devices such as those produced by BlackBerry predated the iPhone, and for market purposes, are often considered early smartphones. This is why Figure 4 has positive sales for the smartphone industry in 2005. As such devices do not fit the modern conception of a smartphone, and would be cumbersome to browse the Internet with, I do not consider them true "smartphones."

[17] Source: `https://www.pewresearch.org/Internet/fact-sheet/mobile/`

the smartphone industry worth $18 billion by 2010 and well over $25 billion the following year. Not only were consumers buying these devices, they were using them on the go. Panel (b) demonstrates the rapid increase in both sales and utilization of mobile broadband subscriptions, with subscription rates exceeding 100 per 100 people by 2014.[18] The amount of mobile data used increased exponentially over the same period, reflecting significantly more intensive use.

4.1 Mobile broadband

Figure 5: Observed and predicted 3G mobile broadband coverage



(a) Observed mobile coverage, December 2010

(b) Predicted year of first coverage

Notes: Panel (a) depicts reported mobile broadband coverage as of December 2010, using State Broadband Initiative (SBI) reports of 3G coverage. Panel (b) shows the estimated year of majority coverage at the ZCTA level, based on state-level random forest predictive models. The year indicated is the first year for which coverage meets or exceeds a threshold of 50% (majority coverage).

The highest-quality current data source for both mobile and fixed-connection broadband coverage in the United States is the Form 477 data collected by the FCC. Since December 2014, mobile broadband providers have been required to submit company coverage records through Form 477, which is then compiled into publicly-released information by the agency. However, older data is available from another source: from 2010 to 2014, the National Telecommunications and Information Administration (NTIA) collected similar data under the State Broadband Initiative. These data are not identical—the NTIA data were voluntarily reported by broadband providers and validated by mobile speed tests (Myers and Hooper, 2026); though some data artifacts exist, most providers did contribute to this initiative.[19] Form 477 data is not independently validated, but is also not voluntary, as providers are required to report their coverage areas.

The mid-to-late 2000s was also a period of significant cell-tower construction, and 3G predates the adoption of smartphones by several years. To accurately determine when a location first gained access to mobile broadband in the form of 3G (or better), I must use supplementary information.

The FCC's Antenna Structure Registration (ASR) system requires entities that own an antenna structure (tower) of over 200 feet in height to file information on that tower's location, height, date of construction,

[18]This reflects business subscriptions, as consumers are unlikely to need more than one mobile broadband connection per device.

[19]Relative to the December 2010 snapshot, the June and December 2011 records seem to indicate *lower* levels of 3G coverage in the majority of the country. This seems unlikely at best.

and any modifications with the FCC. This allows me to use the construction of new towers, combined with information about local topography, to predict 3G coverage in nearby areas. All together, I use NTIA and FCC coverage snapshots, periods of zero mobile broadband coverage in 2000-2004, and ASR information on tower construction to train random forest models and predict 3G coverage levels from 2005 onward.[20] I emulate Hersh, B. J. Lang, and M. Lang (2022) in the variables I include in each model, largely focusing on tower-level information (height, the number of towers within 20 miles of the location) and local characteristics that may interfere with cell coverage (standard deviation of local elevation).[21] Tower-level variables are based on either the newest tower in the grid cell, or the closest tower to the cell's boundaries. Due to differing geographic conditions across different regions of the country, I train 49 separate state-level random forests, including one for the District of Columbia. This allows, for instance, locations in West Virginia to be more accurately predicted as receiving coverage later, despite the Appalachian mountains tending to be not quite as steep as the Rockies in the West.

The random forests use the percent (0-100) of a census block covered by some kind of mobile broadband as the outcome of interest, though I aggregate this by land area up to a 100-square mile grid cell level. This means that the predicted values the models produce (likewise 0-100) can be thought of either as proportions of coverage or as probabilities of coverage. In one sense, an area with a coverage rate of 80 is 80% covered. But likewise, a random location within this area has an 80% chance of being covered. With this in mind, there are different possible thresholds for an area to be considered "treated"; the most intuitive (and my preferred specification) is 50%.

4.2 Random forest validation

The random forests generally behave intuitively, and match tower construction to expansions in coverage. Table 2 includes detailed model performance, comparing different metrics to a potential alternative: a fractional logit predictive model. Because 20% of the sample (containing records from December 2000-2004, 2010, 2012, and 2015) was withheld to prevent overfitting, there are two logical data sets to assess out-of-sample performance. The "Input" data reflects the 80% of the training data used to fit both the random forest and logit models, while the "Holdout" is the remaining 20% left out of the sample. Additionally, the NTIA data includes snapshots from other time periods, of which the highest-quality unused samples occur in June 2012 and June 2013.[22]

In Panel (a) of Table 2, in- and out-of-sample performance for the input and holdout samples are fairly similar, where the root mean squared error is about 13.5 (on a 0-100 scale) and accuracy for a 50% cutoff[23] hovers around 96.5%. The model performs worse on the fully out-of-sample data, though still maintains extremely high levels of accuracy—around 90%—and RMSE values around 18. While a fractional logit

[20] Snapshots are: December 2010 and December 2012 from NTIA, the former depicted in Panel (a) of Figure 5; December 2015 from FCC. "Empty" maps from 2000-2004 are included because 3G was rolled out in the United States starting in around 2005, providing some accurate pre-3G data on which to train models.

[21] I source information on local topography from Amazon Web Services, accessed via the `elevatr` package in `R`.

[22] While the NTIA data has several desirable qualities, it also relies on voluntary reporting to state-level agencies. Because the data also begins near/during a decennial Census cycle, documentation indicates that snapshots from 2010 and 2011 use Census 2000 blocks, while those in 2012 and onward use 2010 blocks. In practice, operating under this assumption yields poor match quality for many snapshots, and many records after 2012 continue to be recorded in Census 2000 blocks. Relative quality of a snapshot was judged by how plausible it looked, as these breakdowns generally result in implausibly small and inconsistent levels of coverage. One example of an untrustworthy sample is June 2011, which sees generally lower coverage than the December 2010 data in Panel (a) of Figure 5. I allow for matches based on the "correct" data snapshot year first (2000 for 2010/2011; 2010 for 2012 onward), only allowing for matches in the other decennial Census year for geographies without immediate matches.

[23] This cutoff determines at what point binary "coverage" activates—under a 50% cutoff, a cell/ZCTA that is 55% covered by 3G would be considered "covered," while it would not be under a 60% cutoff.

Table 2: Model quality metrics

(a) Random forest(s)

			Residual percentile			Model performance			Accuracy at X% cutoff		
	Count	Sample pct.	15%	50%	85%	RMSE	MAE	R^2	40%	50%	60%
Input	205957	80%	-7.009	-0.445	5.249	13.443	7.157	0.908	0.965	0.965	0.960
Holdout	51491	20%	-6.969	-0.429	5.344	13.537	7.222	0.908	0.964	0.964	0.960
Extra sample - 2012	32181	–	-8.488	0.922	18.573	18.514	12.025	0.611	0.903	0.896	0.881
Extra sample - 2013	32181	–	-7.453	0.911	17.078	17.575	11.105	0.608	0.912	0.904	0.893

(b) Fractional logit model

			Residual percentile			Model performance			Accuracy at X% cutoff		
	Count	Sample pct.	15%	50%	85%	RMSE	MAE	R^2	40%	50%	60%
Input	205957	80%	-26.079	-5.558	30.277	31.208	22.614	0.481	0.844	0.853	0.847
Holdout	51491	20%	-26.012	-5.379	31.492	31.594	22.916	0.476	0.840	0.849	0.843
Extra sample - 2012	32181	–	-4.263	18.160	63.948	41.281	31.771	0.103	0.665	0.627	0.590
Extra sample - 2013	32181	–	-3.745	16.828	63.505	40.682	30.799	0.101	0.679	0.642	0.606

Notes: The unit of observation is the grid cell, a 100 mi.^{2} (10 mi. by 10 mi.) square. The fractional logit regression in Panel (b) includes state-level interaction terms for all variables, though predictions do not differ substantially when removing these interactions. See Table A.4 for the coefficients from the logit regression without state-level interactions.

model is a natural alternative given the distributional bounds of coverage, this model does not perform as well as the random forests. Panel (b) shows that the RMSE for in-sample data jumps significantly to over 30, yielding a value much higher than either out-of-sample assessment of the random forest approach. Accuracy also decreases, though not as much, to around 85%. However, out-of-sample predictions—the entire point of using this predictive model—are much worse in the logit setting. Out-of-sample accuracy hovers around 63%, with a higher RMSE of around 41 and a very poor R^2 of 0.10. The gains of using machine learning in this setting more than make up for the reduction in transparency that results from using random forests instead of simple logit predictions.

Using the random forest models, Panel (b) of Figure 5 shows the resulting predicted timing of 3G rollout across the United States. For the purposes of this project, I aggregate up to the zip code level, using whether a consumer lives in an area with 3G coverage to determine whether they are likely to have coverage while shopping. This abstracts away from in-store access, as NielsenIQ only provides 3-digit zip locations for stores. I assume that consumers who live in areas without 3G access are unlikely to get smartphones, making the comparison a bit more apt. For areas predicted to have 3G coverage prior to 2009, because I do not think that this technology should influence consumer behavior in the absence of smartphones, I delay this "access" until 2009, making my measure an intent-to-treat rather than a direct observation of access. Table A.5 shows the cumulative proportion of ZCTAs that are considered to have 3G access during the study period, which steadily increases from roughly a quarter of zip codes in 2005 to over 95% coverage by 2012.

5 Reduced form results

In this section, I evaluate whether smartphones contributed to the broad changes in siloing attested in Section 3.

5.1 Methodology

For reduced-form estimates, I use Callaway and Sant'Anna (2021)'s robust difference in differences estimator. All estimations are performed at the city-quarter level, with the following equivalent two-way fixed-effects formula:

$$y_{ct} = \beta \cdot Smartphone_{ct} + \kappa_c + \tau_t + \varepsilon_{ct} \tag{1}$$

where y_{ct} is a particular outcome and $Smartphone_{ct}$ is a binary variable indicating majority smartphone coverage for city c at year-quarter t.

Estimating shifts in consumer behavior is somewhat complicated by the spatial balance of the NielsenIQ Consumer Panel. The panel strives to be nationally representative, and most Americans live in cities—however, much of the variation in mobile Internet rollout occurs in rural areas. Exposure to mobile broadband/smartphones is determined at the zip code level, but many zip codes have only a few households recorded in the panel. 18.6% of zip codes that appear in the data average less than one household from 2004-2019, and a majority (51.2%) average less than four. A substantial portion of 3G/smartphone rollout variation occurs in these areas, where zip-code fixed effects would absorb a large portion of any variation. This makes household-based regressions much noisier than city-level ones. With this in mind, and for consistency with the city-level demand estimates in Section 6, I estimate reduced-form outcomes at the city level.

5.2 Siloing results

I start by determining whether 3G and smartphone access contributes to the three broad trends outlined in Section 3. Table 3 shows these reduced-form estimates, using the Callaway and Sant'Anna (2021) estimator.

The first aggregate trend I described was a decline in top brands' market share. Smartphone access contributes to this decrease. The proportion that consumers purchase from top brands falls after smartphones become accessible: specifically, the top three and top ten brands see declines of 1.2 and 1.6 percentage points, respectively. This reflects the same pattern of smaller brands faring better in Figure 1, pointing to greater aggregate diversity.

The second trend was greater consumer spending concentration, described using the average household HHI. While Figure 2 shows that household spending concentration is generally trending up, the coefficient associated with smartphone adoption in Table 3 is negative, albeit statistically insignificant. I cannot rule out a small increase in household HHI, but I can rule out a change any larger than approximately 29 points, or 0.36% of the sample mean. This suggests that, counter to what theory would suggest, smartphone access is not significantly contributing to increases in household spending concentration.

The third trend was greater homogenization of brands' consumer bases. To determine whether smartphone access contributes to this trend, I examine consumers' propensity to buy products where they resemble the modal/typical consumer. In age/presence of children and household size, I find significant negative effects, reflecting drops of 9.4 and 3.9 percentage points (respectively) of households being in the largest demographic of a product's consumer base. These shifts are proportionally large, with movements representing 10-13% of the sample mean. Though these are the only significant deviations, six of the eight metrics I examine are negative, indicating that households are buying more "minority" products in general. The coefficient on household income bins is positive and insignificant, while that for female head age is functionally zero.

Taken together, Table 3 indicates that smartphones are *not* driving increases in consumer sorting/siloing.

Table 3: Siloing evidence, brand level

	Estimate	Sample mean
Prop. from top 3 brands (in PMC)	-0.012*	0.546
	(0.006)	
Prop. from top 5 brands (in PMC)	-0.018	0.654
	(0.011)	
Prop. from top 10 brands (in PMC)	-0.016*	0.774
	(0.008)	
Household HHI (brand)	-62.534	8,016
	(46.599)	
In plurality, age/presence of children	-0.094*	0.689
	(0.042)	
In plurality, Hispanic origin	-0.026	0.915
	(0.019)	
In plurality, household size	-0.039*	0.399
	(0.011)	
In plurality, marital status	-0.002	0.551
	(0.039)	
In plurality, female head age	0.000	0.298
	(0.011)	
In plurality, household income	0.045	0.251
	(0.049)	
In plurality, male head age	-0.032	0.355
	(0.037)	
In plurality, race	-0.029	0.831
	(0.022)	

Notes: The estimates presented here are the result of Callaway and Sant'Anna (2021)-style estimations, using the city and year-quarter as area and time fixed effects, respectively. For concision, this table includes only brand-level outcomes; product-level outcomes are similar. Note that "treatment" is determined at the city level, and all cities are considered plurality "treated" by Q2 of 2011. This limits the time horizon to estimate post-treatment effects. Sample means for logged dependent variables are reported in unlogged levels.

