

Demand in the Wake of an Emerging Middle Class and Low-End Entry*

Alon Eizenberg[†]
Department of Economics
Hebrew University of Jerusalem

Alberto Salvo[‡]
Kellogg School of Management
Northwestern University

January 13, 2012

Abstract

The “emerging middle class” has become a force of great economic importance in consumer markets around the globe. This paper examines the impact of a substantial rise in Brazil’s living standards on the development of the country’s large soft-drink market, during a period which saw unprecedented growth in the share of generic soda brands. Relying on a unique combination of richly varying market and consumer-level data, we estimate a novel structural demand model that identifies a mechanism by which a household develops either a “premium brand habit” or a “frugal habit.” We find strong empirical evidence of such persistence in preferences, and demonstrate that controlling for such consumption habits is instrumental for identifying the heterogeneous price sensitivities of different socioeconomic groups. Our model casts light on the nature of competition between premium brands and the expanding competitive fringe in an emerging market. In particular, our demand model provides support for Coca Cola Co’s abrupt price cut in July 1999, whereas models that do not incorporate our persistence structure provide weaker support for this strategic move.

*Preliminary results, comments most welcome. For facilitating access to data, we wish to thank Ricardo Fort, Bruno Gouvea, Isadora Nardy, Claudia Pessoa and Daniela Pisetta. We also thank Alaor Dall’Antonia Junior and Maria Cristina Costa at the National Institute of Meteorology (INMET) for access to their data. We are grateful to David Besanko, Allan Collard-Wexler, David Genesove, Scott Stern and participants of the I.O. day at Tel Aviv for valuable comments. All errors are our own.

[†]aloneiz@mscc.huji.ac.il

[‡]a-salvo@kellogg.northwestern.edu

“A study this year by the United Nations Economic Commission for Latin America and the Caribbean concluded that tens of millions of the region’s inhabitants have risen into the middle class over the past two decades. That’s prompted ‘a notable expansion of the consumer market,’ ...(thanks to) the prospects of los emergentes – the emerging ones – as marketers call the newly minted middle-class members”

Matt Moffett, *The Wall Street Journal*, November 15, 2011

“(G)rowth is what (Procter & Gamble CEO) Mr Lafley is so anxious to defend and extend. He likes beauty—and has concentrated P&G’s acquisitions there—because it is a business that is almost entirely branded. There is little competition from retailers’ ‘own labels’—competition that, he says, threatens to push him further down into ‘commodity hell’...”

The Economist, April 12, 2006

1 Introduction

The “emerging middle class” has become a major economic phenomenon in consumer markets around the globe. Since the mid 1990s, many developing countries, as far-flung and varied as Brazil, China, India, Indonesia and Turkey, are experiencing an important socioeconomic transition, whereby a substantial mass of low-income households emerge from below the poverty line and begin to consume goods and services that they previously could not afford.¹ These “new consumers” can be viewed as a crucial engine of growth for the world economy. It is, therefore, important to develop empirical tools to quantify and assess the impact of such demographic shifts on the demand for consumer goods.

This paper examines the impact of a substantial rise in living standards, among a large proportion of Brazilian households, on consumer choice for carbonated soft drinks (or “soda”). The soda category is an important branded food & beverage market, and the Brazilian market trails only the United States and Mexico by volume. Following a successful economic stabilization plan in 1994, aggregate consumption of soda doubled by 1997, and continued to grow at an annual rate of about 10% through 1999. As is well documented, this growth was fueled by pronounced upward mobility among lower income households, who were no longer forced to pay an “inflation

¹The Economist (2011a), citing Martin Ravallion (a poverty expert with the World Bank) as its source, states that using a broad income definition “(t)he middle classes... trebled in number between 1990 and 2005 in developing Asia to 1.5 billion.” To provide another recent point of reference, the Economist (2011b) cites Nomura Bank when stating that by 2014 Indonesia should boast almost 150m “newly affluent Indonesians (who) are certainly spending.”

tax.”² A contemporary article in the Financial Times (1999), reporting on general changes to the economy but focusing on the soft drink industry, spoke of “the increased purchasing power that came with stable prices (that) allowed about 25m new consumers into the (soft drink) market” (among a total population of 170m at the time). Other consumer markets, as varied as refrigerators, fresh meat, housing and cement (a consumer good in Brazil), saw similar boosts to demand.

One might have expected that established soda producers, namely Coca Cola Co. and Ambev (now part of the AB Inbev group), who in 1996 jointly accounted for a dominant 90% of Brazilian soda expenditure, were best positioned to reap most of the fruits of this increased affluence. Yet changes in the balance of power between these established “premium brand” sellers and a fringe of regionally-focused discount brands (which we label “generics”) were as dramatic as the socio-economic shock. Between 1996 and 1999, the combined volume share of (ultimately) hundreds of local generic brands doubled from 20% to 40%.³ In contrast to the premium brands’ premium prices and heavy advertising, generics focused their “merchandising spending” on securing shelf space via low prices. With the stiff competition slowing down company growth, “Coca-Cola blamed difficulties in developing countries such as Brazil when it shocked Wall Street in December (1998) by announcing a rare drop in quarterly sales” (Financial Times 1999).⁴

Casual observation of this and other markets (see below) suggests that the two phenomena just described—a new middle class and an expanding fringe—are interrelated. To stay with the Financial Times’ narrative, “(t)he new customers...had different priorities...(t)hey were less concerned about expensive TV ads and more interested in value.” In other words, the rising share of shoppers with frugal habits provided fertile ground for the growth of producers bent on competing on price.

Having kept prices broadly constant during the preceding years of expansion and entry in the fringe, in 1999 Coca-Cola Co. abruptly cut prices across its brands by over 20%, a move that was soon matched by Ambev. At this point, the relentless process by which the competitive fringe’s share grew at the expense of the premium brands’ share leveled off. Importantly, the fringe’s pricing did not respond (i.e., deviate from a gradual downward trend), consistent with competitive behavior among generic operators—near if not perfect substitutes for one another. Moreover, the fringe was able to weather the storm in the sense that it did not lose market share

²The “unbanked masses,” with no access to inflation-indexed bank accounts, were the main beneficiaries of the taming of chronically high inflation. To cite a recent article, “...Jose Benevenuto, a 53-year-old Rio de Janeiro bus driver ... still recalls the years in the early 1990s when Brazil’s four-digit inflation forced him to rush to the supermarket as soon as he was paid so he could spend his money before it lost all value” (Wall Street Journal 2011). By 1995, inflation was (sustainably) down to single-digit annual levels.

³This pertains to the dominant market segment of family-size bottles sold through the “self-service outlets” distribution channel (supermarkets with checkouts) in urban areas on which our empirical analysis focuses. Such outlets accounted for over 70% of soda volume purchased for consumption inside the urban home, of which family-size (rather than single-serve) packaging accounted for over 80%.

⁴Of note, Coca-Cola was the quintessential premium brand seller. For example, Interbrand/BusinessWeek’s annual rankings have consistently featured the Coke brand among the world’s most (if not *the* most) valuable. Ambev, who ran the Pepsi franchise in Brazil, was also a heavy advertiser.

to premium producers.

One might speculate whether Brazil’s premium soda sellers were verging on what former Procter & Gamble (another premium brand seller) CEO Alan Lafley has dubbed “commodity hell,” a situation in which a large mass of consumers are not willing to pay premium prices for a highly-advertised, high-margin product. Instead, these consumers adopt the frugal habit of choosing value products—whether retailers’ “private labels” in developed markets or low-end entrants in developing ones—over premium brands. In our setting, an emerging middle class may well have been tilting the balance toward frugal consumption patterns, prompting the established duopoly to defensively cut prices.

In this paper we provide an empirical study of the relationship between Brazil’s sweeping socioeconomic changes and the growing competitive fringe of soft drink producers during a period that spans over six years (December 1996 through March 2003). We extend the literature on estimating differentiated-product demand systems by developing and implementing a model that is well-suited for such fast-changing markets and can be estimated off (predominantly) aggregate (market-level) data. Building on the random-coefficient logit framework, our model displays two novel features.

First, motivated by socioeconomic data used widely by marketing practitioners, consumers belong to either one of three discrete demographic groups: “Established Affluent” (EA), “Poor” (P) or “Newly Affluent” (NA). The EA group corresponds to households who were already “affluent” (according to widely-accepted definitions, as explained in Section 2 below) before the process of upward mobility began, while the NA group represents the emerging middle class. Whereas the EA population is fixed, households in the other two groups may either: (i) “move up” from poverty to newly affluent status, or (ii) fall from grace, returning to poverty after having joined the ranks of the newly affluent. Such downward mobility is apparent toward the end of our sample period, when a fraction of Brazil’s new middle class was hit by a recession following the 1999 Brazilian crisis and the 2000-01 Argentine crisis next door. These aggregate trends—*upward* and *downward mobility*—are incorporated into the estimation of our model. Further, our model also allows for *urbanization*, another important demographic shift.

The second key component of our model is habit formation, of a special kind. In particular, a soda-consuming household develops a habit to consume either premium brands or generic products based on the consumption choice it made in the immediately preceding period.⁵ We refer to this state-dependent consumer attribute as a “premium habit” versus a “frugal habit.” We label this a *Brand Type Persistence* (BTP) model, as it captures a persistence in demand for

⁵As explained below, much of the extant literature on state-dependent preferences assumes that current “habits” were developed in the immediately preceding period. In most of these papers, periods are captured as shopping trips in household-level retail scanner data. In contrast, a period in our setting lasts one or two months, an interval we view as more appropriate for our habit formation model. A more general model would allow habits to evolve as a function of consumption choices over multiple preceding periods.

a certain type of good (i.e., premium or generic).

In our BTP model, developing a premium habit in period $t - 1$ (by consuming, say, Coke) increases the utility from consuming any premium brand (say, Coke, Fanta or Pepsi) in period t . Similarly, recent consumption of a generic brand provides a boost to the current utility from consuming any generic good. Notice that by defining habit as being developed with respect to a class of products (premium vs. generic) as opposed to particular products distinguishes our model from brand loyalty models. Our model captures P&G CEO Lafley’s notion of “commodity hell” as a situation in which a substantial mass of households develop the frugal habit.⁶

A household’s type in our model is determined, therefore, by both its current socioeconomic standing (EA, NA or P) and the current habit (Premium, Frugal or No Soda), formed in the previous period. Crossing these definitions together yields nine household types. We track the population fractions belonging to each of these types over time by combining data on demographic trends with brand choices predicted from our utility framework given any guess of the model’s parameters. As explained below, we develop an estimation procedure that, building on the GMM framework familiar from papers estimating differentiated-product demand systems, incorporates this dynamic updating process into each evaluation of the objective function.

Findings. Our results demonstrate that newly affluent households are significantly less price sensitive than the poor, but are more price sensitive than the established affluents. This ordering is consistent with the Financial Times article cited above that depicts the new customers as focused on value rather than on expensive brands. In addition, our results also provide strong evidence for the persistence mechanism, which is shown to be of both statistical and economic significance.

We employ our structural model in a counterfactual analysis that provides insights into Coca Cola’s drastic price cut in July 1999. Our results demonstrate that, had premium brands failed to cut prices in 1999, they would have seen their market shares suffer substantial declines through 2003. Our BTP model, therefore, provides a strong justification for this strategic move.

One might view this result as trivial, since any reasonable demand model would predict that the failure of a premium brand to cut prices would lead to a drop in market share in an environment characterized by: (i) the rapid growth of a price-sensitive consumer segment, (ii) intensifying competition from generic competitors, and (iii) limited ability to price discriminate. Our analysis, however, demonstrates that it is the key component of *persistence* that we capture in our model that motivates a price cut as deep as that performed by Coca Cola.

The reason that accounting for persistence is a crucial component in analyzing Coca Cola’s strategic considerations is that, once a new middle class develops a habit of choosing cheap

⁶Having convinced a substantial share of consumers in mature markets to pay margins of 50% (say), heavy advertisers like Coca-Cola and Nestlé may “know” how to compete against similarly premium brand sellers such as Pepsico and Unilever. In contrast, competing against “value brand” sellers who challenge their very business model may fall outside their comfort zone.

substitutes over premium brands, it becomes more difficult to convince these consumers to pay as much as 100% more for the premium product. By cutting prices, Coca Cola ensured that a larger share of newly-affluent consumers would buy its premium products and develop a “premium habit” that would help defend premium market shares into the future. Our BTP mechanism, capturing Mr. Lafley’s notion of “commodity hell,” therefore provides a very strong incentive for a premium price cut.

To investigate the significance of the BTP mechanism, we estimate two model variants, each of which considers a variation in the persistence component. The first model variant allows for *Soda Category Persistence*—call this the SCP model. Here, choosing *any* soda in period t (i.e., either premium or generic) over the outside option increases the period- $(t + 1)$ utility associated with *any* soda product. The second model variant, which we refer to as the *No Persistence* (NP) model, shuts down the persistence in preferences altogether.⁷

Comparing estimates of our two model variants to those of the baseline BTP model reveals that the BTP model does a better job in capturing the heterogeneous price sensitivities of the established affluents, newly affluents and poor. As a consequence, it is the BTP model that delivers the compelling prediction according to which the “premium-to-generic” consumption ratio increases with socioeconomic standing: the BTP model predicts these ratios to be 2.1, 1.1 and 0.6 for the EA, NA and P groups, respectively. The SCP and NP models do an inferior job in generating this pattern (e.g., the NP model actually predicts that the poor’s premium-to-generic ratio is higher than that of the newly affluent).⁸

We also employ the SCP and NP variant models in the same counterfactual analysis described above, which asks what would have happened had premium brands failed to cut prices. While these models also show that premium brands would have lost market share, the effect obtained from the SCP and NP models is substantially smaller compared to that obtained from the BTP model. This illustrates our claim, stated above, that incorporating brand-type persistence into the demand model is important for understanding Coca Cola’s strategic considerations in its competition with the generic fringe.