They do not significantly increase household spending concentration, and have the opposite effect on demographic sorting than theory would indicate.

5.3 Other possible causes

The natural next question is whether some other factor might be driving these results. It is theoretically possible that another factor could influence smartphone adoption while also offsetting trends in consumer behavior. For instance, if growing cities were more likely to receive mobile broadband coverage, then ascribing any effects to smartphones/broadband would be a mistake. The primary demographic characteristics I examine in this section are population and median income. I also explore how the number of fixed-connection broadband providers (provided by the FCC through 2012) is related, to provide evidence of what factors matter for broadband expansion.

The estimations in this section still use the Callaway and Sant'Anna (2021) approach to avoid negative weights, but should be interpreted as suggestive results only. These results are not causal—they are intended to indicate a correspondence, not that broadband (does not) change an area's population, income, or broadband provider landscape.

Table 4: Difference-in-differences estimates on potential confounders

	Estimate	Sample mean
Log population	-0.001	1,470,685
	(0.004)	
Log median income	-0.012*	$68,599
	(0.004)	
Log count of broadband providers	0.057	16.733
	(0.063)	
Prop. 3+ person households	0.001	0.385
	(0.001)	

Notes: The estimates presented here are the result of Callaway and Sant'Anna (2021)-style estimations, using the city and year-quarter as area and time fixed effects, respectively. Standard errors are clustered at the city level; asterisks denote statistical significance at or below the 5% level. Population and median income are measured using Census 2000 and ACS data at the ZCTA level from 2011-2019, aggregated up to the city; linear interpolation is used to estimate values from 2001-2010. Provider count data (likewise originating at the zip code level) comes from the FCC, and spans 2000-2012. Note that "treatment" is determined at the city level, and all cities are considered majority "treated" by Q2 of 2011. This limits the time horizon to estimate post-treatment effects. Dependent variables are derived at the city level (aggregating population, averaging otherwise); sample means for logged dependent variables are reported in unlogged levels.

In Table 4, I find no statistically significant relationship between smartphone adoption and local population, broadband provider counts, or households with three or more members (a proxy for families). I do find a negative relationship between smartphones and *lower* city-level incomes, with a magnitude of about 1.2%. While this association is statistically significant, it is practically small. The implication that broadband providers are not targeting high-income areas exclusively is encouraging, though the suggestion that they are targeting lower-income locations seems unlikely. This association may instead reflect regional growth patterns, as areas of the South and Mountain West experienced a great deal of telecommunications investment and have lower income levels than the Northeast and Pacific West. Real incomes during this period were also stagnant or shrinking for large swathes of the country.[24] Regardless of the reasoning behind this association, I correct for any distortions this might cause by controlling for household income in subsequent estimations.[25] I therefore find no evidence to suggest an alternative cause for my findings beyond changes in mobile-phone technology.

Having discussed potential confounders, it bears repeating the initial question: how are smartphones offsetting trends in consumer sorting? There are many aspects of the shopping experience that can change to accomplish this, which I group into three broader categories to walk through the reduced-form implications. First, the price that consumers pay can vary; this includes both the retail price and the final price, after applying any sales or coupons. If consumers have underlying strong brand preferences, they may be more willing to price-shop among retailers after becoming more informed, resulting in greater observed brand sorting. Second, the in-store "market" characteristics can change. This includes the consumer's choice of retailer,[26] but also the options available in the store. Conceptually, increases in product offerings as in Neiman and Vavra (2023) would fall into this category. Third, consumers' decisions can change. This is primarily what is purchased (or not purchased), but can also include *who* purchases particular products, which may be more commonly chosen among a particular demographic. Cereal is more commonly purchased by households with children, for instance.

[24]Source: `https://www.census.gov/library/visualizations/2015/acs/2014-2009-acs-hh-income.html`

[25]Controlling for household income does not significantly alter the coefficients in Table 3; see Table A.2.

[26]The choice of when to shop would also factor in here, as markets are divided by time as well as space, though I do not generally observe this decision.

I begin by testing to what extent prices change as a response to smartphone access.

5.4 Prices

Table 5: Price at purchase

	Estimate	Sample mean
Prop. of entries with a "deal"	0.018	0.205
	(0.010)	
Prop. of entries with coupon value	-0.001	0.062
	(0.003)	
Prop. of spending saved, coupons	0.000	0.025
	(0.001)	
Log total amount spent (USD)	0.028	$929.68
	(0.036)	
Log total amount saved, coupons (USD)	-0.428	$24.91
	(0.555)	
Log net amount spent (USD)	0.029	$904.77
	(0.034)	

Notes: The estimates presented here are the result of Callaway and Sant'Anna (2021)-style estimations, using the city and year-quarter as area and time fixed effects, respectively. Standard errors are clustered at the city level; asterisks denote statistical significance at or below the 5% level. Note that treatment is determined at the city level, and all cities are considered majority treated by Q2 of 2011. Sample means for logged dependent variables are reported in unlogged levels.

If consumers are comparing prices to get the best deal on a particular product, we might expect that to appear in the data. Rather than comparing the final price consumers pay for particular goods, Table 5 shows how usage of coupons and sales changes—or, more aptly, does not change. I find no evidence that consumers are more likely to shop "deals"[27] or use coupons after gaining access to smartphones. The coefficients on coupon use are slightly negative, though not close to statistical significance. The total amount spent or saved is likewise unaffected; consumers are not buying a great deal more or less at the grocery store.

Another way to explore whether prices change involves dealing with individual product markets. Consumer purchases do not reflect all options available to them at the store. If posted prices became further apart, for instance, it would be difficult to determine this based only on what consumers paid. It is infeasible to download all product markets from NielsenIQ directly due to space constraints, so I analyze a subset of 35 markets in Figure 6.[28] Only five of those markets see a statistically significant difference in the standard deviation of posted prices (four of which see an increase), suggesting that neither firms nor consumers are necessarily changing their reaction to prices.

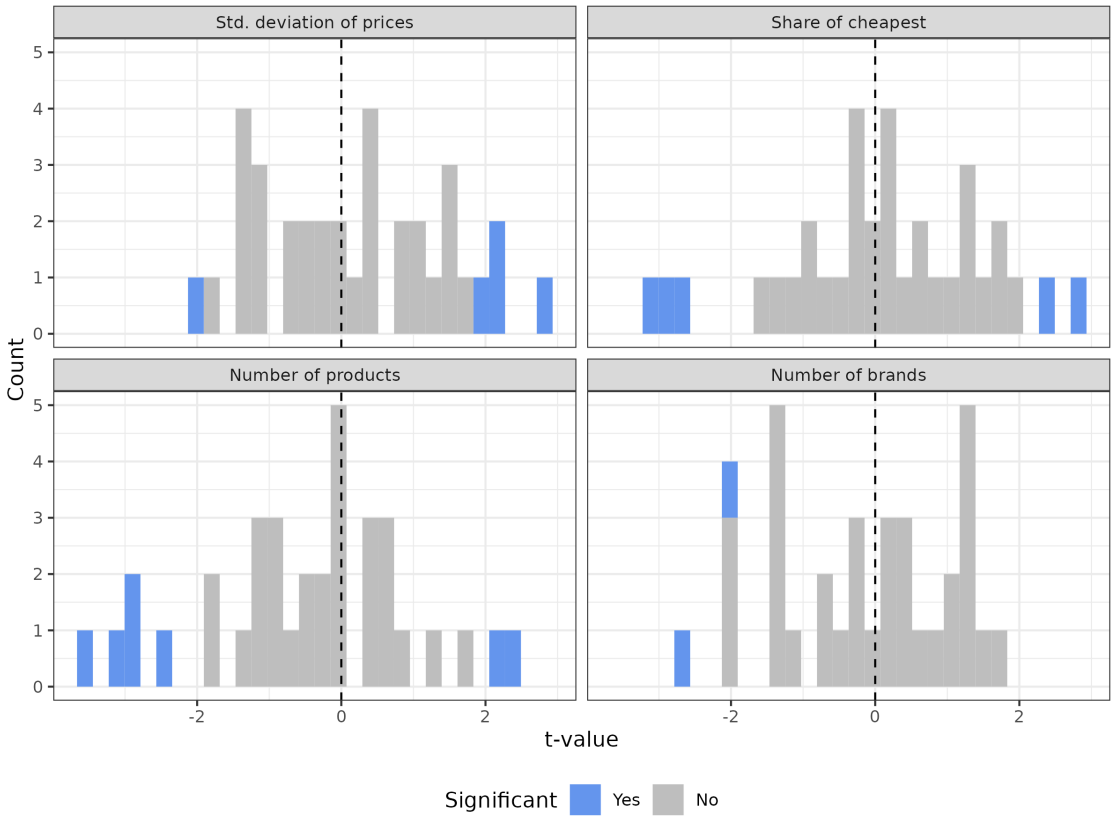
The NielsenIQ Consumer Panel also does not show whether consumers purchased what was cheapest, but the Retail Scanner data do. The proportion of consumer purchases from the cheapest option is also inconsistently affected by smartphone entry, with only five of 35[29] product markets seeing significance; only two of those five represent increases. Together, Table 5 and Figure 6 indicate that smartphone access does not significantly change the prices consumers face or the ways they "shop around."

[27] This is a term used by NielsenIQ, denoting internal sales. This is distinct from coupons, which have a specific monetary value and are recorded as such.

[28] See Table A.1 for a list of these markets.

[29] A total of 36 product markets were examined, but numerical issues leave only 34-35 depending on the outcome.

Figure 6: Product market estimates



Notes: The estimates presented here are the result of many Callaway and Sant'Anna (2021)-style estimations, using the city and year-quarter as area and time fixed effects, respectively. The exclusion of one product market for the share of cheapest good is motivated by numerical issues. Note that treatment is determined at the city level, and all cities are considered majority treated by Q2 of 2011. Due to the number of estimations run, this figure shows the distribution of t-values.

5.5 Market characteristics

Smartphones could potentially change where consumers shop, and/or the options available to them at a given retailer. I find that consumers make more of their trips to their "favorite" (modal) retailer after getting access to smartphones, as Table 6 indicates. This 2.2 percentage point increase is almost a 6% increase over the sample mean. It is difficult to measure market characteristics only from a consumer's perspective, though. Moving back to Figure 6, among 35 product markets, only seven saw a statistically significant change in the number of products on offer (and four in brands), though the majority of these changes were *reductions* in options for consumers. I conclude that there is no consistent evidence that consumers' in-store options are changing, much less increasing, after smartphones become available.

5.6 Purchase decisions

The primary driver of the results in Table 3 is changes in what consumers buy. The first major finding from Table 7 is that, while consumers are buying less from the top brands, they may be purchasing weakly more distinct products and brands. Shifts seem to be driven primarily by purchases in unfamiliar product markets: after gaining smartphone access, the number of product module codes (PMCs) purchased increases by 5.7%, while the number of brands and products per PMC does not significantly change.

Table 6: Market characteristics

	Estimate	Sample mean
Log distinct retailers visited	-0.011	11.291
	(0.033)	
Prop. of visits, modal retailer	0.022*	0.370
	(0.006)	

Notes: The estimates presented here are the result of Callaway and Sant'Anna (2021)-style estimations, using the city and year-quarter as area and time fixed effects, respectively. Standard errors are clustered at the city level; asterisks denote statistical significance at or below the 5% level. Note that treatment is determined at the city level, and all cities are considered majority treated by Q2 of 2011. Sample means for logged dependent variables are reported in unlogged levels.

Table 7: Purchase decisions

	Estimate	Sample mean
Log distinct UPCs bought	0.059	174.567
	(0.031)	
Log distinct brands bought	0.051	104.549
	(0.033)	
Log number of PMCs bought from	0.057*	101.236
	(0.022)	
Log avg. brands per PMC	-0.004	1.011
	(0.023)	
Log avg. UPCs per PMC	-0.001	1.653
	(0.014)	

Notes: The estimates presented here are the result of Callaway and Sant'Anna (2021)-style estimations, using the city and year-quarter as area and time fixed effects, respectively. Standard errors are clustered at the city level; asterisks denote statistical significance at or below the 5% level. Note that treatment is determined at the city level, and all cities are considered majority treated by Q2 of 2011. Sample means for logged dependent variables are reported in unlogged levels.

This further characterizes changes in consumer sorting. Because households are more frequently buying products where most consumers do not look like them, they may be discovering these products thanks to smartphones, or through experimentation induced by smartphones. This sorting is unlikely to be driven by firms (via targeted coupon provision or advertising), as the simplest way for firms to target consumers would be to focus on repeat consumption.

Together with there being no significant changes in prices (paid or posted), this suggests that common narratives of how smartphones affect shopping are not likely to be true. However, these reduced-form estimates are the result of possible shifts in both supply and demand. To better disentangle these separate shifts, I next estimate demand in several product markets.

6 Demand estimation

6.1 Methodology

I estimate demand as in S. Berry, Levinsohn, and Pakes (1995), with the addition of micro moments as in Petrin (2002) and many others, using the `pyBLP` package.

In explaining the model here, I generally follow the notation in Conlon and Gortmaker (2025). The demand model in this setting is identified by two sets of moments. The first are the standard BLP moment conditions, interacting structural errors with instruments for price. These help account for endogeneity in

how prices and shares are determined. Second is a set of micro moments that describe consumer behavior. These moments vary with respect to product characteristics and demographics among those who purchased the inside good, and help pin down variation that market-level data alone obscure.

I express the indirect utility of individual consumer i in market $t = 1, ..., T$ consuming a product $j = 1, ..., J_t$ in the following way:

$$U_{ijt} = \delta_{jt} + \mu_{ijt} + \epsilon_{ijt} \tag{2}$$

where

$$\begin{aligned} \delta_{jt} &= \boldsymbol{x}'_{jt}\beta + \xi_{jt} \\ \mu_{ijt} &= \alpha_{it} \cdot p_{jt} + \boldsymbol{x}'_{jt}(\Pi y_{it} + \Sigma\nu_{it}) \end{aligned}$$

$$\begin{aligned} y_{it} &= \zeta_{it} + Smartphone_{it} + \zeta_{it} \cdot Smartphone_{it} \\ \alpha_{it} &= -exp(\alpha_0 + \Pi y_{it} + \Sigma\nu_{it}) \end{aligned}$$

Here, δ_{jt} is the mean utility for good j in market t, μ_{ijt} is the heterogeneous utility, and ϵ_{ijt} is a consumer's idiosyncratic preference. δ_{jt} is dependent on product characteristics, where an additivity assumption yields $\delta_{jt} = \boldsymbol{x}'_{jt}\beta + \xi_{jt}$ for product characteristics x_{jt} and demand-side structural error ξ_{jt}. μ_{ijt} is more consumer-dependent, so for demographics and preferences $(y_{it},\ \nu_{it})$, I describe it as $\mu_{ijt} = \boldsymbol{x}'_{jt}(\Pi y_{it} + \Sigma\nu_{it})$.

In my setting, the vector of demographics y_{it} includes dummy variables for income, household head age, and smartphone access; I indicate the terms for income and age as ζ_{it}. To allow the effects of income and age to vary with smartphone access, as shown above, I include interaction terms. I provide more details on the derivation of these variables in Section 6.1.2.

I also use a lognormal price coefficient as described in Appendix A of Conlon and Gortmaker (2025), such that α_{it} is the random price component, containing α_0 as a baseline constant. This prevents consumers with very positive unobserved preferences from having an upward-sloping demand curve, as α_{it} is always strictly negative.

Based on these utilities, consumers choose which good j they want to consume (choosing only one at any given time), with their choice d_{ijt} being 1 if a given good yields the highest utility among all $J_t = \{0, 1, ..., J_t\}$ and 0 otherwise. Consumers' outside option, $j = 0$, provides a normalized value of zero.

After all consumers make their respective choices, I aggregate those choices to produce market shares, s_{jt}. Let $\boldsymbol{\delta}_t$, $\boldsymbol{\mu}_{it}$, and $\boldsymbol{\epsilon}_{it}$ be the $J_t \times 1$ vectors with elements δ_{jt}, μ_{ijt} and ϵ_{it}, respectively. Then:

$$s_{jt} = \int d_{ijt}(\boldsymbol{\delta}_t,\ \boldsymbol{\mu}_{it})\mathrm{d}F(\boldsymbol{\mu}_{it}, \boldsymbol{\epsilon}_{it}) \tag{3}$$

If the ϵ_{ijt} follow a type I extreme value distribution (by assumption), then the logit probability that consumer i (or of type i, sometimes) chooses product j is

$$s_{ijt} = \frac{\exp(\delta_{jt} + \mu_{ijt})}{1 + \Sigma_{k \in J_t}\exp(\delta_{kt} + \mu_{ikt})} \tag{4}$$

The goal of demand estimation is to use the known aggregate shares to back out parameter values for the underlying equations. In practice, estimation inverts the demand system and recovers δ_{jt} from the shares

$s_{jt}(\Pi,\ \Sigma,\ \delta_t) = s_{jt}(\theta)$ by matching observed shares $\mathcal{S}_{jt}$. Under a couple of assumptions,[30] there is at most one vector δ such that the σ_j mapping from δ_t and p_t to s_{jt} is satisfied for all j. This results in the following correspondence:

$$s_{jt} = \sigma_j(\delta_t,\ p_t), \quad \forall j$$
$$\delta_{jt} = \sigma_j^{-1}(s_t,\ p_t),\ \forall j$$

The general logic behind the GMM estimator for θ that I construct and use to structurally estimate demand hinges on M_A moment conditions $\mathbb{E}[\xi_{jt} \cdot z_{jt}] = 0$, with price instruments z_{jt}. Note that ξ_{jt} has mean zero and is mean-independent of the M_A instruments z_{jt}. Using a weighting matrix W_A, the optimization operates on

$$\hat{\theta}_A = \arg\min_{\theta}\ \hat{g}_A(\theta)' \cdot \hat{W}_A \cdot \hat{g}_A(\theta) \tag{5}$$
$$\hat{g}_A(\theta) = \frac{1}{N_A} \sum_t \sum_j \underbrace{(\hat{\delta}_{jt}(\Pi,\ \Sigma) - x'_{jt}\beta)}_{\hat{\xi}_{jt}(\theta)} \cdot\ z_{jt}$$

The specific instruments I use here include a Hausman (1996)-like price measure, leveraging differences in pricing for a given product in different local markets, as well as product-characteristic differentiation and BLP IVs using the brand dummies.

While Equation 5 facilitates demand estimation, I extend the setup in (5) with M_M micro moments to the following:

$$\hat{\theta} = \arg\min_{\theta}\ \hat{g}(\theta)' \cdot \hat{W} \cdot \hat{g}(\theta) \tag{6}$$
$$\hat{W} = \text{diag}(\hat{W}_A,\ \hat{W}_M), \qquad \hat{g}(\theta) = \begin{bmatrix} \hat{g}_A(\theta) \\ \hat{g}_M(\theta) \end{bmatrix}, \qquad \hat{g}_M(\theta) = \begin{bmatrix} f_1(\bar{v}) - f_1(v(\theta)) \\ \vdots \\ f_{M_M}(\bar{v}) - f_{M_M}(v(\theta)) \end{bmatrix}$$

which helps better fit the observed shares, and requires fewer price instruments. Borrowing notation from Conlon and Gortmaker (2025), I include two micro moments for each demographic characteristic in $y_{\cdot it}$ and exogenous product characteristic in $x_{\cdot jt}$. These moments are calculated across consumers in a region using the NielsenIQ projection factor weights and take the following form, as all values are binary:

$$\mathbb{E}[x_{\cdot jt} \mid y_{\cdot it} = 1]$$
$$\mathbb{E}[y_{\cdot it} \mid x_{\cdot jt} = 1]$$

Combining all of these observable characteristics yields a total of 32 micro moments, the product of four

[30] Using S. T. Berry and Haile (2014), the assumptions are (1) that ξ_{jt} and x_{jt} only affect the distribution of utilities through δ_t, and (2) that the J goods are substitutes in $-p_t$ and δ_t.

$y_{\cdot it}$ and four exogenous $x_{\cdot jt}$ variables.

6.1.1 Inside good share derivation

Following Backus, Conlon, and Sinkinson (2021), I define market size as the number of trips taken to a retailer in a city-quarter. This is attractive for a number of reasons, the primary one of which is that it abstracts away from potential discrepancies between local population and store popularity. While the Consumer Panel data is rich, the number of trips (which would include consumers entering a store and purchasing nothing) are not directly observed. To approximate a measure of trips, I estimate the following equation:

$$log(q_{kt}) = \beta_1 log(q_{kt}^{milk}) + \beta_2 log(q_{kt}^{eggs}) + \gamma_k + \varepsilon_{kt} \tag{7}$$

where q_{kt} is the quantity sold in a product category at chain k and market t. Milk and eggs are among the most frequently purchased items. Milk in particular is the most-purchased product category, and is a part of 23% of all trips to grocery stores and mass-market retailers in the NielsenIQ Consumer Panel. Eggs are the fourth most purchased product type, purchased in 11.5% of all trips to these stores.[31]

I then use Equation 7 to predict the amount of each good sold in a given market, scaled to match the average inside good share. The resulting inside-good purchasing shares are shown in Figure A.1, with the dashed line indicating the average value across all markets.

6.1.2 Data and estimation

The main data for demand estimation come from the NielsenIQ Retail Scanner. Products are defined at the UPC-version level (the "unique products" column of Table 1); some products are newly introduced, discontinued, or heavily altered over time, so I use the version number provided by NielsenIQ to distinguish them. Observable characteristics are based on brand identity, making the study of additional goods easier. The observable characteristics are all binary variables: three indicate whether the product is from the highest/second-highest/third-highest-selling brand, and a fourth indicates whether the product's brand is in the top ten most sold.

Agent data and micro moments are derived from the consumer panel. To produce consumer-specific demographic draws for each market, I aggregate all panel households up by these demographic characteristics, essentially distilling households down into a particular demographic "type" within a city-year.[32] To improve optimization performance, I take the Kronecker product of these household types with a sparse grid of nodes/weights according to the level-two Gauss-Hermite quadrature rule. The final weights are the product of the quadrature weights and the portion of the household type's weight in the city. That is, if a household type contained 50% of all weights in a given city, the quadrature weights for that type would all be multiplied by 0.5.

To determine these types, I use seven binary demographic characteristics $y_{\cdot it}$, with three of these being interactions with smartphone access. The primary four are summarized in Table 8. Two of these deal with income, indicating whether the household earned between $50,000 and $100,000 USD or over $100,000 USD in the year prior.[33] The third indicates whether the household head's age is above 50, and the fourth,

[31] Number two is bread at 19.2% of all trips, and third is yogurt at 11.6%.

[32] Households are part of the consumer panel for an entire calendar year, so it would be inappropriate to split them out by quarter.

[33] These values are binned by NielsenIQ, so are not inflation-adjusted. However, the $50,000 threshold approximately aligns with the median income during the study period, and a household earning over $100,000 occupies NielsenIQ's top income category in almost all time periods. It is therefore a logical upper bound.

whether the household resides in a zip code that should have access to 3G and smartphones. Cross-time averages are included in Table 8.

Table 8: Summary statistics—panelists

HH Head > 50	Inc. >$100k	Inc. $50-$100k	Smartphone access
0.554	0.199	0.312	0.677
(0.497)	(0.399)	(0.463)	(0.467)

Notes: Author's calculations based on the NielsenIQ Consumer Panel dataset (2007-2017 only). Values are weighted averages/standard deviations based on respondents' projection factor.

3G access is assigned based on availability at the household's residential zip code, and predicted using the random forest model discussed in Section 4. The model for predicted mobile Internet coverage predicts outcomes at a 100-square-mile grid-cell level, which are used to scale down to the pertinent zip code level. For zip codes that are contained completely within one grid cell, their treatment timing is determined by that of the grid cell. For those that overlap, I produce a weighted average depending on how much area is covered by different cells—for instance, if one zip code was 50% covered by two different predicted coverage levels, then the result would be a simple average of those levels. The first "treated" year is the first year for which 3G coverage is predicted to be 50% or higher, after 2009, to account for the availability of smartphones. I include interaction terms for smartphone access and all demographics, as well as product characteristics (brand dummies), to allow estimations to flexibly change as smartphones are introduced.

All demand estimations include year and city-retailer fixed effects, to control for cross-time or cross-location shocks. In the interest of computational efficiency, estimations are performed at the Census subregion level.[34, 35]

6.2 Coefficients

Though the prices consumers pay and the market characteristics of where they shop do not seem to change in a reduced-form way, demand estimation lets us explore further implications of these potential stories. I therefore start by determining whether smartphones make users more price elastic. The common stories in the popular press (Bustillo and Zimmerman, 2010; Reisinger, 2010) focus almost exclusively on smartphones enabling users to price shop. It is therefore worth exploring this implication via demand estimation before completely writing off this narrative.

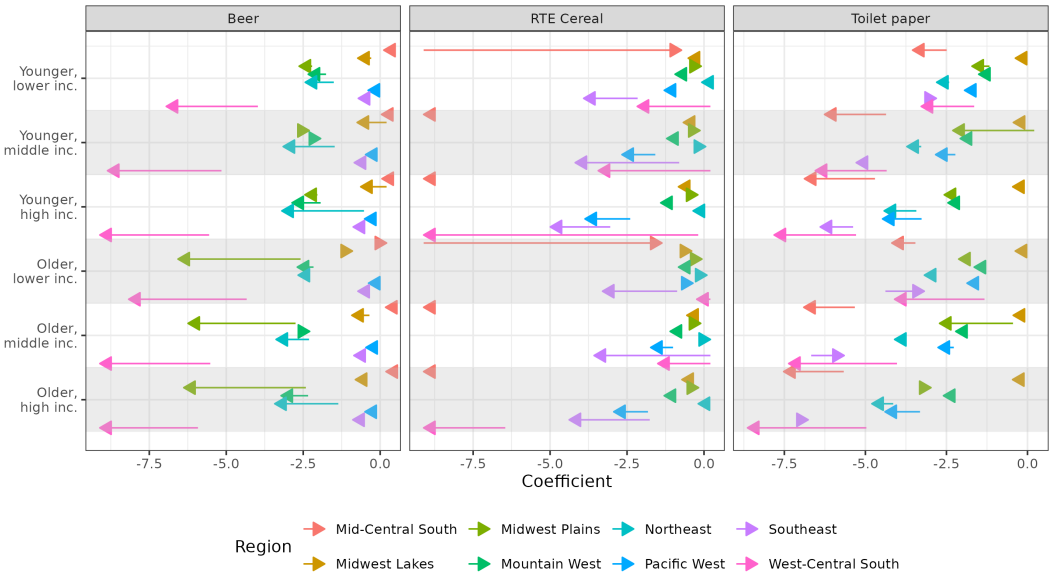
Note that, while the reduced-form analyses in Section 5 dealt with aggregate consumer baskets, demand estimation necessarily works in one product market at a time. I present results for three product markets here; future versions of this paper will expand on these markets, and more information can be found in future appendices.