Identification. Our demand model faces a familiar challenge from the literature on estimating models of habit formation: how can one separately identify consumer heterogeneity from persistence? As we explain in detail in the sections below, our model is identified off of rich cross-sectional and time-series variation in the data, including region-specific social mobility and pricing variation that is likely exogenous to demand *unobservables*. For example, following Salvo (2009), we argue that the magnitude and abruptness of Coca-Cola’s nationwide price cut

⁷Notice that the BTP model and its SCP variant are non-nested, while the NP model is nested in both the BTP and SCP models.

⁸It is also worth noting that the BTP model predicts “inside shares” (i.e., fraction of households buying soda) of 54%, 25% and 7% for the established affluents, newly affluents and poor, respectively. This appears consistent with Argentine economist Nuria Sussmel’s statement to the Wall Street Journal (2011) that “(t)he middle class has the tastes of the wealthy and the salary of the poor.”

halfway into the sample period strongly suggests that it was unlikely to be correlated with any region-specific shocks to demand, making the price cut itself an effective instrument.

Our model implies that price sensitivity varies at the socioeconomic level, allowing this to be identified off of the observed co-variation of the socioeconomic shifts (i.e., upward- and downward-mobility) and the market shares. Identification of the persistence mechanism also follows from data variation: for instance, during the recession that sets in toward the end of our sample period, households fall back from newly-affluent status to the ranks of the poor- but even though prices remain quite stable, soda consumption does not fall.

This latter point is related to another important insight: failing to control for persistence may frustrate the identification of heterogeneous price sensitivities. In the absence of persistence, the model would have to interpret the events of the recession (i.e., the high soda consumption of an increasingly poor population) as evidence that the poor are “*not that price sensitive.*” Indeed, as alluded to above, our NP model does not do a good job in separating out the price sensitivity of the poor from that of the newly affluents.

In Sections 3 and 4 below we provide more detailed arguments for the manner in which persistence manifests itself in our data, allowing us to identify it using observed data variation and the restrictions of the model.

Literature. Our study contributes to different lines of research. Studies that introduce habits or persistence into models of consumer behavior have been explored in the empirical economics and marketing literatures. Eichenbaum, Hansen and Singleton (1988), for instance, use US macroeconomic data to estimate a version of their representative consumer model that allows for non-time-separable consumer preferences (admitting habit formation as a special case). The marketing literature offers an extensive body of work on state-dependent preferences and brand loyalty.⁹ Other papers have examined competition between branded goods and lower cost generics, such as in pharmaceuticals, e.g., Chaudhuri, Goldberg and Jia (2006) in India, Hurwitz and Caves (1988) and Scott Morton (2000) in the United States.

Our work differs from the marketing literature in at least three dimensions. First, these studies tend to rely on micro-level panel data, repeatedly observing individual households’ purchasing behavior (e.g., scanner data). By contrast, our paper demonstrates how a model with persistent preferences can be estimated with a panel of market-level data (in addition to a single cross-section of household-level data). We believe that the methodology developed here is more applicable to emerging markets, where repeated observations on a fixed panel of households may not be available, and—even if they are available—such data are likely to miss the demographic shift that lies at the heart of the study.

Second, the state dependence we model differs from brand loyalty. In our context, its source is

⁹This literature dates back at least to Guadagni & Little (1983). Other selected examples are Erdem (1996), Keane (1997), Shum (2004) and Dubé et al (2008). We thank Maria Ana Vitorino for discussing this literature with us.

the “heterogeneity of business formats/models” that is encountered in several emerging consumer markets today, i.e., the premium vs. generic dichotomy. In emphasizing this aspect of the market, our goal is to capture a potentially important mechanism in emerging markets, rather than to extend the literature on estimating brand loyalty effects.¹⁰

Third, the extant literature on brand loyalty often examines mature markets such as US orange juice or margarine (e.g., Dubé, Hitsch, Rossi and Vitorino 2008). In contrast, our emerging market setting enables us to rely on very rich data variation—in terms of both the composition of demand and the pricing behavior of firms. We argue that this setting is particularly valuable for the separate identification of state dependence from heterogeneous price sensitivities.

Finally, this paper wishes to contribute to our better understanding of agents in emerging market settings, such as the heterogeneity of consumers. Another recent example is Sancheti and Sudhir (2009), who examine the consumption of education in India. The growth such markets have been experiencing in recent years makes us view such settings as exciting areas for applied microeconomic research.

The rest of the paper is organized as follows. Section 2 describes the data and the joint phenomena that motivate our analysis: upward social mobility and the growth of the generic fringe. Section 3 develops our demand model, explaining our estimation strategy. Section 4 discusses estimation results, including the counterfactual analysis that sheds light on the nature of competition between premium and generic brands. Intuitive arguments for identification are provided in Sections 3 and 4. Section 5 concludes.

2 Market and data

2.1 Data sources and observed variables

This study brings together three main datasets and incorporates consumer-level as well as market-level data. We begin by briefly describing, in this subsection, the variables we observe from each of the three datasets (and provide further details in a data appendix). Then, in Subsection 2.2, we describe how these variables shed light on two striking features of the data, which we refer to as *the Rise of Newly Affluent Households*, and *the Growth of the Competitive Fringe*.

Data source 1: Product market data (Nielsen). We observe a panel consisting of total quantities and prices for individual established soft-drink brands and an aggregation of generic brands in the Brazilian market. There are $G = 7$ regions and $T = 57$ time periods, running from

¹⁰Executives of a global “fast-moving consumer goods” firm meeting one of us recently in Delhi stated that “as a company in the *A business* we don’t naturally understand the *B business*, where the *value* proposition is at the heart of it, putting us at a certain disadvantage when selling to the *Bottom of the Pyramid* in the Indian market.” (To be clear, all words—including the terms in italics—are the executives’ own, though in slightly rearranged order without modifying context.) Of interest, the China-based appliance manufacturer Galanz cites the National Bureau of Statistics in claiming that there were “nearly 300 brands in (the) Chinese market” in 2008 (Galanz 2008). Abbott India’s brands Digene, Eptoin and Cremaffin face competition from 211, 327 and 242 “regional” generics, respectively (as shared by the company during a corporate presentation in late 2011).

the December 1996-January 1997 bimonth to the March 2003 month (in 2000 Nielsen raised the frequency of its point-of-sale audits from a bimonthly to a monthly basis), such that there are $G \cdot T = 399$ region-period markets. The geographic markets are urban and, as in Salvo (2009), we consider soft drinks sold through the predominant self-service channel in the predominant (2-liter) family-size bottle.¹¹

Also following Salvo (2009), we aggregate flavors and brands into $J = 9$ brand-groups, indexed by $j \in \mathcal{J}$ where $J = |\mathcal{J}|$. These include five brands of the Coca-Cola Company (Coke, Fanta, Kuat - a guaraná-flavored soda, Diet Coke and “Other Coca-Cola”), three brands marketed by Ambev (Guaraná Antarctica, Pepsi and “Other Ambev), and a ninth group of discount brands which Nielsen aggregates and we refer to as generics or “B brands.” This latter group will be a focus of interest for us as it covers a fringe of small, regionally-focused operators who were able to grow substantially at the expense of the established firms’ brands, as well as the outside good, by virtue of fierce price competition. We label the eight brands owned by Coca-Cola or Ambev as “premium brands,” or “A brands” $j \in \mathcal{A}$, whereas the group of (numerous) B brands is the only element of the set of generics \mathcal{B} .

Much of our interest lies in understanding the process by which both premium and generic brands attempt to tap into the growth potential of the market. We observe the quantity and price, denoted by q_{jgt} and p_{jgt} respectively, associated with brand $j \in \mathcal{J}$ sold in the region-period market gt . To compute market share s_{jgt} , defined as the share of market gt ’s population that consumes brand j , we need to define the size of the potential market. Our base specification defines the “size of market gt ,” denoted \mathcal{S}_{gt} , as six liters per week over the duration of period t multiplied by the number of urban households residing in region g (which we obtain from a fourth data source). One can interpret the six liters per week as three weekly family meals in which a 2-liter family-size bottle of soda might be brought to the table (rather than water, juice, etc.) One informal way to think of this market size is “the maximal amount of carbonated soft drinks that would be sold on this market under any reasonable situation.” We then compute $s_{jgt} = q_{jgt}/\mathcal{S}_{gt}$. Finally, we define the share of the “outside option” (i.e., the population share not consuming soda) as $s_{0gt} = 1 - \sum_{j \in \mathcal{J}} s_{jgt}$.¹²

Data source 2: Socioeconomic composition of Brazil’s urban households (IBOPE).

To track the undercurrent of social mobility in the Brazilian economy, we rely on the proprietary LatinPanel survey from IBOPE, a leading (private-sector) provider of data on consumer demographics.¹³ The survey, widely used by marketing practitioners, profiles urban households

¹¹Brazil is substantially more urbanized than, say, China or India, and its urban households (over four-fifths of total) buy their groceries mostly at supermarkets with checkouts, rather than the traditional “behind-the-counter” stores that are popular in rural areas. That is, “modern” or “organized” retail is prevalent in Brazil’s urban areas. Family-size bottles (also known as “multi-serve” packaging) dominate “single-serve” (300ml) bottles or cans which are sold mostly through the bars and restaurants channel.

¹²We are in the process of exploring alternative assumptions on the market size.

¹³The company’s name is so established among Brazilian households that, as cited in Wikipedia, it is synonymous with research (e.g., see the *Aurélio* Portuguese language dictionary). Coca-Cola Co. kindly shared the data with us for the purpose of this study.

in Brazil’s different regions based on their expenditure on durable goods and services (e.g., ownership of a refrigerator, number of TVs and bathrooms in a residence, current employment of house maids, education attainment). Adopting an industrywide points scale (ABEP 2003), each household is assigned to a “socioeconomic segment.”¹⁴ The IBOPE data that we have access to covers the period 1994-2006 (with 1995 missing) and provides the proportion of urban households that belong in either the “AB,” “C” or “DE” groups (respectively with “high,” “intermediate” or “low” levels of affluence) in each of seven (comprehensive) geographic regions.

The IBOPE data indicate that the demographic composition of urban households: (i) was stable between 1994 and 1996; (ii) was characterized by the proportions of “AB” and “C” households sharply gaining ground between 1996 and 2000; and (iii) this upwardly mobile change in shares was partially reversed thereafter, consistent with a recession setting in at that time. In aggregate, the proportion of “DE” households fell from 50% in 1996 to 33% in 2000, then rose to 44% by 2003 (and, conversely, the proportion of “AB” households rose from 19% in 1996 to 33% in 2000, then fell to 23% by 2003).

An important (secular) demographic trend during the sample period was migration of households from rural to urban regions. Since these migrants tend to be relatively poor, this shift operates in the direction of increasing the share of “DE” urban households reported by IBOPE (as explained below, some of the subsequent reversion of the upward mobility observed in the IBOPE data can be attributed to the rapidly urbanizing, relatively underdeveloped North/Northeast region).

In sum, the striking feature emerging from the IBOPE data is the pronounced upward mobility of Brazilian households during 1997-2000. We provide further insight into this key phenomenon, and lay out how we model it, in Subsection 2.2 below.

Data source 3: Soda expenditure by different socioeconomic group (IBGE). Our third main data source allows us to relate household characteristics to soda consumption choices at the beginning of our period of study. We use an urban household expenditure survey that was conducted between October 1995 and September 1996 by IBGE, a federal agency equivalent to the US Census Bureau and Bureau of Labor Statistics combined. This survey (hereafter HEX 95/96) reports the type of soda brand purchased (as well as the amount spent) for consumption inside the home. Since households in the survey are not classified according to the ABCDE system, we use the detailed “balance sheet” data (e.g., ownership of a refrigerator, number of TVs and bathrooms in the residence, current employment of house maids, education attainment) to assign, like IBOPE does, the household to an (IBOPE) socioeconomic segment.¹⁵

¹⁴The points scale used to classify each household stays clear of income, there being reasons why income-based measures might less accurately reflect changes in the standard of living (Carvalho Filho and Chamon 2011, Economist 2007). That said, to provide perspective, mean annual incomes in 2000 for “C” and “DE” urban households were respectively US\$ 6,100 and US\$ 2,600 (ABEP 2003, based on an IBOPE survey, using nominal 2000 R\$/US\$).

¹⁵For perspective on household size, “ABC” and “DE” households average 3.64 (std. dev. 1.58) and 3.76 (s.d. 1.99), respectively.