Because the price coefficient in these estimations is logged, all consumers must behave as though they have a downward-sloping demand curve, meaning that price elasticities will necessarily be negative. Interpreting the coefficients in Figure 7 correctly requires keeping that in mind. This means that positive/larger coefficients indicate *larger* responses with respect to price, and negative coefficients a smaller one. The logged coefficients on the interaction of price and smartphone access in Figure 7 vary somewhat by product market, region, and demographic group, though generally become more negative following smartphone access. This

[34] As one example, the Pacific West subregion consists of California, Oregon, and Washington. All cities in those areas are included in that estimation; this also means that year fixed effects operate more like year-Census subregion fixed effects.

[35] Optimization uses `scipy`'s `trust-constr` routine with a radius of 1 and tolerances of 1 x 10^{-5} to reduce errors.

Figure 7: Coefficients by demographic group



Notes: This figure depicts the resulting coefficients on price before and after smartphone access for consumers of different demographic groups. For numerical details and statistical significance, see Tables A.6, A.7, and A.8 in the Appendix.

means that smartphones in these markets generally serve to make consumers *less* price elastic, driving their elasticities towards zero. Given the number of parameters that can influence a consumer's ultimate price elasticity, however, I prefer to examine elasticities directly.

6.3 Price elasticities

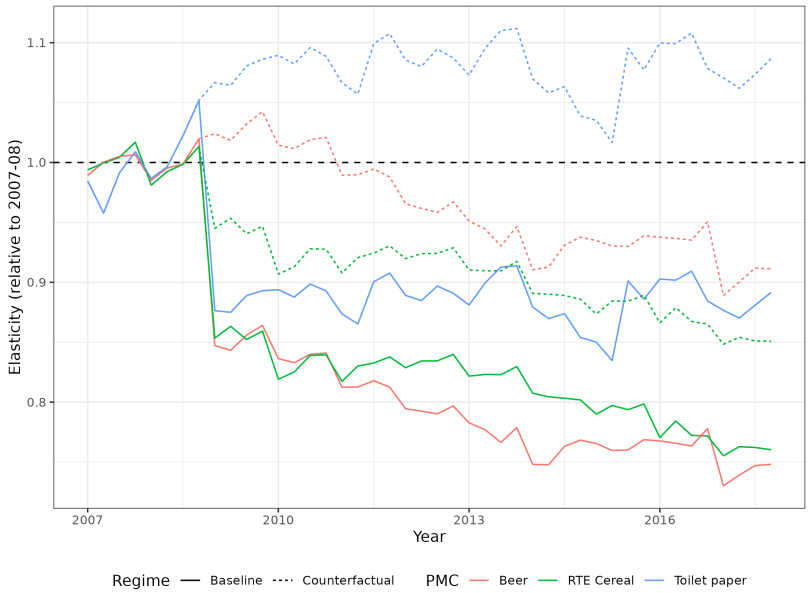
Figure 8 shows that, proportional to 2008, own-price elasticity shrinks in most markets after smartphones become available in 2009. The toilet paper market exhibits the smallest drop examined here of roughly 12%, with larger decreases in the beer and cereal sectors (about 15%).[36]

Because this attenuation of price elasticities is operating on negative numbers, this movement towards zero is also an *increase* in the value. That is, the difference in pre- and post-smartphone price elasticities is generally positive. Figure 9 shows how this change is correlated with pre-period price elasticities. While Figure 7 already established that changes tend to be towards zero (positive), larger absolute changes tend to occur where consumers are more price elastic in the pre-period; for instance, a toilet paper market with an average price elasticity of -5 before smartphones would move to -4.5, while a market with a price elasticity of -2.5 would move less, to -2.25. This results in a downward slope across different product markets in Figure 9, as well as a small convergence of own-price elasticities across markets. Own-price elasticities do not end up at the same point due to smartphones, but they spread out less.

The dashed lines in Figures 8 and 9 indicate the same evolution of price elasticities after removing the observed effects of smartphone access. This is accomplished by replacing agents' smartphone-related demographics with zeroes, then recalculating agents' purchase probabilities to use in deriving elasticities. Note that this does *not* change the choices consumers have access to—retailers, products, and prices are all

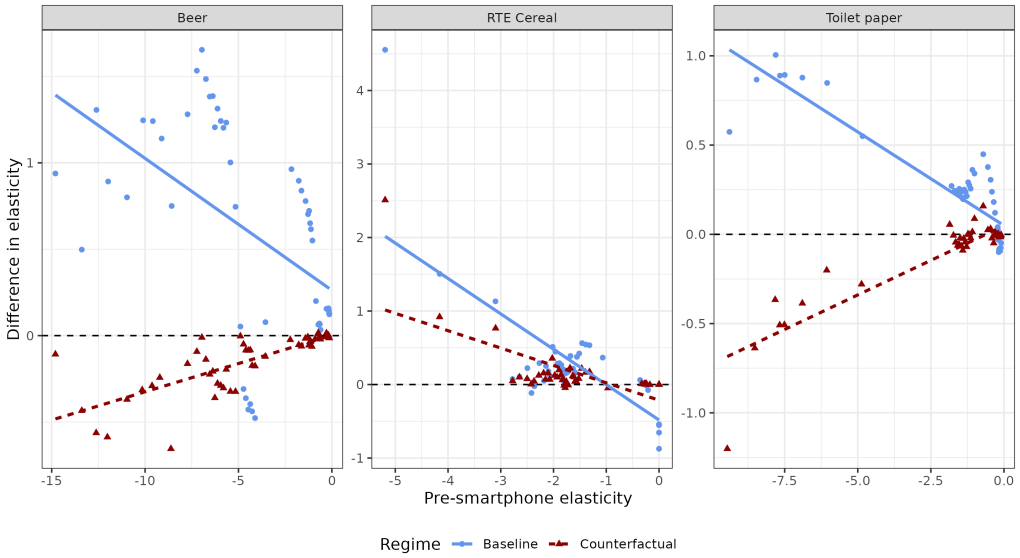
[36]My elasticity estimates tend to be close to those in the broader demand-estimation literature; the average elasticity I estimate for beer is -4.08, relatively close to the -4.22 measured in a working-paper version of Döpper et al. (2025). My average for cereal of -1.56 is also reasonably close to their estimate of -2.29.

Figure 8: Relative changes in price elasticity values



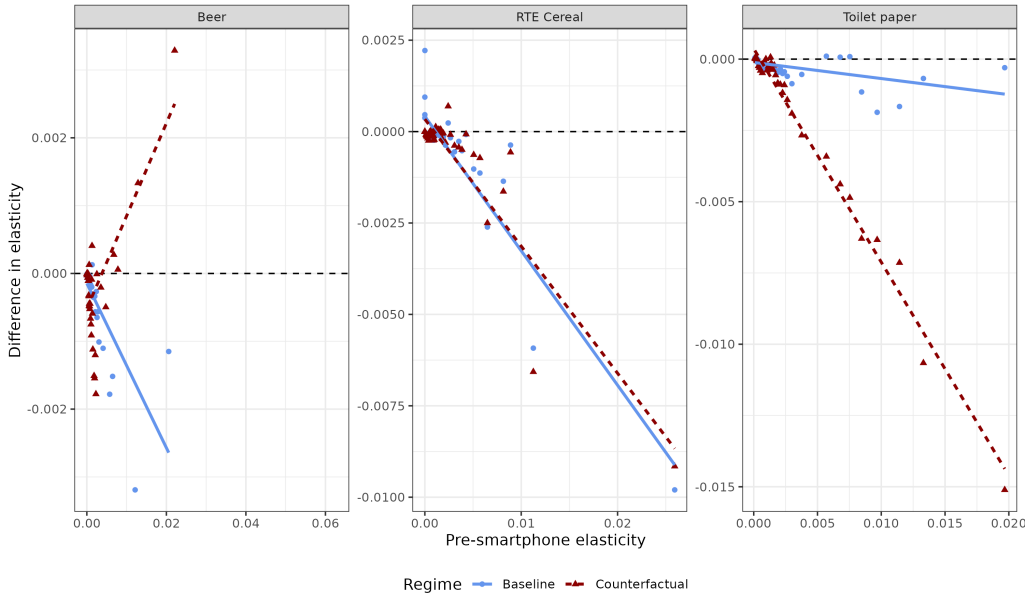
Notes: Price elasticities represent market-level average own-price elasticities, averaged across time by market size (by sales, USD). Elasticity derivation is performed as in Conlon and Gortmaker (2025) for both observed, "baseline" values and counterfactual values.

Figure 9: Changes in price elasticity, vs. counterfactual



Notes: Price elasticities represent market-level average own-price elasticities, derived as in Conlon and Gortmaker (2025) for both observed, "baseline" values and counterfactual values. Differences are calculated using average values for the year before and year after majority mobile coverage, which is modally 2008/2009. Figure includes a 50-point binscatter for both baseline and counterfactual elasticities, and lines are fit linearly.

Figure 10: Cross-price elasticity trends over time



Notes: Price elasticities represent market-level average own-price elasticities, derived as in Conlon and Gortmaker (2025) for both observed, "baseline" values and counterfactual values. Differences are calculated using average values for the year before and year after majority mobile coverage, which is modally 2008/2009. Figure includes a 50-point binscatter for both baseline and counterfactual elasticities, and lines are fit linearly.

unchanged. In the beer and toilet paper markets, the counterfactual line is actually slightly upward sloping, so the "net effect" of smartphones (on the consumer side alone) is even steeper than Figure 9 suggests. In the cereal market, some of the increase in elasticities can be ascribed to changes in product availability or prices, so the effect net of these changes is less steep than depicted in the figure.

6.3.1 Cross-price, brand elasticities

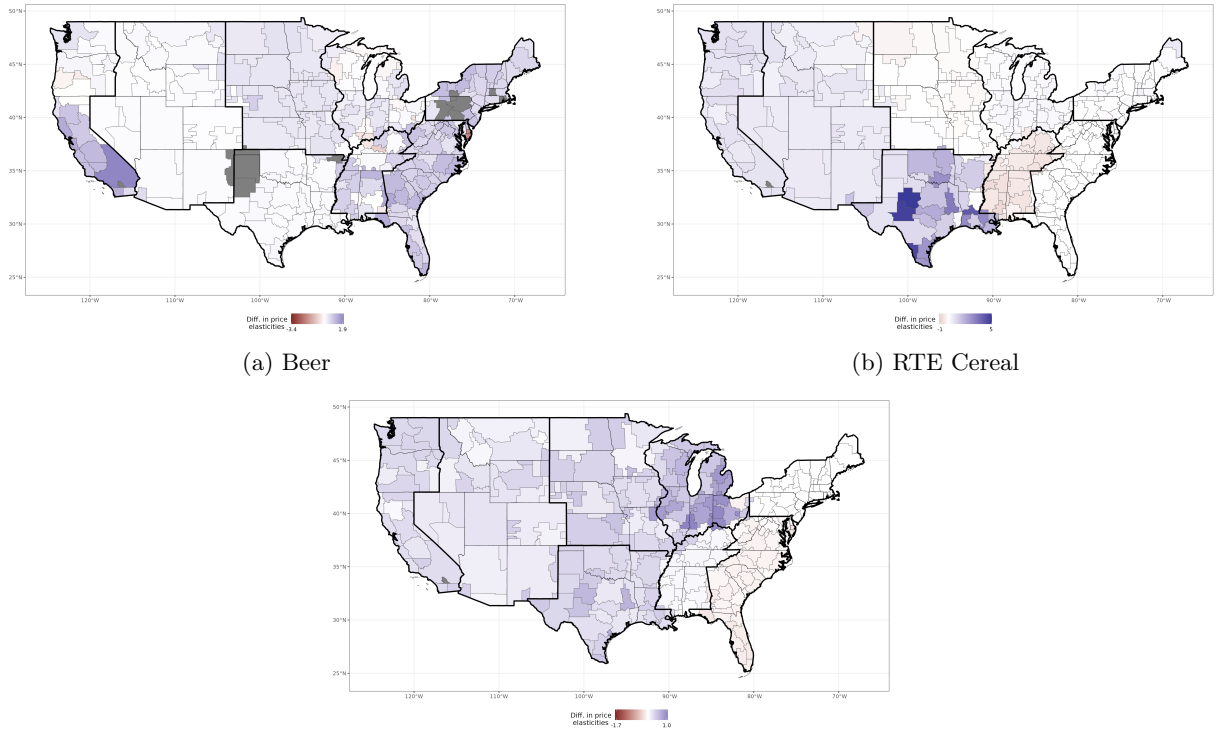
Figure 9 deals only with own-price elasticity; as this moves closer to zero, it could be that consumers focus more on between-product comparisons. With that in mind, I consider how cross-price elasticities might change, aggregating elasticity values across all products in a given market and taking the average. Figure 10 uses the same approach as Figure 9 for cross-price elasticities. Observed changes all seem to result in small decreases. However, the differences in counterfactual lines mean that it is not clear to what extent these decreases are caused or tempered by smartphone access.

I use brand dummies as product characteristics, which enables me to estimate brand elasticities as well. Smartphones similarly have no consistent effect on these brand values, though these characteristics differ in importance more between markets than price does.

6.4 Spatial heterogeneity

Figure 11 shows how the changes plotted in Figure 9 differ by city and Census subregion. Increases (movements towards zero) are generally larger for beer than for toilet paper. Beer markets see larger movements on the coasts, with the largest change taking place around Los Angeles, while the effects for toilet paper are largest in the Midwest around the Great Lakes. The East Coast is practically unaffected, based on the small

Figure 11: Heterogeneity by city, Census region



(a) Beer

(b) RTE Cereal

(c) Toilet paper

Notes: Price elasticities represent market-level average own-price elasticities, derived as in Conlon and Gortmaker (2025) for both observed, "baseline" values and counterfactual values. Differences are calculated using average values for the year before and year after majority mobile coverage, which is modally 2008/2009. Resulting elasticities and differences are averaged based on market size across retailers within city, so larger/higher-selling markets' elasticities count for more. Missing values in Panel (a) occur because there are no recorded stores in NielsenIQ for these cities in 2009, so there is nothing to estimate.

size of the movements for toilet paper.

Toilet paper is a necessity that households could conceivably shop around for, with products that are significantly less differentiated than different beers are. Missing markets in the beer industry due to state control ("alcoholic beverage control" states, or ABC states) likewise reveal a very different societal approach to these products.

Aside from the large impact on beer's price elasticity in Los Angeles, the most pronounced effects do not appear to be consistently in either small or large cities. In all three product markets, New York City does not change significantly more or less than other nearby urban areas.

6.5 Mechanisms

Previous empirical and theoretical arguments largely suggest that consumers' use of smartphones while shopping would change their behavior through relevant uses. Firms can show consumers relevant products and provide coupons, and consumers can compare prices more easily. This paper largely comes to the opposite conclusions. Consumers are not using coupons any more frequently, and are reacting less to price with time. At the same time, consumers with smartphone access are experimenting more with what they

buy, both in terms of what product markets they engage with and the consumer bases of those markets, being an atypical consumer more often.

Broadly, the question seems to be whether smartphone use in this context is likely to be relevant to shopping at all. Considering the effects I observe, it would be reasonable to think that smartphone access might allow consumers to discover products, or get more information about products in the grocery aisle. These arguments seem unlikely. Product discovery is possible, but would leverage the same channels as other forms of targeting—namely, advertising. Moreover, the simplest and most cost-effective way for a firm to promote their product would be to leverage previous customers, information that should be more available with the rise of retailer loyalty programs and/or data mining. Information in the grocery store seems largely irrelevant beyond what the consumer can already see on product packaging, such as ingredients or potential recipes. While online reviews become much more popular over this time period, it seems unreasonable to think that most consumers are engaging with grocery-industry reviews to make different decisions.

I instead turn to non-shopping-relevant smartphone usage as a likely driver. Indeed, Grewal et al. (2018) previously show that consumers who use smartphones (especially older consumers) tend to get distracted. They spend longer in the store and buy more products. Distraction—using smartphones in ways not pertinent to shopping—seems to be a much likelier story than "productive" uses have proven. I corroborate the findings in Grewal et al. and extend them. My findings are much more consistent with distraction being the standard mechanism, rather than smartphones being used for the express purpose of assisting a shopping trip.

7 Discussion

This paper explores how mobile Internet rollout and smartphone adoption changed consumers' in-store shopping behaviors. I first explore changes in consumer behavior, as consumers increasingly sort into distinctive purchasing patterns from one another. I leverage public data and machine learning to derive a high-quality measure of historical mobile Internet coverage. I use these data to conclude that smartphones surprisingly offset increases in aggregate sorting behavior, as consumers buy more distinct goods. Contrary to the common narrative about comparing prices across retailers, consumers also spend a larger proportion of shopping trips at their modal/favorite retailer. Demand estimation provides more details on these changes, showing that consumers are becoming less price elastic. Together, these results are inconsistent with smartphones letting consumers compare prices, but are consistent with mobile devices distracting consumers as they shop.

Consumers that are willing to experiment could also wind up regretting their decisions, with producers possibly taking advantage of this willingness to experiment—or a lack of information as in Albrecht, H. Lang, and Vroman (2002). More work is required to build on the findings in Döpper et al. (2025) and understand the general causes of lower price responsiveness and rising markups, and future versions of this paper will attempt to tackle this issue. However, the implications of any shifts require careful thought.

These results may be unique to the grocery industry. In settings that deal less with necessities, such as entertainment, firms have heavily invested into predictive methods and algorithmic recommendation systems. Many papers discuss the "fragmentation" of pop culture (Paans, 2020; Sugihartati, 2019) and the proliferation of niche hobbies/interests with the rise of Internet culture. Algorithmic recommendation systems can serve either to homogenize consumer tastes or further serve heterogeneous preferences. The divergence of consumers in other realms may serve as a counterbalance to the high-level trends described here—outside of the early COVID-19 pandemic, for instance, it is unlikely that toilet paper will become "trendy" and subject to product differentiation in the same way as other consumer goods.

This paper is one of the first to study offline impacts of online access. Given the importance of the Internet to daily life, further study of this topic is well warranted.

References

Albrecht, James, Harald Lang, and Susan Vroman (2002). "The effect of information on the well-being of the uninformed: what's the chance of getting a decent meal in an unfamiliar city?" In: *International Journal of Industrial Organization* 20.2, pp. 139–162.

Ataman, Berk et al. (2023). "Advertising's Impact on Brand Price Elasticity". In: *Boston University Questrom School of Business Research Paper* 4694207.

Backus, Matthew, Christopher Conlon, and Michael Sinkinson (2021). *Common ownership and competition in the ready-to-eat cereal industry*. Tech. rep. National Bureau of Economic Research.

Bar-Isaac, Heski and Sandro Shelegia (2023). "Search, showrooming, and retailer variety". In: *Marketing Science* 42.2, pp. 251–270.

Bart, Yakov, Andrew T Stephen, and Miklós Sárváry (2012). *A Field Study of the Determinants of Mobile Advertising Effectiveness*. INSEAD.

Benkard, C Lanier, Ali Yurukoglu, and Anthony Lee Zhang (2026). "Concentration in product markets". In: *American Economic Journal: Microeconomics* 18.2, pp. 268–292.

Berry, Steven, James Levinsohn, and Ariel Pakes (1995). *Automobile Prices in Market Equilibrium*. URL: http://www.jstor.org/stable/2171802 (visited on 10/21/2026).

Berry, Steven T and Philip A Haile (2014). "Identification in differentiated products markets using market level data". In: *Econometrica* 82.5, pp. 1749–1797.

Biswas, Shirsho, Hema Yoganarasimhan, and Haonan Zhang (2025). "How does online shopping affect offline price sensitivity?" In: *arXiv preprint arXiv:2506.15103*.

Book, Theodore and Dan S Wallach (2015). "An empirical study of mobile ad targeting". In: *arXiv preprint arXiv:1502.06577*.

Bustillo, Miguel and Ann Zimmerman (Dec. 15, 2010). "Phone-Wielding Shoppers Strike Fear Into Retailers". In: *The Wall Street Journal*. URL: https://www.wsj.com/articles/SB10001424052748704694004576019691769574496 (visited on 08/10/2026).

Callaway, Brantly and Pedro HC Sant'Anna (2021). "Difference-in-differences with multiple time periods". In: *Journal of econometrics* 225.2, pp. 200–230.

Chamberlain, Gary (1987). "Asymptotic efficiency in estimation with conditional moment restrictions". In: *Journal of econometrics* 34.3, pp. 305–334.

Chan, Jason et al. (2021). "The role of physical stores in the digital age: Quasi-experimental evidence from product level analysis". In: *Available at SSRN 3762286*.

Chu, Junhong, Pradeep Chintagunta, and Javier Cebollada (2008). "Research note—A comparison of within-household price sensitivity across online and offline channels". In: *Marketing science* 27.2, pp. 283–299.

Clifford, Stephanie (Mar. 11, 2009). "Advertisers Get a Trove of Clues in Smartphones". In: *The New York Times*. URL: https://www.nytimes.com/2009/03/11/business/media/11target.html (visited on 08/10/2026).

Conlon, Christopher and Jeff Gortmaker (2025). "Incorporating micro data into differentiated products demand estimation with PyBLP". In: *Journal of Econometrics*, p. 105926.

Dagorret, Anna and Fang Guo (2023). "The impact of 3G internet on political accountability: evidence from the US House of Representatives". In: *Work. Pap., Stanford Grad. Sch. Bus., Stanford, CA:. https://fangguo34.github.io/files/Political_Accountability_Draft_May_2023. pdf*.

Döpper, Hendrik et al. (2025). "Rising markups and the role of consumer preferences". In: *Journal of Political Economy* 133.8.

Erdem, Tülin, Michael P Keane, and Baohong Sun (2008). "The impact of advertising on consumer price sensitivity in experience goods markets". In: *Quantitative Marketing and Economics* 6.2, pp. 139–176.

Farahat, Ayman and Michael C Bailey (2012). "How effective is targeted advertising?" In: *Proceedings of the 21st international conference on World Wide Web*, pp. 111–120.

Gensler, Sonja, Scott A Neslin, and Peter C Verhoef (2017). "The showrooming phenomenon: it's more than just about price". In: *Journal of Interactive Marketing* 38.1, pp. 29–43.

Grewal, Dhruv et al. (2018). "In-store mobile phone use and customer shopping behavior: Evidence from the field". In: *Journal of Marketing* 82.4, pp. 102–126.

Hansell, Saul (Nov. 3, 2009). "AT&T Declares Color War on Verizon". In: *The New York Times*. URL: `https://archive.nytimes.com/bits.blogs.nytimes.com/2009/11/03/att-declares-color-war-on-verizon/` (visited on 08/13/2026).

Hausman, Jerry A (1996). "Valuation of new goods under perfect and imperfect competition". In: *The economics of new goods*. University of Chicago Press, pp. 207–248.

Hersh, Jonathan, Bree J Lang, and Matthew Lang (2022). "Car accidents, smartphone adoption and 3G coverage". In: *Journal of Economic Behavior & Organization* 196, pp. 278–293.

Hill, Kashmir (Feb. 16, 2012). "How Target Figured Out A Teen Girl Was Pregnant Before Her Father Did". In: *Forbes*. URL: `https://www.forbes.com/sites/kashmirhill/2012/02/16/how-target-figured-out-a-teen-girl-was-pregnant-before-her-father-did/` (visited on 08/10/2026).

Hoch, Stephen J et al. (1995). "Determinants of store-level price elasticity". In: *Journal of marketing Research* 32.1, pp. 17–29.

Kaul, Anil and Dick R Wittink (1995). "Empirical generalizations about the impact of advertising on price sensitivity and price". In: *Marketing Science* 14.3_supplement, G151–G160.

Kim, Byung-Do, Kannan Srinivasan, and Ronald T Wilcox (1999). "Identifying price sensitive consumers: the relative merits of demographic vs. purchase pattern information". In: *Journal of Retailing* 75.2, pp. 173–193.

Krishnamurthi, Lakshman and Sethuraman P Raj (1985). "The effect of advertising on consumer price sensitivity". In: *Journal of Marketing Research* 22.2, pp. 119–129.

Lau, Yan (2020). *A brief primer on the economics of targeted advertising*. Bureau of Economics, Federal Trade Commission.

Leung, Justin H and Zhonglin Li (2026). "Rising Concentration of Household Shopping, Superstar Firms, and Implications for Retail Markups". In: *American Economic Journal: Applied Economics* 18.1, pp. 160–194.

MacKay, Alexander and Marc Remer (2026). "Mergers in Markets with Consumer Inertia". In: *Available at SSRN 4912975*.

Morton, Fiona Scott (2006). "Consumer benefit from use of the internet". In: *Innovation policy and the economy* 6, pp. 67–90.

Myers, Caitlin K and Ezekiel Hooper (2026). *Is the iPhone Birth Control? Causal Evidence from AT&T's 2007–2011 Carrier Monopoly*. Tech. rep. National Bureau of Economic Research.

Neiman, Brent and Joseph Vavra (2023). "The rise of niche consumption". In: *American Economic Journal: Macroeconomics* 15.3, pp. 224–264.

Paans, Otto (2020). "Postmodernity and the Politics of Fragmentation". In: *Borderless Philosophy* 3, pp. 264–310.

Petrin, Amil (2002). "Quantifying the benefits of new products: The case of the minivan". In: *Journal of political Economy* 110.4, pp. 705–729.

Reisinger, Don (Nov. 22, 2010). "Survey: Smartphones altering holiday-shop habits". In: *CNET*. URL: https://www.cnet.com/home/smart-home/survey-smartphones-altering-holiday-shop-habits/ (visited on 08/10/2026).

Smith, Dominic A and Sergio Ocampo (2025). "The evolution of US retail concentration". In: *American Economic Journal: Macroeconomics* 17.1, pp. 71–101.