We then cross the household’s socioeconomic group (**ABC** or **DE**) with their spending on soda for inside-the-home consumption, namely: whether the household purchased soda during the week of reference, and, if so, which brand was purchased. Table 1 reports the extensive margin on purchases for inside-the-home consumption of soft drinks by socioeconomic group in the 10 cities covered by the HEX data (see the appendix for details on the mapping of these 10 cities to the seven Nielsen regions). Thus, for example, 34.5% of the metropolitan area of São Paulo’s (region 4) **ABC** households in 1996 purchased soda for home consumption while only 19.8% of **DE** households did so (see the appendix for evidence and a discussion of the “intensive margin” of household soda consumption).

We also compute, for each of the $g = 1, \dots, 7$ regions, the share of **ABC** households who consume a premium brand, the share of **ABC** households who consume a generic soda product, and similar figures for **DE** households. These reflect the covariation between a household’s socioeconomic standing and its consumption choices, which will help identify heterogeneous price sensitivity effects for different consumer segments.¹⁶ Our subsequent modeling of households at each point in time as either premium or frugal (or outside good) shoppers in the soda category, but not “hybrids,” is consistent with the HEX data. As explained below, the fact that the HEX survey was conducted shortly before the beginning of our Nielsen market data allows us to use this information as an “initial condition” for the evolving relationship between socioeconomic standing and consumption choices.

Additional data sources. As detailed in the appendix, our analysis draws on further data sources. In particular, we use: (i) the population of urban households by region from IBGE’s expanded annual household demographics surveys (PNAD), (ii) proprietary McCann-Erickson data on advertising intensity at the brand-market level, (iii) proprietary temperature data (a natural demand shifter) from the National Institute of Meteorology, and (iv) data on cost shifters such as the prices for sugar, plastic and electricity.

2.2 Market dynamics: the emerging middle class and the growing fringe

2.2.1 The Rise of Newly Affluent Households

As described above, the studied period is characterized by a potent process of upward mobility, by which households “move up” from **DE** to **ABC** status, with gains being partially reversed toward the later part of the sample. What likely explains these demographic shocks?

Economic reforms in the early 1990s, including trade liberalization and, most notably, the *Real* stabilization plan of 1994 which brought very high and chronic inflation under control, were

¹⁶As Table 1 indicates, say for São Paulo Metro, among every 1000 **ABC** households in 1996, 331 purchased premium brands and 14 purchased generic soda; comparable figures for every 1000 **DE** households are 173 premium-soda buyers and 26 generic-soda buyers. Across all cities, **DE** soda-consuming households were more likely to purchase generics over premium brands (12 : 163) relative to **ABC** soda-consuming households (11 : 339).

followed by strong consumption growth across the Brazilian economy, particularly among lower-income households. To illustrate, Figure 1 reports per capita consumption between the mid 1980s and mid 2000s in two different sectors—beverages (soft drinks) and housing (cement); a similar temporal pattern leading up to 2000 is present.¹⁷ The Boston Consulting Group (2002), reporting on its own household survey, spoke of the emergence of a middle class with “very strong consumer potential,” whereas Fátima Merlin, chief economist for the Brazilian Association of Supermarkets (ABRAS), referencing the same IBOPE data that we use, stated that “following the *Real* Plan, thanks to price stability and real growth in workers’ earnings, consumer markets experienced entry by households previously outside such markets, *with upward migration from the ‘E’ and ‘D’ segments of the population to the ‘C’ segment, as the IBOPE data indicate*” (SuperHiper 2003; emphasis added). Also recall, in this context, the above-cited contemporaneous 1999 Financial Times article which referenced Brazil’s “25m new consumers (in the aftermath of an) economic plan in 1994.”

As for the partial reversion of upward mobility reported by IBOPE over 2001-2003, economic episodes that may have dampened investor and consumer sentiment include the 1997-98 Asian crisis, the 1999 Brazilian crisis (with the Brazilian Real, denoted R\$, initially devaluing 70% against the US Dollar) and the 2000-01 Argentine crisis.¹⁸ One should also note that continued urbanization, most marked in the relatively less urbanized North/Northeast region, implies that, all else equal, the proportion of **DE** households was bound to grow over time, as rural migrants arriving at the city are typically the least affluent.¹⁹

We seek to incorporate these demographic shifts, as captured in the IBOPE data described above, into a consumer choice model. To this end, we define three socioeconomic groups: “Established Affluent” (EA), “Newly Affluent” (NA) and “Poor” (P).

The “Established Affluent” group consists of urban households who were already in **ABC** status in 1996, i.e., before the process of upward mobility captured in our data takes off (see Subsection 2.1 above). This group’s size (i.e., the number of such households) is thus fixed, in each of the seven regions, at the initial level of **ABC** households in that region, computed using the IBOPE data.²⁰ Denote the number of Established Affluent households in the region-period market gt by $M_{EA,gt}$. To emphasize, each region’s $M_{EA,gt}$ series is held constant over all sample periods $t = 1, \dots, 57$. We define the size of the “Poor” group in each region by the

¹⁷See Carvalho Filho and Chamon (2011) and Salvo (2009, 2010) for further discussion of the consumption effects of reforms in the 1990s. See also Neri (1995).

¹⁸A similar temporal pattern of prosperity can be detected in earnings data in IBGE’s monthly survey of earnings and employment, conducted in 6 large cities, though the turning points in the series tend to occur sooner than 2000. Details are available from the authors upon request.

¹⁹In the Northeast region (Nielsen region 1), 66% of the population in 1994 lived in urban areas, rising to 73% by 2006. This compares to a higher (and slower-changing) urban share for the other six regions of 83% in 1994 and 87% in 2006. Over this period, the Northeast gained 3.0 million urban households (5.6m to 8.6m) and all other regions combined gained 10.4m (22.3m to 32.7m).

²⁰See the appendix for further details, including how we interact IBOPE’s urban socioeconomic distributions (proportions) with the number of urban households from IBGE’s annual household demographics surveys, as well as consistency checks between IBOPE and IBGE survey data.

contemporaneous level of urban households who (based on IBOPE) belong to socioeconomic segment **DE**, denoting this measure by $M_{P,gt}$.

Finally, we define the size of the “Newly Affluent” group in market gt as the difference between the contemporaneous number of **ABC** households (momentarily denote this by $M_{ABC,gt}$) and the region’s initial (i.e., 1996) number of **ABC** households (which fixed the value of $M_{EA,gt}$ for all $t = 1, \dots, 57$ as explained above). In the data, this difference $M_{ABC,gt} - M_{EA,gt}$ is strictly positive for all markets $t > 1$ and zero for $t = 1$; denote this mass of newly affluent households by $M_{NA,gt}$. To be clear, any number of households who are in **ABC** status in period $t > 1$ in excess of the number of **ABC** household at the beginning of the studied period are considered time- t newly affluent households.^{21 22}

One should note that the zero (by definition) number of newly affluents in period 1 is justified by the fact that, in the IBOPE data, the process of upward mobility takes off just before our Nielsen sample of the soda market begins (late 1996). This modeling assumption is also consistent with press and trade articles from the time. For example, our $M_{NA,gt}$ measure for mid 1999, summed across 7 Nielsen regions, translates into 20m consumers, a notch below the Financial Times’ (June 1999) count of “(Brazil’s) 25m new consumers,” noting that Nielsen does not cover the northern states and that our study relates to urban areas.

Figure 2 plots the evolution of socioeconomic composition by region, i.e., the population fractions of established affluent, newly affluent and poor households. The figure clearly demonstrates the emergence of a new middle class. The increase, toward the end of the sample period, in the fraction of the poor at the expense of the fraction of newly affluents reflects the joint effects of the recession and the urbanization process as discussed above. There are large regional disparities, with region 1 (states in the Northeast) being the most poor (65% of its urban households are initially poor) and region 4 (São Paulo Metro) being the most affluent (36% of urban households are initially poor).

2.2.2 The Growth of the Competitive Fringe

Brazil’s soft drink industry was characterized by two “dominant firms”—the Coca-Cola Company and Ambev—facing a competitive fringe (Ambev was formed through the merger of Antarctica and Brahma in 2000). At the start of our sample, Coca-Cola Co’s flagship brand, Coke, commanded a 33% Nielsen volume share of the soda category across all 7 regions, and an additional 3% with Diet Coke. At that time, Ambev’s leading brands were Guaraná Antarctica and Pepsi,

²¹As noted above, in addition to upward and downward mobility, another demographic force affecting $M_{EA,gt}$, $M_{NA,gt}$, $M_{P,gt}$ is (net) rural-to-urban migration. We capture this process via changes in the size of urban population numbers, assuming that households migrating to the city are initially “Poor.” (See Assumption 1 below.)

²²To illustrate our computations by way of an example from the IBOPE data, in the South region there were: (i) in $t = 1$ (Dec-96/Jan-97), 3149 (thousand urban) **ABC** households and 2116 **DE** households, and (ii) in $t = 2$ (Feb/Mar-97), 3238 **ABC** households and 2045 **DE** households. Thus $M_{\text{South},1} = (M_{EA,\text{South},1}, M_{NA,\text{South},1}, M_{P,\text{South},1}) = (3149, 0, 2116)$ and $M_{\text{South},2} = (3149, 89, 2045)$.

each with an 8% Nielsen volume share (recall that we examine sales of family-size bottles through the self-service channel). Both firms heavily advertised their brands, and distributed them nationally (Salvo 2009 describes the industry in further detail, including vertical relations.)

In contrast to the established duopoly, fringe players ran small-scale operations, in most cases individually covering a fraction of a state, and focusing their “merchandising” efforts on securing shelf space via low prices, as they were quite readily substitutable by retailers. Having hovered around a 15% volume share of the soda category at least since 1980 (Salvo 2009), the fringe began growing strongly in the mid 1990s, in the wake of the country’s stabilization plan and the ensuing demographic shift.

Trade accounts report that a shift in distribution technology in the early 1990s also lowered barriers to entry, with the inexpensive non-returnable 2-liter PET bottle replacing the returnable proprietary glass bottle (this was returned to the bottler for reuse, requiring a certain level of sophistication and scale). No census of fringe operators exists, but industry sources suggest that following three years of substantial entry, the number of firms selling generic soda may have reached 500 by 1999.²³

This growth in the competitive fringe is depicted by Figure 3 which, for illustration purposes only, aggregates activity across the seven regions and groups the eight premium brands together. The upper left panel reports total monthly quantities sold, in million liters (with consumption peaking in the summer), and the upper right panel shows (mean share-weighted) prices for premium brands ($j \in \mathcal{A}$) and generics ($j \in \mathcal{B}$), in R\$ per liter. Throughout the paper, R\$ prices are reported at constant Brazil CPI March 2003 terms. (A useful rule-of-thumb to convert R\$ values into rough current US\$ values is to divide the former by 2.)

Over the first 30 months of the sample, from December 1996 to May 1999, the growth of the fringe markedly outpaces that of their established rivals. While premium brands held prices broadly flat, at R\$ 1.15 (per liter), prices in the fringe declined gradually but relentlessly, from R\$ 0.90 to R\$ 0.60.²⁴ Falling generic prices are consistent with substantial entry and capacity expansion in the fringe, as competitive firms passed efficiency gains through to consumers.

In mid 1999, the Coca Cola Co. cut its prices by more than 20%, a move that was soon matched by Ambev. Fringe prices did not respond in the sense that they did not deviate from their trend, consistent with competitive behavior. Concomitant with this abrupt premium price cut, the generic share stopped growing, having *doubled* in the preceding 30 months, as indicated

²³The Financial Times (1999), to further quote their narrative, speaks of “over 900 companies in the industry” thanks to “(t)he introduction of plastic bottles (that) also made lower-cost production easier” (in addition to the “new customers”). Ambev’s annual report in 2003 to Brazil’s version of the US Securities and Exchange Commission refers to “700 ‘low-price brands’ (known as *tubainas*), which are sold at large discounts relative to premium brands... in part due to low distribution costs from disposable PET bottles. Non-returnable packaging accounted for about 12% of the country’s soft drink sales in 1993, having risen fast to 90% by 2003” (Ambev 2003, figures refer to all channels including bars where single-serve returnable glass bottles persist).

²⁴Given that we convert prices to constant R\$ (using the ‘CPI-br’ published by the Fundação Getúlio Vargas) what this means in practice is that nominal prices in the fringe fell 24% compared with the overall price level in the economy (the CPI) growing by 14% over these 30 months, i.e., $.6/.9 \simeq (1 - .24) / (1 + .14)$.

in the two lower panels of Figure 3—the left panel reports volume shares of the soda category (with premium and generic shares summing to one), whereas the right panel corresponds to our construction of the “share of the family dinner table” (e.g., $\frac{1}{G} \sum_g \sum_{j \in \mathcal{A}} s_{jgt}$ for premium).

3 The model

In this section, we develop a model of consumer demand that allows us to study the intertwined processes of upward socioeconomic mobility and growth of generic brands. A salient feature of our approach is that consumer utility, and hence demand, exhibits a persistence attribute with respect to type of brand consumed (“premium” versus “generic”). In Section 4 we will find that controlling for this feature is essential to eliciting heterogeneous price sensitivities that accord with intuition, specifically that newly affluent households display an intermediate level of sensitivity when compared with their established affluent and poor counterparts.