Stein, Joel (Mar. 10, 2011). "Data Mining: How Companies Now Know Everything About You". In: *Time*. URL: https://time.com/archive/6595428/data-mining-how-companies-now-know-everything-about-you/ (visited on 08/10/2026).

Sugihartati, Rahma (2019). "The identity fragmentation of youths as fans of global popular culture". In: *Pertanika Journal of Social Sciences & Humanities* 27.2, pp. 1007–1022.

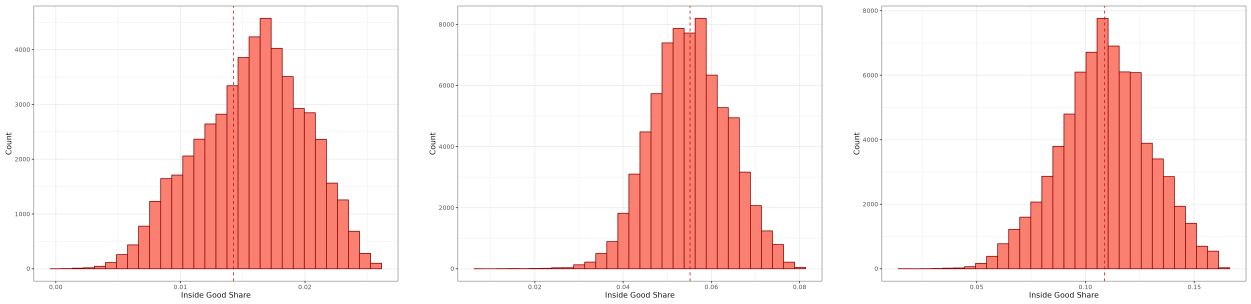
Tran, Uyen (2026). "The rise of broadband and the retail landscape: Evidence from consumer packaged goods". In: *Journal of Marketing* 90.5, pp. 99–124.

Wood, Zoe (Dec. 10, 2010). "Smartphones changing the way we shop". In: *The Guardian*. URL: https://www.theguardian.com/technology/2010/dec/10/smartphones-changing-way-we-shop (visited on 08/10/2026).

Yan, Jun et al. (2009). "How much can behavioral targeting help online advertising?" In: *Proceedings of the 18th international conference on World wide web*, pp. 261–270.

A Supplemental information

Figure A.1: Inside shares



(a) Beer (b) Toilet paper (c) RTE Cereal

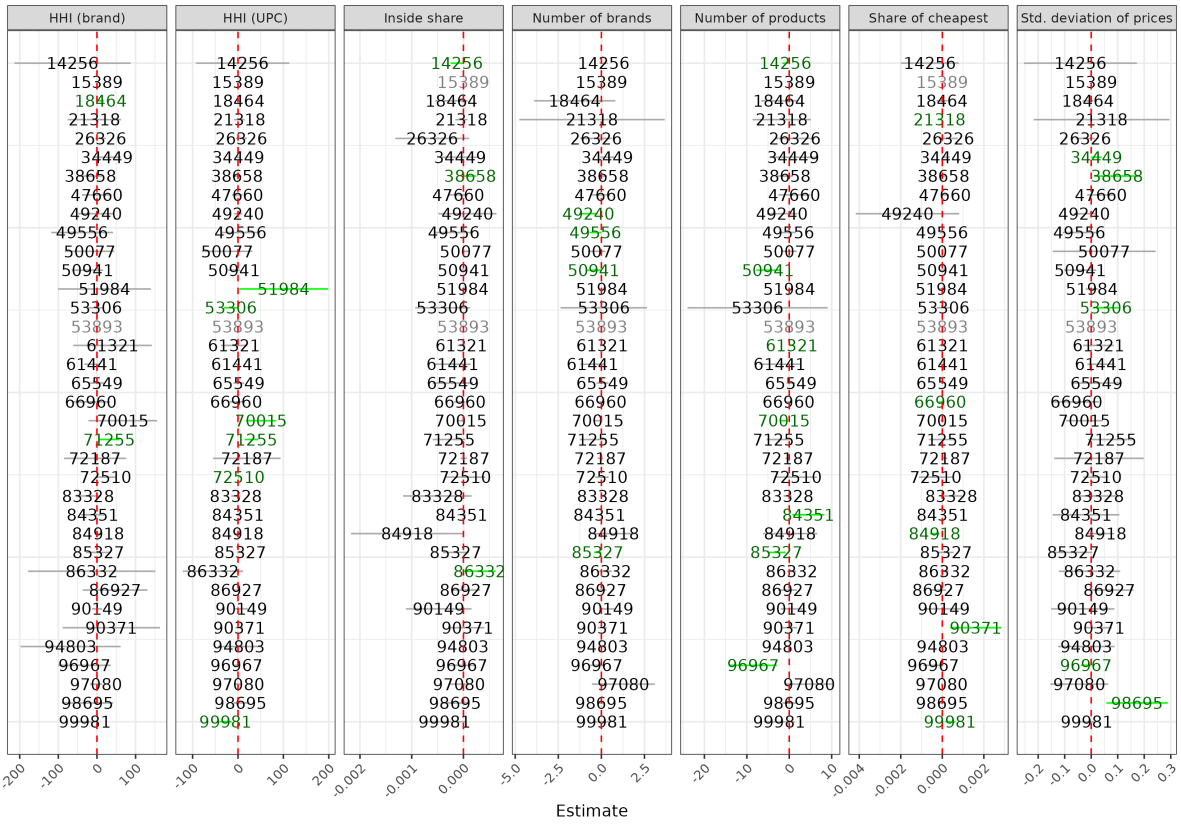
Notes: The above figure shows the distribution of total inside shares across markets for the three main goods analyzed here. The dashed line indicates the mean inside share value, used to normalize the predicted total shares as in Backus, Conlon, and Sinkinson (2021).

Table A.1: Product module code reference

PMC	Description
14256	SUGAR-GRANULATED
15389	WINE-IMPORTED DRY TABLE
18464	CEREAL - READY TO EAT
21318	BEER
26326	SOFT DRINKS - CARBONATED
34449	ICE CREAM - BULK
38658	DETERGENTS - HEAVY DUTY - LIQUID
47660	TEA - LIQUID
49240	BAKERY - BREAD - FRESH
49556	NUTS - CANS
50077	DOG FOOD - DRY TYPE
50941	TEA - BAGS
51984	BAKING SODA
53306	COOKIES
53893	TEA-HERBAL PACKAGED
61321	BAGS - SANDWICH
61441	FRUIT DRINKS-OTHER CONTAINER
65549	CAT FOOD - DRY TYPE
66960	BAGS - TALL KITCHEN
70015	COFFEE - LIQUID
71255	WATER-BOTTLED
72187	FLOUR-ALL PURPOSE-WHITE WHEAT
72510	YOGURT-REFRIGERATED
83328	TOILET TISSUE
84351	DISPOSABLE DIAPERS
84918	CANDY-CHOCOLATE
85327	CAT FOOD - WET TYPE
86332	FRESH MEAT
86927	CHEESE - SHREDDED
90149	SOFT DRINKS - LOW CALORIE
90371	PRECUT FRESH SALAD MIX
94803	SALT - TABLE
96967	SOUP-CANNED
97080	NUTRITIONAL SUPPLEMENTS
98695	BATTERIES
99981	SNACKS - POTATO CHIPS

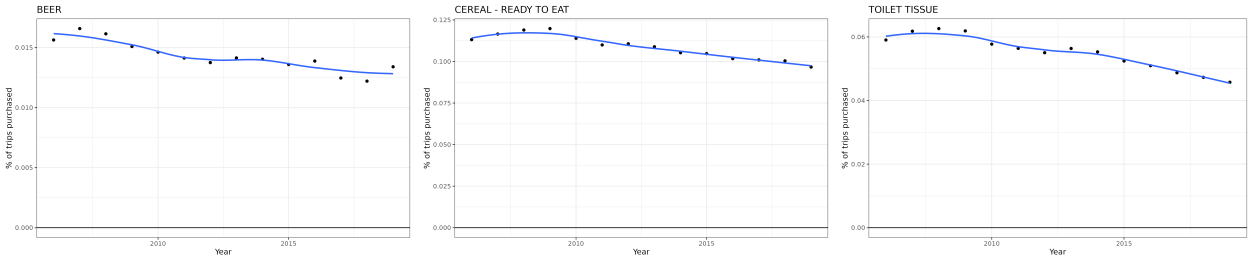
Notes: Descriptions come from the NielsenIQ data codebook. Product module codes are anonymized by the author. Herbal tea is grayed out as estimations failed due to numerical issues.

Figure A.2: Product-market reduced form



Notes: This figure plots the estimated effect of a Callaway and Sant'Anna (2021) estimation at the product module code (PMC) level, using city-retailer and year-quarter fixed effects. The position of the PMC indicates the estimated effect, with the horizontal lines indicating the 95% confidence interval. Statistically significant results appear in green, while insignificant ones are in black/gray. Product module code descriptions can be found in Table A.1. To preserve data confidentiality, codes have been anonymized.

Figure A.3: Changes in inside-good share over time



(a) Beer (b) RTE Cereal (c) Toilet paper

Notes: Averages are calculated using the NielsenIQ Consumer Panel, using the product module codes for beer, cereal, and toilet paper. These are simple averages of the proportion of trips that include a purchase in a particular category. Note that the current construction of inside good share is based on a cross-time average, not the (slightly decreasing) values in both of the above categories. Note also the different y-axis scales between these product categories.

Table A.2: Siloing evidence, brand level

	Estimate	Sample mean
Prop. from top 3 brands (in PMC)	-0.011*	0.546
	(0.005)	
Prop. from top 5 brands (in PMC)	-0.022*	0.654
	(0.007)	
Prop. from top 10 brands (in PMC)	-0.019*	0.774
	(0.007)	
Household HHI (brand)	-66.321	8,016
	(45.712)	
In plurality, age/presence of children	-0.097*	0.689
	(0.036)	
In plurality, Hispanic origin	-0.032	0.915
	(0.019)	
In plurality, household size	-0.034*	0.399
	(0.012)	
In plurality, marital status	-0.011	0.551
	(0.028)	
In plurality, female head age	-0.002	0.298
	(0.011)	
In plurality, household income	0.047	0.251
	(0.042)	
In plurality, male head age	-0.019	0.355
	(0.023)	
In plurality, race	-0.022	0.831
	(0.014)	

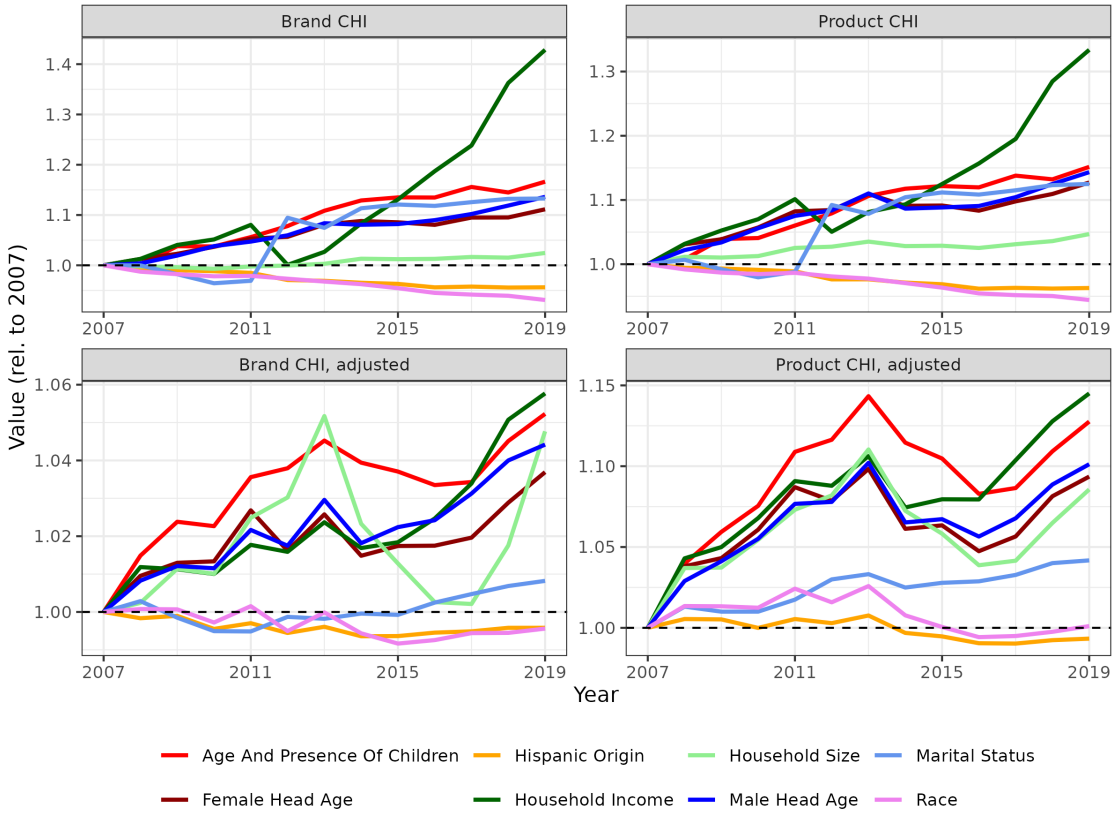
Notes: The estimates presented here are the result of Callaway and Sant'Anna (2021)-style estimations, using the city and year-quarter as area and time fixed effects, respectively. This table also controls for local median income values from the ACS. For concision, this table includes only brand-level outcomes; product-level outcomes are similar. Note that "treatment" is determined at the city level, and all cities are considered majority "treated" by Q2 of 2011. This limits the time horizon to estimate post-treatment effects.