3.1 Model setup

In every period t , a household belongs to one of three socioeconomic groups (EA, NA, P) (again, established affluent, newly affluent or poor) and, consistent with the data, we allow for mobility across these groups over time. The J soft drink brands are vertically classified as either premium brands (referred to for brevity as “A” brands) or generics (“B” brands). In practice, as explained above, we have eight A brands and a single aggregate of generics, classified as a B brand.

A household’s current preferences over substitute soft drink brands $j = 1, \dots, J$ depend on its current socioeconomic standing and, as we motivate below, also depend on the household’s previous-period consumption by virtue of a persistence mechanism. Specifically, we differentiate households who, in the preceding period, consumed: (i) a premium A brand, (ii) a generic B brand, or (iii) did not consume soda at all (label the outside option set \mathcal{O}). Combining the three socioeconomic groups and the three types of recent brand history, there are nine household types, which we denote by:

$$r \in \{EA^A, EA^B, EA^{\mathcal{O}}, NA^A, NA^B, NA^{\mathcal{O}}, P^A, P^B, P^{\mathcal{O}}\}.$$

Thus, for example, a current-period household who is newly affluent and chose generics in the previous period is of type $r = NA^B$, whereas an established affluent household who consumed a premium brand in the preceding period is of type $r = EA^A$.

We denote the fractions of the nine types out of (region-period) market gt ’s household population by:

$$\mathcal{F}_{gt} = \{F_{EA^A,gt}, F_{EA^B,gt}, F_{EA^{\mathcal{O}},gt}, F_{NA^A,gt}, F_{NA^B,gt}, F_{NA^{\mathcal{O}},gt}, F_{P^A,gt}, F_{P^B,gt}, F_{P^{\mathcal{O}},gt}\},$$

such that $\sum_r F_{r,gt} = 1$, and we refer to \mathcal{F}_{gt} as the *type-distribution vector* for market gt .

The indirect utility of household i of type r in market gt from consuming brand j is given by:

$$u_{i \in r,j,gt} = \delta_{jgt} + \alpha_r \cdot p_{jgt} + \lambda \cdot 1\{persistence\}_{jr} + \epsilon_{ijgt}, \quad (1)$$

where the first term δ_{jgt} denotes the (market-specific) household-invariant “base utility” from brand j :²⁵

$$\delta_{jgt} = x'_{jgt}\beta + \alpha \cdot p_{jgt} + \xi_{jgt},$$

with: (i) vector x_{jgt} containing elements capturing non-price characteristics such as brand-region fixed effects, region-specific time trends, seasonal effects, brand advertising covariates and market temperature, a natural demand shifter for soft drinks; (ii) price p_{jgt} as defined above; (iii) ξ_{jgt} reflecting a (brand-market specific) utility shock that is observed by firms and consumers, but unobserved to the econometrician; and (iv) (α, β) are coefficients to be estimated. Notice that α is a common, “base” price sensitivity parameter, and that we next allow different socioeconomic groups to differ in price sensitivity.

The second and third terms in (1) introduce household-type heterogeneity (in that they depend on the type r) as follows. The parameter α_r is a household-type specific shifter of the base price sensitivity α , which in practice we define as follows:

$$\alpha_r := \begin{cases} \alpha_{EA} & \text{if } r \in \{R^A, R^B, R^O\} \\ \alpha_{NA} & \text{if } r \in \{NA^A, NA^B, NA^O\} \\ 0 & \text{otherwise} \end{cases}$$

This implies that while α is the price sensitivity of poor households, the sums $(\alpha + \alpha_{EA})$ and $(\alpha + \alpha_{NA})$ are the price sensitivities of the established affluent and the newly affluent, respectively. We expect $\alpha < 0$, $\alpha_{EA} > \alpha_{NA} > 0$ and $(\alpha + \alpha_{EA}) < 0$ (which jointly imply that $(\alpha + \alpha_{NA}) < 0$).

The variable $1\{persistence\}_{jr}$ in (1) is defined as

$$1\{persistence\}_{jr} := \begin{cases} 1 & \text{if } r \in \{EA^A, NA^A, P^A\} \text{ and } j \in \mathcal{A} \\ 1 & \text{if } r \in \{EA^B, NA^B, P^B\} \text{ and } j \in \mathcal{B} \\ 0 & \text{otherwise} \end{cases} \quad (2)$$

The intuition underlying this term in the utility function is that households’ affinity toward a current brand choice are affected by fairly recent consumption patterns. The persistence we capture here is the tendency, beyond that captured by one’s price sensitivity α_r , to consume a given type of brand: one either has a “habit” of consuming premium soda or a habit of consuming

²⁵We avoid the familiar terminology “mean utility” since, with discrete types, this interpretation is not true.

generic soda. Specifically, consuming *any A brand* $j \in \mathcal{A}$ in the previous period increases one’s utility from consuming *any A brand* $j \in \mathcal{A}$ in the current period by a magnitude of λ . Similarly, consuming *any B brand* $j \in \mathcal{B}$ in the previous period shifts one’s utility from consuming *any B brand* $j \in \mathcal{B}$ in the current period by λ .

It is important to note the difference between this model and models of “brand loyalty,” where past consumption of a particular brand affects the current utility of consuming that particular brand. Such models have been explored in the economics and marketing literatures as outlined in the Introduction. In contrast, our framework offers a different type of habit formation: the habit to pick an expensive, highly-advertised item—or a cheap substitute—in a given product category. As detailed below, we offer a novel estimation strategy to capture this effect, and demonstrate its importance in the context of a fast-growing market.

We interpret former P&G CEO Alan Lafley’s notion of “commodity hell” as a situation where a large mass of consumers “became used to” paying low prices for, in our setting, soft drinks (P&G being a premium brand seller and heavy advertiser). It is then hard to persuade these consumers to pay as much as 100% more for a heavily advertised premium brand. In the opposite case, if one is used to paying premium prices for an “A brand,” that habit may also persist and increase the future utility from consuming premium brands. For parsimony, we allow both of these “habits” (a habit for premium and a habit for generic) to have the same magnitude λ , but a richer specification can allow these magnitudes to differ.

The notion captured by this model is that a non-trivial portion of the A-brand price premium extracts surplus that derives from “non-physical” product characteristics, stemming from differential advertising, and the strength of this “brand type image” in a consumer’s mind depends on recent consumption choices. A household who “goes generic” rather than “stays premium” realizes that the quality of a generic is not *that* shoddy, and the *relative* utility from consuming a B brand over an A brand in the next period rises by 2λ .

It seems natural that, as a lower middle class emerges, there is a growing mass of “new consumers” who are affluent enough to buy soft drinks, yet more price-sensitive than the established affluents. These households are therefore more likely to, at least, try out a generic substitute and develop a habit for “going generic.” For premium brands to win such customers over from the generic habit is, as a figure of speech, “ 2λ more difficult” in this model than in a model that does not have this persistence feature.

We refer to this baseline specification as the *Brand Type Persistence* (BTP) model, and it is on this specification that we base our empirical work. Below we describe two variants of this model that we estimate for comparison purposes. One of these variants allows for what we shall refer to as “soda category persistence (SCP),” where the habit that develops is toward consuming soda in general, rather than of a particular brand type. The second model variant (“no persistence,”

NP) does not include a persistence feature.

Note also that what we have in mind by “preceding period” is a time interval such as a month (or two) of consumption, consistent with our product market data, i.e., likely longer than the interval between shopping trips. While this is an obvious simplification (in particular, we do not allow for the full history of consumption choices to affect current utility), this is consistent with the extant marketing and economics literature of which we are aware (e.g., Dubé, Hitsch and Rossi 2010).

Finally, we model consumers as making static, current-period choices to maximize current-period utility. In other words, consumers are not forward-looking and do not internalize the effect of consumption today on their future utility. Given that soft drinks are: (i) non-durable goods, and (ii) relatively inexpensive items over which decisions are likely made with less-than-perfect foresight, we view these static consumer decisions as an appropriate modeling choice.²⁶ Nonetheless, the demand patterns implied by this model do reflect a dynamic “persistence” feature, and we seek to understand how this feature interacts with the dynamics in demography (i.e., socioeconomic mobility), and its implications for the competition between premium and generic brands.

The last term in the utility function, ϵ_{ijgt} , represents household and product-specific shocks that follow the Type I Extreme Value distribution and are i.i.d. across households, brands and markets. We complete the utility specification by defining the utility from the outside option, $u_{i \in r, j=0, gt} = \epsilon_{i, 0, gt}$.

The model’s parameters to be estimated are, therefore, denoted $\theta = \{\beta, \alpha, \alpha_{EA}, \alpha_{NA}, \lambda\}$. In the tradition of the literature on estimating random-coefficient logit models, we classify these into two subsets: the “linear parameters” are $\theta_1 = \{\beta, \alpha\}$ and the “non-linear parameters” are $\theta_2 = \{\lambda, \alpha_{EA}, \alpha_{NA}\}$ (Nevo 2000). As explained below, we extend this literature by showing how to incorporate into the estimation routine a dynamic demographic process that requires evaluating, for each guess of the model’s non-linear parameters, a prediction for the evolution of the type-distribution \mathcal{F}_{gt} . This process involves inverting the market-share function in each period in order to recover the type-distribution vector for the next period, and to follow this process forward over the 57 periods in each evaluation of the GMM objective function (see Subsection 3.3.1 below).

The share of type- r households consuming brand j in market gt is given by the logit formula:

$$s_{j,r,gt}(\theta) = \frac{\exp(\delta_{jgt} + \alpha_r \cdot p_{jgt} + \lambda \cdot \mathbf{1}\{persistence\}_{jr})}{1 + \sum_{\ell=1}^J \exp(\delta_{\ell gt} + \alpha_r \cdot p_{\ell gt} + \lambda \cdot \mathbf{1}\{persistence\}_{\ell r})} \quad (3)$$

²⁶An alternative is to follow the dynamic estimation literature and model consumers as maximizing an infinite-horizon utility function, making perfect predictions about the future path of prices and utility levels associated with soft-drink brands. Given the nature of the product, we view this as both an unnecessary and unjustified extension.

This notation reflects the fact that $s_{j,r,gt}(\theta)$ are model predictions and they depend on the value of the model's parameters.

The aggregate share of brand j in market gt is the weighted sum of the shares of brand j among household types $r = 1, \dots, R$, where the weights are the fractions of the market's population that belong in these types:

$$s_{jgt}(\theta) = \sum_{r=1}^R F_{r,gt} \cdot s_{j,r,gt}(\theta), \quad (4)$$

In the tradition of the literature on estimating random-coefficient logit models of differentiated product demand (Berry, Levinsohn and Pakes 1995, BLP), we estimate θ by matching the model predictions for aggregate market shares $s_{jgt}(\theta)$ from equation (4) with the observed values of these market shares in the Nielsen data.

Our work follows the literature on demand estimation in which consumer heterogeneity takes the form of discrete types (e.g., Berry, Carnall and Spiller 1996, Kalouptsi 2010). As is usual in such models, we do not observe type-specific shares that can be matched with the model predictions from (3). However, once θ that match *aggregate* shares are estimated, predictions for type-specific shares and elasticities follow. It is also worth noting that in this literature, the population fractions of the various types are often unobserved and are estimated as additional model parameters. In contrast, as explained below, our framework generates the fractions $F_{r,gt}$ by combining data on socioeconomic mobility with model predictions for consumption choices.

Of particular interest are the household-type specific elasticities. The household-type- r specific own-price ($j = k$) and cross-price ($j \neq k$) elasticities of demand for brand j are computed from:

$$\eta_{jk,r,gt} = \frac{\partial s_{j,r,gt}}{\partial p_{kgt}} \frac{p_{kgt}}{s_{j,r,gt}},$$

where, for brevity we omit the argument θ , and,

$$\frac{\partial s_{j,r,gt}}{\partial p_{kgt}} = \begin{cases} (\alpha + \alpha_r) s_{j,r,gt} (1 - s_{j,r,gt}) & \text{if } j = k \\ -(\alpha + \alpha_r) s_{j,r,gt} s_{k,r,gt} & \text{if } j \neq k \end{cases}$$

The aggregate price elasticities of demand for brand j are:

$$\eta_{jk,gt} = \frac{\partial s_{jgt}}{\partial p_{kgt}} \frac{p_{kgt}}{s_{jgt}} = \frac{p_{kgt}}{s_{jgt}} \sum_{r=1}^R F_{r,gt} \cdot \frac{\partial s_{j,r,gt}}{\partial p_{kgt}}.$$

Two additional variants on the modeling of persistence in tastes. To develop further

insight into our baseline demand model just described, we also consider two model variants. These variants change only the structure of the persistence component in the utility function.

Model variant 1 (“Soda Category Persistence,” SCP) This model variant captures persistence in the tendency to consume soda of any kind, as opposed to persistence in the tendency to consume premium versus generic brands. It defines the persistence term as:

$$1\{persistence\}_{jr} := \begin{cases} 1 & \text{if } r \in \{EA^A, EA^B, NA^A, NA^B, P^A, P^B\} \text{ and } j \in \mathcal{A} \cup \mathcal{B} \\ 0 & \text{otherwise} \end{cases}$$

In this model variant, the current-period utility from consuming *any* soda brand (premium or generic) is shifted by λ if the household consumed *any* soda in the previous period, i.e., the household has “adopted” soda consumption as part of its lifestyle.²⁷

Model variant 2 (“No Persistence,” NP) A third model we estimate constrains λ to zero, disallowing state dependence in preferences. We estimate this model to provide further insight into the implications of controlling for our hypothesized persistence mechanism when analyzing demand in the market we consider.

3.2 Dynamic type evolution

Having characterized the current-period utility function for each household type, and the type-specific and aggregate demand for each brand, we now describe the data-driven process of mobility, across the nine types, that households experience over time. We explain how we compute the “type-distribution” vector \mathcal{F}_{gt} defined in Subsection 3.1. Recall that this nine-element vector captures the fraction of market gt ’s household population that belongs in each of the nine types.

We begin by computing \mathcal{F}_{g1} for regions $g = 1, \dots, 7$ from the information in our Data Source 3 (HEX 95/96). As explained in Section 2, that survey allows us to link a household’s socioeconomic standing (which we classify, following IBOPE, as either **ABC** or **DE**) to its consumption choice (A brand, B brand, or no soda). By virtue of the survey having been conducted shortly before the start of our Nielsen market-level sample, we use these fractions as “initial values” for our type-distribution vector. Recall that for period $t = 1$, for each region, our model sets the number of established affluent and poor households equal to the number of **ABC** and **DE** households, respectively, with the number of newly affluents set to zero. To be clear, for each region $g = 1, \dots, 7$, the fraction of HEX **ABC** (resp., **DE**) households who consume an A brand, a B brand

²⁷As a motivating example, imagine a child who has consumed soda at home in recent weeks and nags the parent if soda is no longer present at the table on weekends. The child is less likely to nag the parent if he has experienced only juice or tap water over the preceding month.

or no soda is used to set the values of $F_{EA^A,g1}, F_{EA^B,g1}, F_{EA^O,g1}$ (resp., $F_{PA,g1}, F_{PB,g1}, F_{PO,g1}$), whereas $F_{NA^A,g1}, F_{NA^B,g1}, F_{NA^O,g1}$ are set to zero.

We now explain how, *for a given guess of the model's parameters*, these fractions are updated forward for periods $t = 2, \dots, 57$. In the next subsection we explain the procedure of estimating this model using GMM. As we clarify there, since every guess for the model's parameters predicts a different evolution path of \mathcal{F}_{gt} over $t = 2, \dots, 57$, such predictions are made as part of each evaluation of the GMM objective function.

For the purpose of illustration, let us fix region g and period t (i.e., the gt market) and assume that the type-distribution vector \mathcal{F}_{gt} is given. We now explain how, given a guess for the model's parameters, we can update forward and find the type-distribution vector $\mathcal{F}_{g(t+1)}$. Repeating this updating process for $t = 1, \dots, 56$ (noting from above that \mathcal{F}_{g1} is known) yields the full trajectory of the type distribution over the sample period.

Intuitively, a guess of the model's parameters yields a prediction of the shares (and masses) of type- r households who purchase any premium or generic brand in period t via equation (3). Had we not allowed social mobility, the computation of $\mathcal{F}_{g(t+1)}$ would be straightforward by simply summing, across the three types in each socioeconomic group, the number of individual households who in period t consumed a given brand type (premium or generic), and dividing this sum by the period $(t + 1)$ total household population. For instance, $F_{EA^A,g(t+1)}$ would have been obtained by:

$$\frac{\# \text{ of period-}t \text{ } EA^A \text{ buying A} + \# \text{ of period-}t \text{ } EA^B \text{ buying A} + \# \text{ of period-}t \text{ } EA^O \text{ buying A}}{\text{Period } (t + 1) \text{ household population}}$$

The social mobility process, however, complicates these computations substantially. For instance, some of period- t poor will become period- $(t + 1)$ newly affluent, if aggregate upward mobility is detected between the two periods in the IBOPE data. We may also witness aggregate downward mobility, in which case some of period t 's newly affluent will join the ranks of period $(t + 1)$'s poor. Complicating matters even further is migration from rural areas to the urban markets that our three datasets cover.²⁸

To incorporate information on *aggregate social mobility*, we need a set of assumptions on this process. These are summarized below:

Assumption 1. *The following patterns characterize the process of mobility across the nine household types:*

²⁸Had we observed individual-level data over time (as opposed to a single micro-level cross-section at the beginning of the studied period), such issues would have been trivial. On the other hand, such repeated observations on a fixed pool of households may miss the very dynamic nature of this market that we seek to investigate (e.g., rural-to-urban migration).

(i) The share of households moving “up” from POOR to NEWAFF who consumed an A brand (resp., B brand) in the previous period is equal to the overall share of POOR who consumed an A brand (resp., B brand) in the previous period;

(ii) The share of households moving “down” from NEWAFF to POOR who consumed an A brand (resp., B brand) in the previous period is equal to the overall share of NEWAFF who consumed an A brand (resp., B brand) in the previous period;

(iii) Any increase (resp., decline) in the region’s urban household population is interpreted as migration from (resp., migration to) the rural area outside the urban region;

(iv) Households migrating into the urban region join the POOR group and did not consume soda in the previous period, prior to migrating (i.e., these migrants join P^O types);

(v) Households migrating out of the urban region leave the POOR group, and consumed soda in the previous period according to the shares of A brands, B brands and the outside good consumed by the POOR in that period (i.e., these migrants leave P^A , P^B and P^O types).

The appendix illustrates this detailed process by way of an example.

Assumption 1 is needed for “accounting” purposes, i.e., to allow us to keep track of the masses (and fractions) of various household types over time. One may wonder whether these assumptions are too strong, in particular, part (iv) which states that rural migrants into the urban areas have a “no soda” habit upon their migration. We believe such an assumption to be reasonable since we model purchases in self-service outlets (supermarkets) that are prevalent in Brazil’s urban areas. Since we model consumption habits, not endowing rural migrants with habits to either pick a premium brand or grab a generic off the supermarket shelf seems plausible.²⁹

Notice that, in general, one is not likely to have reliable and comprehensive micro-level data that keeps track of households’ consumption choices as well as their socioeconomic standing over time, particularly in a market that is in a state of flux. We thus view the tools developed here—which rely on rich aggregate demographic data—as a methodological contribution.

3.3 Identification and estimation

3.3.1 The estimation procedure

We now outline the procedure that estimates the utility parameter vector θ . This procedure is a novel contribution to the literature in that it incorporates into the familiar method of estimating random coefficient logit models the feature that tracks the evolution of region-specific type distributions over time.

While there is a developed literature on estimating such models in which consumers are classified into discrete types, these types are often abstract groupings of “similar” consumers for

²⁹Just the same, we are in the process of confirming the robustness of our results to these assumptions.

which population weights are parameters to be estimated.³⁰ In our method, these population weights are computed from a combination of data (the IBOPE data that classifies households into different socioeconomic groups) and model predictions regarding the fractions within each socioeconomic group that consumed premium (resp., generic) brands in the preceding period and therefore have a premium (resp., generic) habit in the current period. The fact that model predictions for period t affect the type distribution in period $t + 1$ requires us to incorporate a dynamic updating routine into each evaluation of the GMM objective function.

Conditional on a type-distribution vector \mathcal{F}_{gt} , each guess of the model’s parameters yields a prediction for the period- t brand shares for each household type, and thus a prediction for aggregate shares. Inverting the aggregate market share function (using the BLP contraction mapping) allows us to recover period- t base utilities for each brand δ_{jgt} and obtain values for the disturbances ξ_{jgt} . We then use the type-specific period- t shares, data on social mobility, and Assumption 1, to predict a type-distribution vector for the next period ($t + 1$), i.e., $\mathcal{F}_{g(t+1)}$.

Starting at period $t = 1$, the process (including the inversion of the market share function) is repeated for periods $t = 2, \dots, 57$ and allows us to compute the demand unobservables ξ for every brand in every region-period market. Note the unique feature that the market-share inversion for period $t + 1$ actually depends on information obtained from the inversion of that equation in period t . As a consequence, these inversions cannot be performed independently but rather they must be performed in a recursive fashion.

Once the unobservables ξ are computed for each brand in each region-period market, given the current guess of the model parameters, we proceed along familiar lines from the literature on estimating random coefficient logit models. Specifically, as discussed in detail in Subsection 3.3.2. below, we make the identifying assumption that these demand unobservables are mean-independent of a set of instruments Z . This allows us to construct a GMM objective function that is minimized at parameter values that minimize the correlation between the instruments and the demand unobservables.

We now outline the algorithm’s steps, leaving the exact form of the objective function and other details to the appendix.

1. Pick a generic guess for $\tilde{\theta}_2$ (non-linear parameters).
2. For each region $g = 1, \dots, 7$, perform the following: given $\tilde{\theta}_2$ and \mathcal{F}_{g1} , invert J market-share equations (4) to retrieve $\Delta_{g1} = (\delta_{1g1}, \dots, \delta_{Jg1})'$, the period-1 vector of base utilities for each of the J brands. Stack all these first-period base utilities in a $(7J \times 1)$ vector Δ_1 .
3. For each region $g = 1, \dots, 7$ and type $r = 1, \dots, R$, use household-type specific brand-share

³⁰For example, Nair (2007) models video-game consumers as either “high valuation” or “low valuation” types, and estimates their relative population fractions.

equation (3) to predict the shares of type- r households who purchase each type of brand (A brands or B brands) in period 1.

4. Use these period-1 type- r A-brand and B-brand shares, the data on how the size of each socioeconomic group within each region evolves, and Assumption 1 to update forward, for each region $g = 1, \dots, 7$, the proportion of households in period 2 who belong to each of the R types, \mathcal{F}_{g2} .
5. Repeat steps 2-5 for periods $t = 2, \dots, T$.
6. Stack the base utility vectors Δ_t for all periods, regions and brands in the vector $\delta(\tilde{\theta}_2)$. Evaluate the GMM objective at the guess $\tilde{\theta}_2$.³¹
7. Update the guess $\tilde{\theta}_2$ and repeat the above steps until we find the value of θ_2 that minimizes the GMM objective.

3.3.2 Identification

How can a persistence feature be identified in a random-coefficient logit model? In the absence of household-level panel data, it is not *a priori* clear how persistence can be separately identified from mean-utility levels as either characteristic can imply high, persistent market shares for popular brands. We now argue that the rich variation in our data enable us to identify not only persistence but also consumer heterogeneity under the persistence structure of the baseline model. In Section 4, we will show that controlling for persistence in the way we do actually helps us identify heterogeneity in household price sensitivity; in particular, the structure we build into our baseline model yields the compelling result that newly affluent households display levels of price sensitivity that are intermediate between the established affluent and the poor.

We next explain how we exploit two kinds of cross-regional and time-series variation in the data: socioeconomic transitions and variation in price.

Socioeconomic transitions. The first type of data variation we exploit is the shifting demographics—both the growth of the middle class from 1996/97 on and the subsequent partial reversion during the recession that started around 2000/01. These shifts provide a key source of variation. Importantly, these socioeconomic transitions occurred at differential rates across regions. This intra-region temporal variation is important since our inclusion of brand-region fixed effects should soak up fixed differences across regions in the tastes for brands and for soda; households in some regions may on average value certain brands, or soda as a category, more than households in other regions, e.g., for cultural or historical reasons.

³¹Following Nevo (2000), the linear parameters θ_1 that minimize the objective can be concentrated out analytically conditional on the guess for θ_2 .

To illustrate how shifting demographics help identify the parameters in our model, consider two regions that vary widely as to how their socioeconomic composition evolved over the sample period: region 1 (states in the Northeast) and region 4 (São Paulo Metro). The following tables depict some socioeconomic and product market data points for these two regions:

% Urban households	Dec/Jan 1997		Dec/Jan 2000		Dec/Jan 2003	
	Poor	New. Affl.	Poor	New. Affl.	Poor	New. Affl.
Region 1 (Northeast)	65%	—	44%	24%	57%	15%
Region 4 (São Paulo Metro)	36%	—	23%	16%	23%	19%

Sources: IBOPE, IBGE

Shares: No Soda & Generics	Dec/Jan 1997		Dec/Jan 2000		Dec/Jan 2003	
	$s_{0,gt}$	$s_{B,gt}$	$s_{0,gt}$	$s_{B,gt}$	$s_{0,gt}$	$s_{B,gt}$
Region 1 (Northeast)	87%	0.3%	82%	5.0%	80%	8.4%
Region 4 (São Paulo Metro)	61%	6.6%	62%	12.6%	59%	11.9%

Sources: Nielsen, IBGE. See Section 2 for definition of shares.

At the start of our sample period, region 1 is substantially poorer than region 4 (65% of region 1’s urban households are poor vis-à-vis 36% for region 4) and, at the same time, exhibits lower soda penetration relative to its wealthier counterpart (87% of region 1’s households do not consume soda against 61% for region 4). Notice that this cross-sectional variation can in principle be explained not only by the poor being more price sensitive than the established affluent, but also by region 1 potentially exhibiting a lower preference for soda relative to region 4. This latter variation will thus be captured by the brand-region fixed effects. It is, therefore, the intra-region temporal variation in the socioeconomic composition of households that helps us pin down the heterogeneity in price sensitivities.