Table A.3: Demographic bin breakdown

Category	Levels
Age and Presence of Children	Under 6 only, 6-12 only, 13-17 only, under 6 and 6-12, under 6 and 13-17, 6-12 and 13-17, under 6 and 6-12 and 13-17, no children under 18
Hispanic Origin	Yes, No
Household Size	1, 2, 3, 4, 5, 6, 7, 8, 9+
Marital Status	Married, widowed, divorced/separated, single
Head Age (Male/Female)	<25, 25-29, 30-34, 35-39, 40-44, 45-49, 50-54, 55-65, 65+, no relevant head
Household Income	Under $5000, $5000-$7999, $8000-$9999, $10,000-$11,999, $12,000-$14,999, $15,000-$19,999, $20,000-$24,999, $25,000-$29,999, $30,000-$34,999, $35,000-$39,999, $40,000-$44,999, $45,000-$49,999, $50,000-$59,999, $60,000-$69,999, $70,000-$99,999, $100,000+
Race	White, Black, Asian, Other

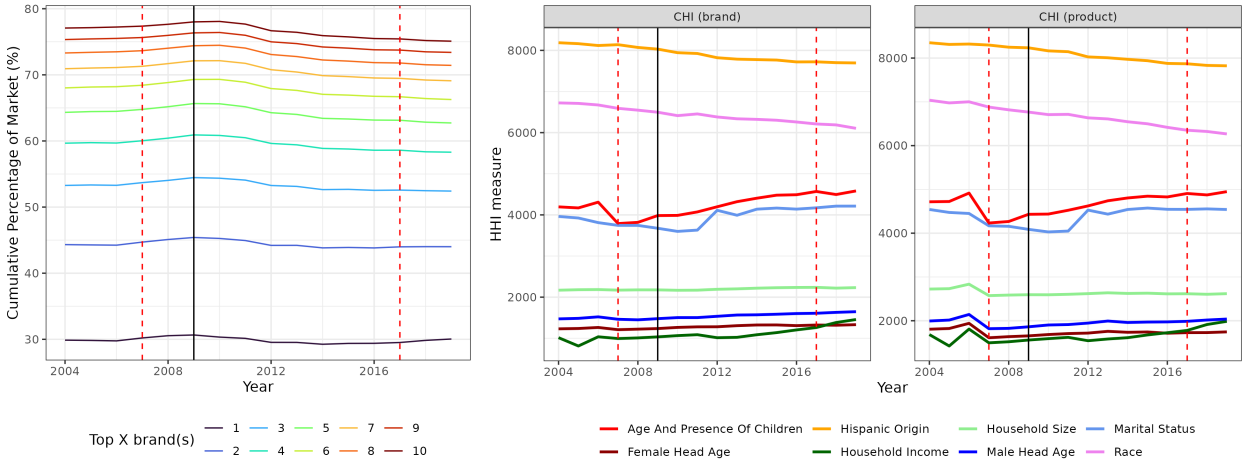
Notes: Information sourced from the NielsenIQ consumer data codebook.

Figure A.4: Consumer sorting, homogeneity



Notes: The consumer homogeneity index (CHI) is calculated identically to the HHI, but employs the demographic bins calculated by NielsenIQ rather than firm/brand identity. For instance, a brand that received 85% of its sales from White customers and 15% from Black customers would have a CHI of 7,450. Due to noise prior to 2007, values are relative to 2007, and preceding values are omitted. "Corrected" values in the lower two panels account for changing population shares of these demographic categories, using that adjusted share as the input to the CHI rather than a raw proportion. These cross-industry values are weighted across product/brand by total sales of that product/brand.

Figure A.5: Consumer sorting (absolute)



(a) Share of top X brands

(b) CHI (product/brand-level)

Notes: This figure includes absolute versions of the data presented in Figures 1 and 3. Both panels are the author's calculations using the NielsenIQ Consumer Panel, 2004-2019. Note that Panel (a) only includes product categories with at least ten brands, so as to avoid conflating observed differences with differences in selection.
While the Consumer Panel does not include as many goods as the Retail Scanner Dataset, there are two advantages to its use in this setting. First, the Consumer Panel includes purchases from stores that are not in the Scanner Dataset, allowing for a fuller picture of consumer purchasing behavior. Secondly, as it includes all purchases, it is simpler to get a suggestive picture across all purchases; equivalent measures from the Scanner Dataset would require working with several terabytes of data.

Table A.4: Fractional logit model on random-forest training data

Model:	(1)
Variables	
Constant	-697.2***
	(10.2)
Year Built x Elev. SD (m) w/in 20 mi	-0.0004***
	(1.3×10^{-5})
Year Built x Overall Height Above Ground (m)	-0.0002***
	(5.16×10^{-5})
Towers w/in 20 mi x Overall Height Above Ground (m)	3.36×10^{-5}***
	(5.14×10^{-6})
Towers w/in 20 mi squared	-7.77×10^{-6}***
	(7.85×10^{-7})
Towers w/in 20 mi	0.0008*
	(0.0005)
Year Built	0.348***
	(0.005)
Elev. SD w/in 20 mi	0.883***
	(0.026)
Towers w/in 20 mi x Elev. SD w/in 20 mi	1.44×10^{-5}***
	(1.54×10^{-6})
Overall Height Above Ground (m) squared	-6.33×10^{-6}***
	(1.08×10^{-6})
Overall Height Above Ground (m)	0.445***
	(0.103)
Height Above Local Elevation (m)	-0.0004***
	(1.34×10^{-5})
Distance to Nearest Tower (m)	0.0004***
	(7.31×10^{-5})
Distance to Nearest Tower (m) squared	-0.066***
	(0.002)
Fit statistics	
Observations	205,957
Squared Correlation	0.470
Pseudo R^2	0.323
BIC	171,955.8

Clustered (grid-cell) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Table A.5: 3G rollout dates

Year	Proportion of ZCTAs with 3G
pre-2005	0.004
2005	0.274
2006	0.489
2007	0.619
2008	0.736
2009	0.834
2010	0.929
2011	0.951
2012	0.964
2013	0.972
2014	0.979
2015	0.984
2016	0.985
2017	0.986

Notes: Author's calculations based on random forest predictions of 3G coverage.

Table A.6: BLP coefficient estimates—beer

	Midwest Lakes	Midwest Plains	Mountain West	Northeast	Pacific West	Southeast	Mid-Central South	West-Central South
1 (σ)	0.198	0.011	0.046	0.489	0.099	0.058	0.125	0.017
	(24.817)	(1918.720)	(86.858)	(18.688)	(53.728)	(2432.666)	(2639.625)	(500.845)
-Price (σ)	0.272	0.008	1.043	0.000	0.149	0.247	0.161	1.375
	(44.061)	(14.033)	(2.811)	(5447.058)	(2.651)	(2.944)	(71.371)	(18.398)
#1 brand (σ)	0.305	0.125	0.134	0.338	0.198	0.196	0.103	0.058
	(29.415)	(1112.109)	(5.121)	(215.597)	(14.068)	(377.102)	(852.731)	(538.178)
#2 brand (σ)	0.150	0.138	0.236	0.864	0.317	0.270	0.167	0.235
	(70.998)	(4648.616)	(2.410)	(76.664)	(8.472)	(223.083)	(928.806)	(160.411)
#3 brand (σ)	0.134	0.027	0.049	0.309	0.102	0.156	0.360	0.043
	(49.562)	(30323.499)	(6.968)	(239.509)	(29.797)	(91.784)	(149.175)	(1105.550)
Top-10 brand (σ)	0.322	0.488	0.191	0.659	0.303	0.225	0.075	1.211
	(20.233)	(68.354)	(2.092)	(27.044)	(6.950)	(199.974)	(835.664)	(27.457)
1 x Annual income > \$100,000	1.615	0.740	0.315	0.459	0.001	0.007	0.937	0.569
	(2.819)	(1.468)	(1.172)	(0.550)	(0.274)	(59.114)	(49.016)	(0.585)
1 x Annual income \$50-100,000	0.876	0.468	0.503	0.077	-0.006	-0.004	0.320	0.266
	(8.706)	(0.716)	(0.503)	(0.425)	(0.237)	(53.996)	(50.292)	(0.378)
1 x Household head > 50	-1.197	-1.081	-0.283	-0.609	-0.014	-0.029	0.438	-0.151
	(12.452)	(0.588)	(0.530)	(0.343)	(0.168)	(21.439)	(76.626)	(0.217)
1 x Smartphones	-0.182	-0.101	-0.068	-2.042*	-0.003	-0.003	0.184	-0.144
	(77.225)	(0.825)	(0.254)	(0.529)	(0.870)	(161.789)	(156.021)	(0.792)
-Price x 1	-0.290	-2.210*	-1.759	-1.507*	-0.229	-0.432	1.269	-3.973
	(55.450)	(0.401)	(2.477)	(0.453)	(2.719)	(18.308)	(179.684)	(11.529)
-Price x Annual income > \$100,000	0.508	0.172	-0.176	0.981	-0.074	-0.126	1.623	-1.583
	(33.309)	(1.017)	(1.801)	(0.913)	(0.050)	(14.157)	(573.007)	(4.786)
-Price x Annual income \$50-100,000	0.642	-0.163	-0.443	0.029	-0.046	-0.092	2.267	-1.183
	(20.158)	(0.538)	(0.795)	(0.227)	(0.057)	(15.127)	(46.093)	(3.858)
-Price x Household head > 50	-0.701	-0.376	-0.407	-0.834*	-0.001	-0.002	-1.089	-0.363
	(51.206)	(0.712)	(0.868)	(0.203)	(0.037)	(14.141)	(117.255)	(1.669)
-Price x Smartphones	-0.438	-0.423	-0.578	-0.928*	-0.162*	-0.278	-1.151	-2.976
	(27.660)	(1.110)	(0.477)	(0.355)	(0.061)	(11.064)	(124.006)	(9.933)
-Price x Household head > 50 x Smartphones	0.528	-3.559	0.050	0.603	0.014	-0.010	1.220	-0.857
	(40.470)	(6.263)	(0.199)	(0.316)	(0.044)	(17.310)	(146.259)	(34.760)
-Price x Annual income > \$100,000 x Smartphones	-0.412	0.015	-0.351	-1.740*	-0.040	-0.041	-1.682	-0.709
	(27.208)	(1.909)	(0.643)	(0.864)	(0.047)	(6.998)	(579.576)	(255.285)
-Price x Annual income \$50-100,000 x Smartphones	-0.662	0.490	0.846*	-0.746*	-0.037	-0.039	-2.340	-0.726
	(21.180)	(0.907)	(0.157)	(0.270)	(0.062)	(8.157)	(29.184)	(88.102)
#1 brand x 1	0.160	-0.459	0.206	0.371	0.106	0.156	-2.596	-0.121
	(9.807)	(61.513)	(0.263)	(26.447)	(0.755)	(45.742)	(268.603)	(8.977)
#1 brand x Smartphones	-0.064	-0.185	-0.244*	-0.352*	0.038	0.052	2.513	-0.340
	(16.448)	(0.141)	(0.057)	(0.082)	(0.217)	(23.531)	(20.238)	(0.779)
#2 brand x 1	0.139	-1.071	0.246	0.865	0.238	0.233	-0.435	-0.074
	(29.254)	(283.630)	(0.204)	(20.699)	(0.572)	(49.589)	(304.995)	(10.871)
#2 brand x Smartphones	0.260	-0.487*	-0.480*	-0.436*	0.134	0.130	-0.686	-0.493
	(5.018)	(0.068)	(0.053)	(0.140)	(0.275)	(24.397)	(61.287)	(1.048)
#3 brand x 1	-0.394	-0.271	0.361*	0.595	-0.012	0.114	2.131	0.011
	(25.945)	(360.647)	(0.131)	(27.028)	(0.683)	(9.832)	(441.934)	(13.851)
#3 brand x Smartphones	-0.470	-0.418*	-0.403*	-0.501*	0.003	0.046	-2.505	-0.222
	(4.591)	(0.108)	(0.049)	(0.092)	(0.264)	(15.298)	(75.886)	(1.394)
Top-10 brand x 1	0.929	-0.022	-0.061	0.243	0.203	0.187	1.749	-0.216
	(21.383)	(15.660)	(0.135)	(6.924)	(0.576)	(5.318)	(143.769)	(10.716)
Top-10 brand x Smartphones	-0.938	0.103	0.094*	0.146*	0.102	0.109	-2.799	0.187
	(31.768)	(0.100)	(0.046)	(0.063)	(0.139)	(9.112)	(10.674)	(0.265)

Notes: All estimates are the results of a BLP estimation, using `pyBLP` by Conlon and Gortmaker (2025). Estimates are performed in two stages; the second stage uses first-stage estimates as starting values, but using Chamberlain (1987)-style optimal instruments.