To see this, notice that from 1997 to 2000, region 1 boasted stronger upward mobility relative to region 4: by 2000, 24% of region 1’s households were newly affluent compared with 16% of region 4’s households. Over these same years, region 1’s soda penetration, $1 - s_{0,gt}$, grew substantially, from 13% to 18% (in proportionate terms, 5 points in 13), while soda penetration in region 4 was about flat. In particular, generic brands in region 1 enjoyed a huge gain in share, $s_{B,gt}$, from 0.3% to 5.0% (in region 4, the share of generics “only” doubled).

Notice that our model allows price sensitivity to vary by socioeconomic standing, i.e., different price sensitivities can characterize the established affluents, the newly affluents and the poor. Variation in the sizes of such consumer groups thus changes the aggregate price sensitivity within each region. As a consequence, we use the co-variation of these shifts in aggregate price sensitivities, on the one hand, and soda market shares, on the other hand, *to identify price sensitivity*

effects.

Another important demographic shift is the reversal in upward mobility that we observe after 2000, and the differential effects of this recession in the different regions. As the tables above indicate, region 1 saw the proportion of newly affluent households shrink considerably, from 24% in 2000 to 15% by 2003, yet the penetration of soda consumption did not fall.³² In fact, the share of soda, and particularly the share of generics, continued rising, albeit at lower rates than previously.

Importantly, the recession was not accompanied by declining soda prices (see Figure 3). As a consequence, households' tendency to continue consuming soda despite being hit by a recession is suggestive of persistence in preferences. This data variation is thus quite effective in identifying the parameter λ in our model. It nicely exemplifies how *persistence manifests itself directly in our data*—as opposed to being an artifact of our econometric strategy. We provide another such example below.

While the above discussion explained how social mobility helps identify both price sensitivities and persistence in preferences, we argue that controlling for one of these features actually helps us identify the other. In particular, imagine that we did not allow for a persistence feature in preferences. Our model would then interpret the recessionary period in the data as evidence that the poor are “not that price sensitive” when it comes to soda consumption (since all one would see through the lens of such a model is an increasingly poor population consuming stable amounts of soda). In this model, it would be difficult to elicit greater price sensitivity among poor households compared to the more affluent groups. Indeed, in Subsection 4.2 below we show that our “no persistence” model variant is outperformed by our baseline model in terms of picking up the heterogeneity in price sensitivities among the different groups.

Controlling for persistence, therefore, is not only important in its own right, but plays an important role in identifying other dimensions of household preferences. To our knowledge, this is the first paper that makes this point. We view this modeling ingredient as particularly important when studying demand in emerging markets.

It is also worth noting that shifting demographics can help identify demand elasticities even in markets which experience price stability, such as the later part of the sample period (though the data *does* exhibit helpful price variation as discussed next). The within-region co-movement between the socioeconomic composition of households and aggregate market shares, when prices remain unchanged, is then explained by the heterogeneous household behavior, i.e., by shares $s_{j,r,gt}$ which vary across household types r by virtue of the parameters θ_2 .

Price variation. The second kind of data variation that we exploit for identification stems from price variation which we argue to be exogenous to *unobservable* demand shocks at the

³²Also recall figures 1 and 2, and note the suggestive contrast between soda and cement consumption.

region-period level ξ_{jgt} .

First, consider the gradual but relentless price reduction in the competitive fringe over the period 1996 to 2000. The close substitutability among each of these generic brands suggests that their price should, to a large extent, stay close to their marginal cost of production and distribution. This view suggests that the decline in fringe prices was driven predominantly by supply (cost) side considerations rather than being a response to unobserved demand shocks. In particular, one may imagine that expanding capacity and scale, learning effects and exit of less-efficient generic producers may all have contributed to declining costs, and thus prices, within the competitive fringe.

Through most of the period of declining fringe prices, generics were able to grow substantially at the expense of the premium brands and the outside good. To stay with the examples from the data (see the table of shares above), between 1997 and 2000: (i) in region 1, $s_{B,gt}$ grew by 5 points, with $s_{A,gt}$ (i.e., $1 - s_{B,gt} - s_{0,gt}$) holding up quite well; and (ii) in region 4, $s_{B,gt}$ grew by 6 points, and at the expense of $s_{A,gt}$. Intuitively, such co-movement in prices and shares can be picked up by the heterogeneous price sensitivity and persistence that our baseline model allows for. Of note, it is in region 1 that the reduction in the price of generics is most pronounced, concomitant with that region’s strong upward mobility and steep growth in soda penetration, particularly of the generic type.³³

We further argue, following Salvo (2009), that the premium brands’ abrupt and substantial price cut in mid 1999 was exogenous to demand unobservables. The argument rests on the notion that this large and sudden price drop was plausibly a response to the demographic shifts and expansion of the fringe that we *observe* over 1996-1999, and *not* a response to some sudden unobserved mid-1999 demand shock caught up in the residual (noting that we also control for advertising intensity, weather shocks and different forms of drift).

Recall that, following this abrupt price cut, the growth of generics petered out—but it did not reverse. This can be seen in the tables above, for two regions, and in Figure 3, for all regions combined. We view this feature of the data as further evidence of persistence in preferences, and, specifically, as evidence for our Brand Type Persistence (BTP) mechanism. That is, a habit for “going generic” among a share of households made it more difficult for Coca Cola to win them over, even via a drastic price cut. The price cut did, however, help protect premium brands from *further* market share losses.

While the above paragraphs might suggest that major sources of price variation in our sample (i.e., the steady price decline in the fringe and the abrupt premium price cut) could be viewed as exogenous to demand unobservables, in general prices are treated as endogenous in our framework. We address this endogeneity in a manner consistent with the bulk of the demand

³³In region 1, $p_{B,gt}$ starts in Dec/Jan 1997 at R\$ 1.15, compared to a mean R\$ 0.85 across the other six regions, but by 2000 drops to a similar level to that of the other six regions.

estimation literature: we instrument for the endogeneity of prices. In particular, as explained in the description of the estimation procedure in Subsection 3.3.1 above, we interact the demand unobservables ξ with instruments for which there is a strong case for exogeneity and which effectively shift prices via shifting markups and costs. We then search for parameters that induce the correlation between the demand unobservables and the instruments to be as close to zero as possible. Additional details on this GMM procedure can be found in the appendix.

We adopt the demand instruments used by Salvo (2009).³⁴ In what follows, we briefly explain how we control for potential price endogeneity, referring the reader to Section III(i) of Salvo (2009) for more detailed discussion. We specify three classes of demand instruments (and, in the appendix, further consider the robustness of our identification strategy). The first class of instruments is afforded by the premium brands’ abrupt price cut, as justified in the preceding paragraph. In practice, we generate a dummy variable which takes on the value 1 for all time periods after July 1999, and interact it with brand-region fixed effects, thus allowing the effects of this *supply-side* shift to vary by brand within each region.

The second class of demand instruments borrows from Hausman, Leonard and Zona (1994). Specifically, we use each brand’s price in the other regions to instrument for its price in a given region. The identifying assumption is that a brand’s prices across the different regions are correlated through a common cost structure or through common shifts in the way firms strategically interact, not least the mid 1999 premium price cut (which was synchronized across regions). We believe that common demand unobservables (e.g., Bresnahan 1997a, 1997b) are of lesser concern in our setting, due to two key reasons: (i) the inclusion of region-specific brand-level advertising controls (which are often absent in demand studies), on top of other demand observables; and (ii) the considerable regional variation in demand and thus the very local nature of Brazilian soft drink distribution and promotion. It is worth noting that, perhaps unsurprisingly, the penetration of national retailers is still limited relative to the United States.

Factor prices comprise the third and final set of exclusion restrictions. These classic cost-shifters, assumed to be uncorrelated with unobserved demand shocks, include refined sugar prices, plastic and other packaging prices, and different energy prices (electricity, road transport).

4 Estimates and counterfactuals

4.1 Baseline model: Brand Type Persistence

Table 2 reports estimates for the baseline model, which allows for “brand-type persistence.” Varying across the three reported specifications is the inclusion or not of linear time trends in

³⁴Salvo (2009) estimates a continuous-choice (AIDS) model, a different approach to the discrete-choice model we offer in this paper, which emphasizes the rich consumer heterogeneity. Just the same, instrumenting for price endogeneity is similarly relevant to both frameworks.

the vector of base utility covariates x : the specification in column 1 does not include trends, column 2 specifies (seven) region-specific time trends, and column 3 allows these time trends to also vary at the brand level (i.e., $9 \times 7 = 63$ brand-and-region trends are included). Trends in column 2, for example, thus account for the possibility that unobservable tastes for soda vary linearly over time at the regional level. In all three specifications, x additionally includes: (i) (63) brand-and-region fixed effects; (ii) five bimonthly seasonal dummies (Dec/Jan being the peak summer season, and Oct/Nov being the omitted category); (iii) market gt 's mean temperature (to illustrate within-season variation, winter temperatures in the southern region 6 averaged $15.1^\circ C$ in July 2001 against $12.1^\circ C$ in July 2000); and, finally, (iv) region-specific effects of brand-level media advertising (again to illustrate, Coke brand's advertising intensity in São Paulo Metro region 4 amounted to 2199 Gross Rating Points in Dec-2000 rising to 3587 GRP in Dec-2001; Pepsi's GRP were 351 and 598 respectively in these same markets).

We find that across all specifications, estimates are intuitively signed and highly statistically significant.³⁵ In particular, we obtain that established affluent households, while still caring about prices (i.e., $(\alpha + \alpha_{EA}) < 0$), are less price sensitive than the newly affluent ($\alpha_{EA} > \alpha_{NA}$), who in turn are less price sensitive than the poor ($\alpha_{NA} > 0$, $\alpha < 0$). The estimated brand-type persistence parameter λ of this model (i.e., a taste or inertia for consuming a certain brand type—premium or generic) is significantly positive, a mechanism we will subsequently interpret in light of the above-mentioned model variants. Further, base utility estimates are significantly increasing in: (i) the peak summer season; (ii) temperature, controlling for seasonality; and (iii) advertising intensity, controlling for seasonality.

The lower portion of Table 1 summarizes the heterogeneous price sensitivities first by way of aggregate own-price elasticities, $\eta_{jj,gt}$, for a handful of brands j , followed by household-type- r specific own elasticities, $\eta_{jj,r,gt}$ (in both cases, the table reports elasticities averaged over markets gt). In column 2 (unless noted otherwise, we hereafter refer to this specification), a 1% increase in the price of Coke lowers aggregate demand for Coke by 1.6%, compared with slightly higher elasticities of -1.9 to -2.0 for the other three leading premium brands, Guaraná Antarctica, Fanta and Pepsi. As the table illustrates for Coke, price sensitivity is decreasing in the household's socioeconomic standing (e.g., $|\eta_{Coke,PA}| = 4.1 > |\eta_{Coke,NA^A}| = 2.2 > |\eta_{Coke,EA^A}| = 1.3$), and, within socioeconomic group, the demand for a certain type of brand (i.e., premium or generic) is less elastic for households who have developed a habit to consume that particular brand type (e.g., $|\eta_{Coke,NA^B}| = 3.0 > |\eta_{Coke,NA^A}| = 2.2$). Thus, past consumption of premium brands lowers a household's elasticity for Coke (whether the household is established affluent, newly affluent or poor) and, similarly, past consumption of generics lowers its elasticity for generics. To further illustrate, Figure 4 plots the temporal evolution of own elasticities, by household type, for Coke

³⁵There is only *one* exception, $\hat{\alpha}_{NA}$ in the specification without time trends (column 1), where the point estimate, while positive, is not statistically different from zero.

brand in region 4.

As for generics, the mean estimated aggregate own-price elasticity is -0.75 . While this value may seem low, remember that this “brand” is actually an aggregation of the many similar generic brands available in the market, and that each of these individual generic brands faces fierce competition from each other and much more elastic demand than the fringe of generics as a whole.³⁶

The model indicates that, compared to the established affluent, newly affluent households consume a disproportionate amount of discounted generics relative to premium brands, a result that follows intuitively from the estimated price sensitivities. Similarly, the poor consume a disproportionate amount of generics relative to premium brands when compared to the newly affluent. Taking means across all regions and time periods, the estimated premium and generic shares are: (i) 37% and 17%, respectively, among established affluent households (with the outside good accounting for the remaining 46% share), i.e., a 2.1:1 premium-to-generic ratio; (ii) 14% and 12%, respectively, among newly affluent households, or a 1.1:1 ratio; and (iii) 2.6% and 4.1%, respectively, among poor households, or a 0.6:1 ratio. The fitted baseline model supports casual observation that the new middle class consumes more than the poor (an inside share of 25% against 7%), yet is not as price insensitive as the established affluent (with an inside share of 54%). These premium and generic shares by socioeconomic group are reported in Table 5, below, alongside shares estimated under each of the model variants (described below).

4.2 Counterfactual: What if premium brands failed to cut prices?

One of the striking features of the data is the premium brands’ sharp price cut, led by Coca-Cola, almost halfway through the period. As the solid lines in the left panel of Figure 5 show (for region 5, but the variation is similar for other regions), prices in the fringe declined gradually from R\$ 0.80 to R\$ 0.55 between late 1996 and mid 2000, while premium brand prices stayed broadly flat at about R\$ 1.15 until mid 1999, then dropping—abruptly—to R\$ 0.90 and staying at this lower level. (All prices per liter, and recall that in constant terms.)

We employ the baseline model to predict how shares would have evolved had the established firms *not* cut prices by over 20%, instead keeping them at their early 1999 levels. This counterfactual price path is marked by the dashed line in the left panel of Figure 5. In this counterfactual, we keep fringe prices equal to the ones observed in the data (i.e., those observed in the presence of the premium price cut). This assumption is justified by the competitive nature of the fringe, implying that fringe prices cannot deviate much from marginal costs (in further support of this

³⁶For the sake of comparison, Salvo (2009, Table III, column IV), estimating an AIDS demand system from the same Nielsen sample and with the same set of demand instruments but, importantly, without access to demographic or household-level data, obtains a (somewhat similar) own-price elasticity of demand of -2.1 for Coke, -2.5 for Guaraná Antarctica and (a substantially larger) -2.8 for the aggregation of B brands.

competitive view, note that there was no apparent price response from the fringe to the actual price cut).

The right panel reports the estimated impact on aggregate market shares for premium brands and for generic brands (counterfactual share paths are again marked by dashed lines). In contrast to the actual scenario in which their decline was stemmed, the established brands’ joint share—by not cutting prices—would have hit rock bottom at 12% by the winter of 2000, *lower* than the share of their generic counterparts.

It is worth noting that much of the counterfactual loss in premium brands’ shares would go to the outside good, rather than to generics. For example, inside shares at actual (reduced premium brand) prices over 2001-02 average 46% (28% premium plus 18% generic) to be compared with inside shares of 39% (17% premium plus 22% generic) at (higher) counterfactual prices. This is suggestive of substantial market segmentation, likely amplified by the BTP mechanism that limits the scope for “business stealing” across product types. Still, the 4 percentage point growth in the generic share would have represented almost a one-quarter increase (+4.0/17.6) in the fringe’s penetration. We further discuss this below.

4.3 Model variants: Different persistence structures

Tables 3 and 4 report estimates for the two variations of the baseline BTP model. In Table 3, the persistence shock to utility from consuming a soda brand applies in the event of any type of brand—irrespective of whether premium or generic—having been consumed in the previous period. In Table 4, there is no persistence: previous soda consumption does not shift current utility. Relative to the baseline model, price coefficients in both model variants are estimated to be lower (in absolute value). In particular, these estimates suggest that the poor are less price sensitive than in the baseline model, and we can now no longer reject (with region-level time trends included—again shown in column 2 in each table) that the newly affluent and the poor are equally price sensitive.

For intuition on why we obtain such differential results across models, recall that one source of identification for λ is the recession that began around 2000, which saw a non-negligible proportion of newly affluent households fall back below the poverty line, yet soda consumption did not fall back to former levels (recall our earlier intuitive arguments for identification, as well as Figures 1 and 2). In both the baseline model and the first variant, this co-movement is picked up by λ , whereas in the second variant λ is set to zero. Thus, in this no-persistence variant, it would appear that the newly affluent who fell from grace and returned to poverty, *but who continued consuming soda into the recession*, are not price sensitive. This explains why $\hat{\alpha}$ might be biased toward zero in the no-persistence model variant. The baseline model and the first model variant enjoy degrees of freedom which the more restrictive “no persistence” model does not.

As for the differential results across the baseline model and the first model variant, recall another key source of identification in the data: the 1999 premium brand price cut (one can make a similar argument on the basis of the 1996-2000 price reduction in the fringe). It is clear that the premium versus generic share response to such pronounced price variation is captured differently in the baseline model, with its *brand-type* specific persistence mechanism, relative to the variant in which persistence does not distinguish between previous premium or generic consumption.

We know from our Data Source 3 (the HEX 95/96) that the poor consume a disproportionate amount of discounted generics relative to premium brands *as compared to the established affluent* (and, from the baseline model, this is also the case for the newly affluent relative to the established affluent). Upon premium brands cutting prices and winning share, they attracted newly affluent and poor households who had previously consumed the outside good (or retained past consumers), given that the share of generics held up relatively well (i.e., it did not collapse)—households who had been consuming generics tended to continue doing so, as they were still not willing to pay for Coke at reduced prices (R\$ 0.90) when generics at R\$ 0.55 had been proving to their taste buds to be of reasonable quality. In the baseline model, this covariation is captured by a strong brand-type specific λ , freeing up the α 's to work “properly.” In the any-soda variant, the low generic share response to the premium price cut is confounded as lower price sensitivities for the poor and the newly affluent (who are disproportionately attracted to generics over established brands relative to the established affluent), i.e., $\hat{\alpha}$ and $\hat{\alpha}_{NA}$ are biased toward zero.

Table 5 reports soda penetration and premium-to-generic consumption ratios by socioeconomic group predicted under each of the model variants, and compares these to the predicted levels and ratios that we discussed earlier for the baseline Brand Type Persistence (BTP) model. Comparing the Soda Category Persistence (SCP) model with BTP, the penetration of soda among poor households—particularly of premium brands—increases and that of newly affluent households decreases; under SCP, the poor are as likely to consume premium brands as they are to consume generics (a 1.0 premium-to-generic ratio), behavior which is similar to that displayed by the newly affluents (a 1.2 ratio). In the No Persistence (NP) model, the poor are predicted to consume *more* soda—and of either brand type—than the newly affluent, a result that is economically unappealing.

Figure 6 reports shares in the same counterfactual experiment we implemented under the baseline model (premium brands not cutting prices), but this time under the Soda Category Persistence model variant, in the left panel, or the no-persistence variant, in the right panel (we stick to the region 5 illustration of Figure 5, for the baseline model). Inspection of the two panels (model variants) relative to Figure 5's right panel agrees with the intuition we just developed. The premium brands' counterfactual loss from not cutting prices, or equivalently, the premium

brands' actual gain from cutting prices, is larger in the baseline model than it is were the data generating process to accord with either model variant.

Though examining the premium brands' pricing policy is the subject of a sequel paper, this finding does lend tentative support to the hypothesis that Coca-Cola acted rationally on cutting prices, particularly insofar that current market share is an "asset," predictive of future profit.³⁷ By not cutting prices, premium shares are predicted by either model variant to counterfactually reach 15% in winter 2000, compared with a predicted "rock-bottom" 12% in the baseline model. In the same vein, the first and second model variants similarly predict that generics, when no longer faced with a price cut from their premium rivals, would hardly have gained ground (whereas the baseline model predicts that the penetration of generics would have risen from 15% to 17% by winter 2000).

These results suggest that it is the BTP model that provides the strongest support for the observed premium price cut. While all models show that premium brands would have suffered a substantial market share loss had they not cut prices, it is the BTP model that predicts the deepest damage to the premium brands' market position.³⁸ The intuition underlying this result has been developed extensively in this paper: it is in the "commodity hell" world of the BTP model that premium brands must act quickly and cut prices, to avoid having "too many" consumers develop a generic habit, leading to substantial difficulties convincing them thereafter to pay a much higher price for a premium brand.

5 Concluding remarks

In this paper we have offered an empirical analysis of two salient features of the Brazilian soft drink market: the marked growth of a "new middle class," on the one hand, and the rapid growth of the generic fringe, on the other hand. Using unique data with very rich cross-sectional as well as temporal variation, we estimate a model that highlights two features which we deem highly important in this market: the heterogeneous price sensitivities of different socioeconomic groups, and persistent household preferences—of a special kind. The persistence we motivate and demonstrate empirically has to do with households developing a habit of "shopping premium" or "going generic."

This mechanism captures a world in which premium brands have to act quickly in the wake of an emerging middle class. If they wait too long, a substantial mass of the "new middle class" might be captivated by the generic habit. It may then prove to be much more difficult to convince

³⁷See Bronnenberg, Dhar and Dubé (2009) on the persistence of brand market shares.

³⁸A preliminary back-of-the-envelope calculation indicates that it is the BTP model that also predicts the deepest \$ damage to the premium brands had these not cut prices. Using information gleaned from Ambev's local SEC filings, conversation with industry insiders, among other sources we estimate that the premium brands' combined variable profit (excluding fixed costs) during the first three years after the price drop amounted to an actual R\$ 1.9 billion, to be compared to R\$ 1.5 billion had the price cut not occurred. The damage in profit had premium brands not cut prices is predicted to be lower by the two variant models.

these consumers to pay much higher prices for a highly advertised premium brand.

As we motivate extensively, this paper wishes to add to several strands of literature. We develop a new model of habit formation and offer a new methodology for its estimation, relying mostly on more widely available aggregate data (as opposed to repeated observations at the household level which may not be available in an emerging market, or may fail to capture the dynamics of social mobility, such as rural-to-urban migration).

While our application focuses on the very concrete example of the Brazilian soft drink market, we view the issues that are tackled in this work as likely characterizing many consumer goods markets in the developing world. Consumers in emerging markets are viewed as an engine of growth for multinational firms and for the global economy as a whole. Understanding the microeconomics of competition in such markets (in particular the tough match multinational premium brands face from local “value players”), and particularly the features of demand in such markets, should be of great interest for policymakers and firms alike.

References

- [1] ABEP, Brazilian Association of Market Research Firms (2003). Critério de classificação econômica Brasil: Sistema de pontos. ABEP Publication, São Paulo. (Points system for classifying households into economic segments: Brazil criteria, in Portuguese.)
- [2] Ambev (2003). Informações anuais (IAN). Comissão de Valores Mobiliários, CVM. (Annual report to Brazil’s Securities and Exchange Commission, in Portuguese.)
- [3] Boston Consulting Group (2002). Mercados pouco explorados: Descobrimo a classe C. BCG Publication, São Paulo. (Under-explored markets: Discovering the “C” class, in Portuguese.)
- [4] S. Berry, J. Levinsohn, and A. Pakes (1995). Automobile prices in market equilibrium. *Econometrica* 63, 841-90
- [5] Berry, S., M. Carnall, and P. T. Spiller (1996). Airline hubs: Costs, markups and the implications of consumer heterogeneity. NBER Working Paper 5561
- [6] Bresnahan, T. F. (1997a). Valuation of new goods under perfect and imperfect competition: A comment, in Bresnahan, T. F. and R. Gordon (eds.), *The Economics of New Goods, Studies in Income and Wealth* Vol. 58. National Bureau of Economic Research, Chicago
- [7] Bresnahan, T. F. (1997b). The Apple Cinamon Cheerios war: Valuing new goods, identifying market power, and economic measurement. Stanford University, Mimeo

- [8] Bronnenberg. B., S. Dhar and J.-P. Dubé (2009). Brand history, geography, and the persistence of brand shares. *Journal of Political Economy* 117, 87-115
- [9] Carvalho Filho, I. and M. Chamon (2011). The myth of post-reform income stagnation: Evidence from Brazil and Mexico. *Journal of Development Economics*, in press
- [10] Chaudhuri, S., P. K. Goldberg and P. Jia (2006). Estimating the effects of global patent protection in pharmaceuticals: A case study of quinolones in India. *American Economic Review* 96, 1477-1514
- [11] Dubé, J.-P., G. J. Hitsch, P. E. Rossi and M. A. Vitorino (2008). Category pricing with state-dependent utility. *Marketing Science* 27, 417-429
- [12] Dubé, J.-P., G. J. Hitsch and P. E. Rossi (2010). State dependence and alternative explanations for consumer inertia. *RAND Journal of Economics* 41, 417-445
- [13] Economist (2007). Adiós to poverty, hola to consumption: Faster growth, low inflation, expanding credit and liberal trade are helping to create a new middle class in Latin America. Print edition of August 18, 2007
- [14] Economist (2011a). Politics in emerging markets: The new middle classes rise up. Print edition of September 3, 2011
- [15] Economist (2011b). Indonesia's middle class: Missing BRIC in the wall. Print edition of July 21, 2011
- [16] Eichenbaum, M. S., L. P. Hansen and K. J. Singleton (1988). A time series analysis of representative agent models of consumption and leisure choice under uncertainty. *The Quarterly Journal of Economics* 103, 51-78
- [17] Erdem, T. (1996). A dynamic analysis of market structure based on panel data. *Marketing Science* 15, 359-78
- [18] Financial Times (1999). Brazil's regional drinks makers slake thirst for value: The tax regime and growing demand have penalised leading brands. Edition of June 16, 1999
- [19] Galanz (2008). Galanz household electric appliances: the unruly competition in the market will hopefully end. 'Corporate News' posted on the company's website, dated September 1, 2008
- [20] Guadagni, P. and J. Little (1983). A logit model of brand choice calibrated on scanner data. *Marketing Science* 2, 203-38

- [21] Hausman, J. A., G. K. Leonard and J. D. Zona (1994). Competitive analysis with differentiated products. *Annales D'Economie et de Statistique* 34, 159-80
- [22] Hurwitz, M. A. and R. E. Caves (1988). Persuasion or information? Promotion and the shares of brand name and generic pharmaceuticals. *Journal of Law and Economics* 31, 299-320
- [23] Kalouptsi, M. (2010). From market shares to consumer types: Duality in differentiated product demand estimation. *Journal of Applied Econometrics*.
- [24] Keane, M. P. (1997). Modeling heterogeneity and state dependence in consumer choice behavior. *Journal of Business & Economic Statistics* 15, 310-27
- [25] Nair, H. (2007). Intertemporal price discrimination with forward-looking consumers: Application to the US market for console video-games. *Quantitative Marketing and Economics* 5, 239-92
- [26] Neri, M. (1995). Sobre a mensuração dos salários reais em alta inflação. (About the measurement of real wages under high inflation, in Portuguese.) *Pesquisa e Planejamento Econômico* 25, 497-525
- [27] Nevo, A. (2000). A practitioner's guide to estimation of random coefficients logit models of demand. *Journal of Economics & Management Strategy* 9, 513-548
- [28] Salvo, A. (2009). Cut-throat fringe competition in an emerging country market: Tax evasion or the absence of market power? *Journal of Industrial Economics* 57, 677-711
- [29] Salvo, A. (2010). Inferring market power under the threat of entry: The case of the Brazilian cement industry. *RAND Journal of Economics* 41, 326-350
- [30] Sancheti, S. and K. Sudhir (2009). Education consumption in an emerging market. Yale University, Mimeo
- [31] Scott Morton, F. M. (2000). Barriers to entry, brand advertising, and generic entry in the US pharmaceutical industry. *International Journal of Industrial Organization* 18, 1085-1104
- [32] SuperHiper (2003). Survey presents consumer purchasing habits at retail. Trade publication, edition of 03/2003
- [33] Wall Street Journal (2011). A rags-to-riches career highlights Latin resurgence. Front page article in the print edition of November 15, 2011

A Data sources and treatment [TO BE COMPLETED]

Provide IBOPE proportions over time by region; explain urban areas with population 20,000+ (later 10,000+); explain IBOPE proportions times IBGE PNAD’s total number of urban households; linear interpolation to increase the frequency of the resulting series from annual to monthly / bimonthly frequency.