Table A.7: BLP coefficient estimates—cereal

	Midwest Lakes	Midwest Plains	Mountain West	Northeast	Pacific West	Southeast	Mid-Central South	West-Central South
1 (σ)	0.097	0.052	0.010	0.048	0.019	0.068	11.575	0.042
	(8.713)	(22.539)	(63.419)	(13.935)	(25.519)	(4.809)	(1202.170)	(5.703)
-Price (σ)	0.509	0.437	0.000	0.856	0.004	2.232*	1.190	0.002
	(0.264)	(0.577)	(290.936)	(1.444)	(22.378)	(0.866)	(3843.765)	(7.283)
#1 brand (σ)	0.246	0.236	0.008	0.237	0.086	0.018	0.811	0.000
	(0.616)	(0.653)	(501.678)	(0.544)	(8.781)	(18.705)	(4636.249)	(4114.447)
#2 brand (σ)	0.234	0.269	0.225	0.297	0.221	0.187	2.355	0.105
	(0.791)	(0.706)	(72.448)	(0.873)	(106.509)	(2.811)	(1211.463)	(118.441)
#3 brand (σ)	0.098	0.052	0.011	0.125	0.019	0.104	3.223	0.003
	(1.627)	(2.593)	(3469.173)	(1.472)	(330.599)	(3.409)	(1123.601)	(1580.640)
Top-10 brand (σ)	0.241	0.183	0.379	0.061	0.036	0.000	1.672	0.021
	(0.421)	(0.436)	(4.854)	(2.128)	(20.214)	(569.567)	(1340.157)	(23.981)
1 x Annual income > \$100,000	-0.099	0.180	-0.318	-0.033	-1.405*	0.200	-0.573*	-2.053*
	(0.199)	(0.172)	(0.203)	(0.157)	(0.355)	(0.202)	(0.093)	(0.097)
1 x Annual income \$50-100,000	-0.068	0.126	-0.338*	0.606*	-1.137*	0.454*	-0.771*	-0.635*
	(0.155)	(0.134)	(0.120)	(0.074)	(0.319)	(0.175)	(0.086)	(0.052)
1 x Household head > 50	-0.069	-0.139	-0.057	0.229*	0.641*	0.118	-0.601*	0.710*
	(0.143)	(0.110)	(0.155)	(0.075)	(0.115)	(0.386)	(0.046)	(0.113)
1 x Smartphones	-0.271*	0.147	-0.402*	-0.068	-0.501*	0.510	13.097	-0.168
	(0.056)	(0.099)	(0.051)	(0.050)	(0.110)	(0.568)	(3045.269)	(0.228)
-Price x 1	-0.412	-0.163	-0.698*	0.044	-1.083*	-2.156	-25.191*	6.529
	(0.382)	(0.587)	(0.099)	(1.376)	(0.263)	(1.674)	(0.314)	(23.313)
-Price x Annual income > \$100,000	-0.206	-0.102	-0.436*	-0.115	-1.314*	-0.889	-10.485	-6.717
	(0.120)	(0.064)	(0.115)	(0.092)	(0.302)	(1.472)	(33.223)	(23.388)
-Price x Annual income \$50-100,000	-0.104	-0.097	-0.291*	-0.094*	-0.499*	1.350	-23.945	8.656*
	(0.089)	(0.051)	(0.066)	(0.045)	(0.188)	(0.895)	(397.362)	(0.045)
-Price x Household head > 50	-0.034	0.061	0.070	0.027	0.576*	1.283	-16.744*	-6.268
	(0.085)	(0.047)	(0.076)	(0.047)	(0.069)	(2.304)	(0.044)	(23.362)
-Price x Smartphones	-0.100*	0.073*	-0.242*	-0.115*	-0.208*	-1.753	24.460*	-8.698
	(0.036)	(0.032)	(0.036)	(0.041)	(0.070)	(1.557)	(0.363)	(23.317)
-Price x Household head > 50 x Smartphones	0.147*	-0.038	0.041	0.138*	0.358*	-0.670	16.109*	8.201
	(0.032)	(0.022)	(0.032)	(0.043)	(0.067)	(1.041)	(0.044)	(23.346)
-Price x Annual income > \$100,000 x Smartphones	-0.112*	-0.008	-0.029	-0.167*	-1.260*	-0.195	-5.870	-6.400
	(0.039)	(0.020)	(0.052)	(0.035)	(0.336)	(0.974)	(33.223)	(121.041)
-Price x Annual income \$50-100,000 x Smartphones	-0.054	0.057*	0.020	0.219*	-0.875*	-1.625*	13.848	-9.923*
	(0.034)	(0.023)	(0.024)	(0.043)	(0.151)	(0.476)	(397.362)	(0.044)
#1 brand x 1	-0.282*	-0.264*	-0.933	-0.293*	-0.512	-0.073	3.234	7.054*
	(0.062)	(0.042)	(1.901)	(0.023)	(0.365)	(0.131)	(–)	(0.907)
#1 brand x Smartphones	-0.254*	-0.119*	-0.053*	-0.116*	-0.359*	-0.451*	-0.741*	-5.255*
	(0.018)	(0.021)	(0.017)	(0.025)	(0.019)	(0.130)	(0.017)	(0.108)
#2 brand x 1	0.246*	0.262*	0.216	0.277*	0.314	0.319	0.257	16.595*
	(0.062)	(0.062)	(7.489)	(0.023)	(11.466)	(0.187)	(–)	(5.878)
#2 brand x Smartphones	0.130*	0.166*	0.388*	0.161*	0.126	0.069	0.109*	0.557
	(0.034)	(0.033)	(0.046)	(0.044)	(0.090)	(0.270)	(0.037)	(0.445)
#3 brand x 1	-0.098	-0.003	-0.211	-0.756*	0.314	-0.472*	-1.418*	5.376*
	(0.058)	(0.042)	(18.516)	(0.023)	(3.023)	(0.136)	(0.065)	(2.327)
#3 brand x Smartphones	-0.128*	-0.240*	0.029	0.592*	-0.191*	-0.453*	-0.410*	-4.049*
	(0.022)	(0.021)	(0.100)	(0.030)	(0.029)	(0.169)	(0.026)	(0.164)
Top-10 brand x 1	0.214*	0.163*	1.559	0.205*	0.294	-0.838*	-5.435	1.356*
	(0.034)	(0.024)	(0.823)	(0.012)	(0.325)	(0.073)	(–)	(0.200)
Top-10 brand x Smartphones	0.073*	0.136*	0.143*	-0.100*	0.341*	0.181*	0.215*	5.111*
	(0.013)	(0.011)	(0.019)	(0.018)	(0.015)	(0.092)	(0.014)	(0.133)

Notes: All estimates are the results of a BLP estimation, using `pyBLP` by Conlon and Gortmaker (2025). Estimates are performed in two stages; the second stage uses first-stage estimates as starting values, but using Chamberlain (1987)-style optimal instruments. Missing values are due to numerical concerns with the estimation, and will be fixed in subsequent iterations of this paper.

Table A.8: BLP coefficient estimates—toilet paper

	Midwest Lakes	Midwest Plains	Mountain West	Northeast	Pacific West	Southeast	Mid-Central South	West-Central South
1 (σ)	0.018	0.015	0.194	0.016	0.024	0.008	0.167	0.005
	(593.532)	(302.538)	(28.461)	(61.684)	(208.612)	(529.401)	(16.487)	(1468.899)
-Price (σ)	0.100	0.003	0.422	1.818*	0.238	0.005	1.928	0.040
	(0.396)	(51.246)	(1.022)	(0.600)	(1.383)	(128.535)	(1.944)	(16.332)
#1 brand (σ)	0.067	0.000	0.161	0.106	0.034	0.073	0.136	0.146
	(18.043)	(15872.607)	(2.166)	(2.428)	(64.048)	(36.970)	(6.708)	(31.395)
#2 brand (σ)	0.070	0.741	0.100	0.388	0.457	0.314	0.902	0.230
	(22.015)	(11.792)	(9.497)	(2.616)	(10.089)	(8.316)	(2.281)	(16.360)
#3 brand (σ)	0.142	0.011	0.283	0.553	0.541	0.212	0.273	0.218
	(10.351)	(13214.818)	(1.200)	(1.609)	(5.686)	(43.628)	(3.628)	(54.173)
Top-10 brand (σ)	0.413	0.335	0.504	0.156	0.197	0.999	0.349	0.877
	(13.552)	(11.833)	(1.697)	(3.459)	(15.777)	(5.564)	(4.002)	(4.124)
1 x Annual income > \$100,000	-0.008	-0.586*	-1.076*	0.010	-1.293*	-0.213	0.050	-0.065
	(0.269)	(0.277)	(0.285)	(0.376)	(0.586)	(0.118)	(0.353)	(0.200)
1 x Annual income \$50-100,000	-0.015	-0.113	-0.812*	-0.044	-0.783	-0.274*	-0.006	-0.241
	(0.184)	(0.189)	(0.171)	(0.332)	(0.414)	(0.103)	(0.340)	(0.124)
1 x Household head > 50	-0.022	-0.485*	-0.235*	-0.173	-0.086	-0.219*	-0.171	-0.252*
	(0.064)	(0.161)	(0.098)	(0.232)	(0.159)	(0.046)	(0.187)	(0.078)
1 x Smartphones	-0.006	2.756	-0.169	-0.225	-0.037	-0.327*	0.400	1.447*
	(0.425)	(2.425)	(0.102)	(1.390)	(0.214)	(0.118)	(0.371)	(0.559)
-Price x 1	-0.196	-1.173*	-1.269	-2.416	-1.612*	-2.882*	-2.495	-1.644*
	(0.336)	(0.420)	(0.717)	(4.895)	(0.733)	(0.220)	(2.516)	(0.664)
-Price x Annual income > \$100,000	-0.044	-1.175*	-0.841*	-1.013	-1.656*	-2.496	-2.211	-3.644
	(0.036)	(0.212)	(0.188)	(1.385)	(0.457)	(2.906)	(2.679)	(4.137)
-Price x Annual income \$50-100,000	-0.033	1.429	-0.651*	-0.863	-0.613*	-2.289	-1.868	-2.700
	(0.031)	(1.441)	(0.110)	(1.330)	(0.306)	(2.000)	(2.408)	(1.725)
-Price x Household head > 50	0.000	-0.703*	-0.109	-0.715	-0.048	-1.500*	-0.963	0.318
	(0.026)	(0.138)	(0.067)	(1.169)	(0.156)	(0.577)	(1.307)	(0.545)
-Price x Smartphones	-0.178*	-0.539*	-0.237*	-0.386	-0.329	0.070	-1.050	-1.634*
	(0.035)	(0.154)	(0.075)	(5.842)	(0.172)	(0.234)	(1.324)	(0.717)
-Price x Household head > 50 x Smartphones	-0.004	0.293*	-0.030	0.336	0.124	1.128*	0.323	-1.140*
	(0.031)	(0.108)	(0.042)	(0.802)	(0.121)	(0.549)	(0.775)	(0.504)
-Price x Annual income > \$100,000 x Smartphones	-0.038	0.320	-0.106	-0.607	-0.874	-1.084	-1.125	-0.893
	(0.029)	(0.243)	(0.083)	(0.633)	(0.691)	(3.983)	(7.681)	(34.798)
-Price x Annual income \$50-100,000 x Smartphones	-0.029	-2.020	0.083	-0.055	-0.280	-0.181	-0.847	-0.567
	(0.034)	(1.358)	(0.054)	(0.847)	(0.178)	(1.980)	(3.439)	(6.152)
#1 brand x 1	-0.067	1.654	0.674*	-0.151	-0.051	1.106	-0.326	-0.243
	(0.419)	(1.391)	(0.118)	(0.153)	(1.030)	(0.908)	(0.488)	(1.662)
#1 brand x Smartphones	0.023	0.095	-0.021	0.170*	0.107*	0.221*	0.076*	0.037
	(0.061)	(0.067)	(0.033)	(0.033)	(0.051)	(0.019)	(0.034)	(0.062)
#2 brand x 1	0.070	0.331	0.742*	0.279	0.164	0.088	0.219	0.012
	(0.589)	(3.693)	(0.365)	(0.550)	(2.023)	(0.868)	(0.680)	(1.433)
#2 brand x Smartphones	0.037	0.050	0.162*	0.369*	0.410*	0.149*	0.329*	0.226*
	(0.055)	(0.069)	(0.028)	(0.058)	(0.036)	(0.022)	(0.134)	(0.071)
#3 brand x 1	0.142	0.171	0.136	0.451	-0.161	0.151	0.118	0.083
	(0.540)	(73.264)	(0.112)	(0.470)	(1.376)	(3.099)	(0.526)	(4.350)
#3 brand x Smartphones	0.071	0.336*	0.429*	0.373*	0.430*	0.391*	0.398*	0.674*
	(0.072)	(0.059)	(0.040)	(0.049)	(0.058)	(0.030)	(0.051)	(0.117)
Top-10 brand x 1	0.413	0.852	0.670	0.137	0.206	0.346	0.291	0.219
	(2.304)	(0.863)	(0.428)	(0.280)	(1.357)	(1.639)	(0.727)	(1.146)
Top-10 brand x Smartphones	0.191*	0.092	-0.074	-0.413*	-0.785*	0.037	-0.106	-0.621*
	(0.069)	(0.159)	(0.052)	(0.102)	(0.095)	(0.030)	(0.055)	(0.177)

Notes: All estimates are the results of a BLP estimation, using `pyBLP` by Conlon and Gortmaker (2025). Estimates are performed in two stages; the second stage uses first-stage estimates as starting values, but using Chamberlain (1987)-style optimal instruments.