HEX 95/96: Carvalho Filho and Chamon (2011) discuss this survey in detail. List brand descriptions that we entered as “premium” and “generic”;³⁹ describe “majority volume share” rule (including proportion of households consuming both premium and generic); describe how we map the cities surveyed in the HEX to the Nielsen regions; mention the good fit between HEX and IBOPE household data; discuss the intensive margin of household soda consumption.

Describe additional data sources: advertising, meteorology, factor prices.

B Additional details

B.1 Dynamic type evolution

Provide an example, as stated following Assumption 1 (TO BE COMPLETED).

B.2 The GMM objective

Steps 1-6 of the estimation algorithm detailed in subsection 3.3.1 yield a $N \times 1$ vector $\delta(\theta_2)$ (where $N = 7 \cdot 9 \cdot 57$) containing the base-utility levels of all brands in all regions in all time periods, given any generic value of the nonlinear parameters θ_2 . Conditioning on the full parameter vector $\theta = (\theta_1, \theta_2)$, one can obtain a vector $\xi(\theta)$, of the same dimension, that captures the base-utility unobservables. This is done by subtracting the systematic portion of the base-utility, for any brand j in any market gt :

$$\xi_{jgt} = \delta_{jgt} - x_{jgt}\beta - \alpha \cdot p_{jgt}$$

³⁹For perspective, examples of the former include the brand descriptions (codes): *Coca-Cola* (9301), *Pepsi* (9302), *Guaraná* (9303), *Fanta laranja, wa, limão* (9304), *Soda limonada* (9307), *Mirinda* (9308), *Sukita* (9315), *Pop laranja* (9316) and *Refrigerante água tônica* (9349). Examples of the latter include: *Refrigerante tubaína* (9318), *Refrigerante laranja exceto Fanta, Sukita, Pop, Crush* (9339), *Refrigerante cola exceto Coca-Cola e Pepsi-Cola* (9340), *Refrigerante cajú qualquer marca* (9346) and *Refrigerante Goianinha* (9355). There were four ambiguous brand descriptions (codes): *Refrigerante água natural* (9310), *Refrigerante gasosa* (9319), *Refrigerantes não especificado* (9335) and *Refrigerante dietético* (9360). These ambiguous codes were designated premium status among households for whom the identifiable share of premium in total soda expenditure exceeded 50% and the identifiable share of generics was lower than 10% (and, similarly, with respect to designating these ambiguous codes generic status). Soda expenditure that remained ambiguous was then allocated in proportion to the household’s expenditure on identified premium brands and identified generic brands. Finally, for the few households who purchased soda and no code was identified, we allocated expenditure among premium and generic according to (identified) expenditure shares within the same socioeconomic group and city.

and stacking all these unobservables together. In what follows, it would be useful to denote the dimensions of θ_1 and θ_2 by K_1 and K_2 , respectively, and write this as:

$$\xi(\theta) = \delta(\theta_2) - X\theta_1 \tag{5}$$

where X is a $N \times K_1$ matrix containing the base-utility covariates (notice that one of them is the price).

Let Z be a $N \times L$ matrix of instruments. Our instruments contain all the elements of X except for price, as well as constructed instruments (see subsection 3.3.2 for a full description of the three classes of instruments that we use) where $L > (K_1 + K_2)$. We let z'_i be a row vector with the instrument values corresponding to observation i .⁴⁰

$$Z = \begin{bmatrix} z'_1 \\ \vdots \\ z'_N \end{bmatrix}_{N \times L}$$

Let $W_{L \times L} = (Z'Z)^{-1}$. The GMM objective is defined by:

$$Q_N(\theta) = \xi(\theta)'ZWZ'\xi(\theta)$$

A substantial reduction in computation time is achieved by noting (see BLP 1995, Nevo 2000) that, conditional on a guess for θ_2 , there is a closed-form solution for the parameters θ_1 that minimize the objective:

$$\theta_1^*(\theta_2) = (X'ZWZ'X)^{-1}X'ZWZ'\delta(\theta_2)$$

B.3 Robustness

[TO BE COMPLETED]

⁴⁰To be clear, an “observation” here is a particular brand-region-period combination.

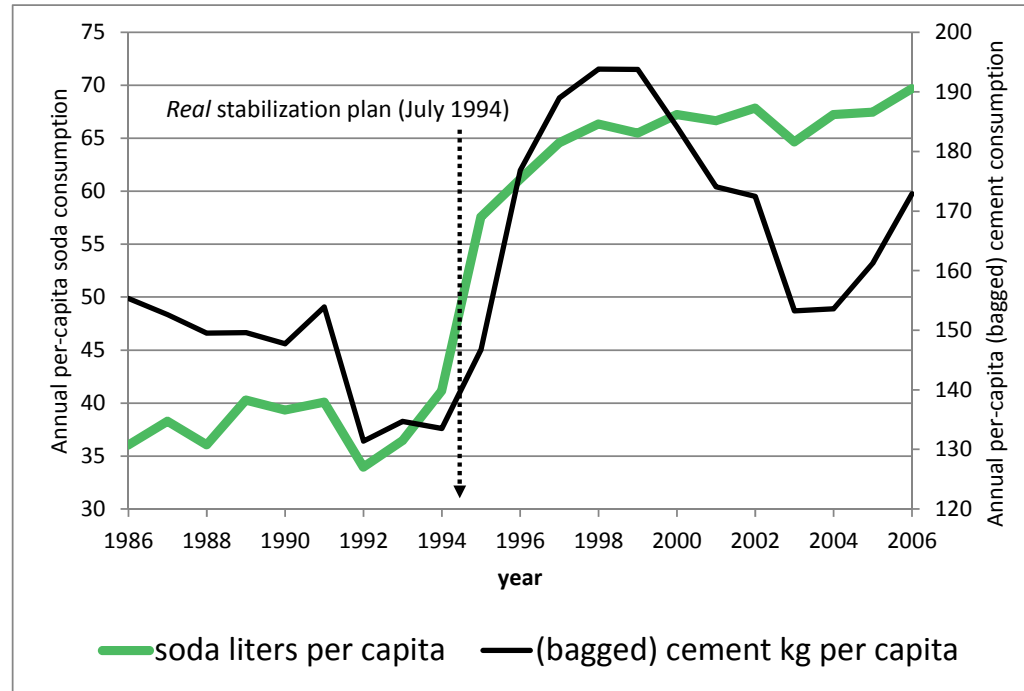


Figure 1: Annual aggregate per capita consumption of soft drinks (in liters per person) and of bagged cement (in kg per person). (Cement sold in bags, as opposed to sales in bulk, filter out large-scale construction activity, such as government spending on infrastructure.) Source: Brazilian trade associations for soft drink makers and for cement producers, ABIR and SNIC respectively.

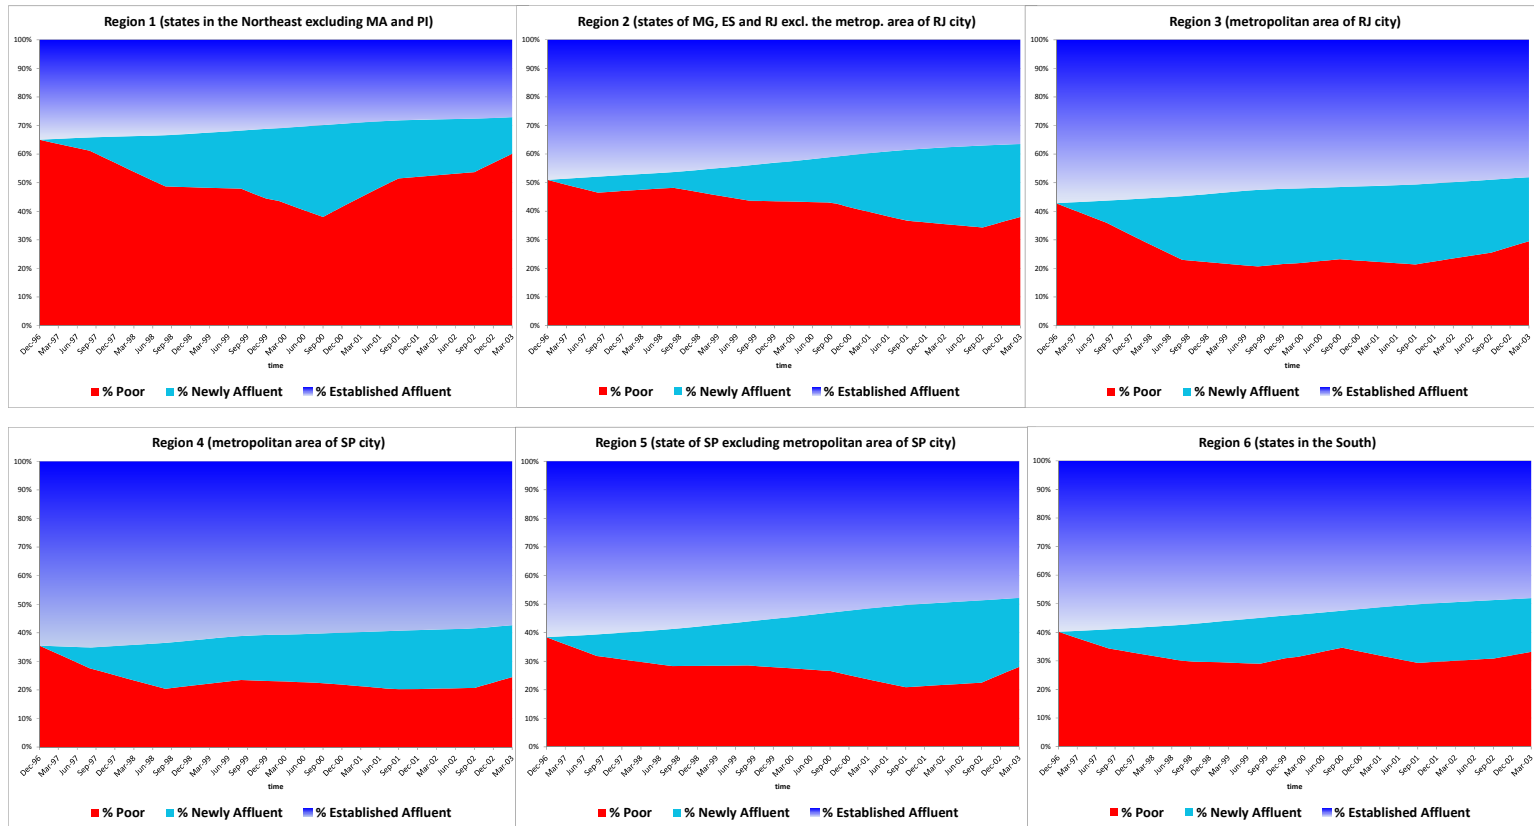


Figure 2: The rise of “newly affluent” households: Proportion of urban households in each of three constructed socioeconomic groups (“Established Affluent,” “Newly Affluent” and “Poor”), as defined in the text, by region in the period Dec-96 to Mar-03. The smallest region by number of households, region 7 (Federal District and states of GO and MS), is not shown for lack of space (the pattern is similar to region 2). Source: IBOPE LatinPanel and IBGE PNAD.




Figure 3: The evolution of quantities, prices and shares by type of brand (Premium versus Generic), for soda sold in family-size bottles through the self-service channel across the seven Nielsen regions, in the period Dec-96 to Mar-03. Upper left (panel): monthly total quantities (in million liters); upper right: mean share-weighted prices (in constant R\$/liter); lower left: volume shares of soda category (summing to one); lower right: mean shares of six liters per week per urban household, as defined in the text. Source: Nielsen and IBGE PNAD.

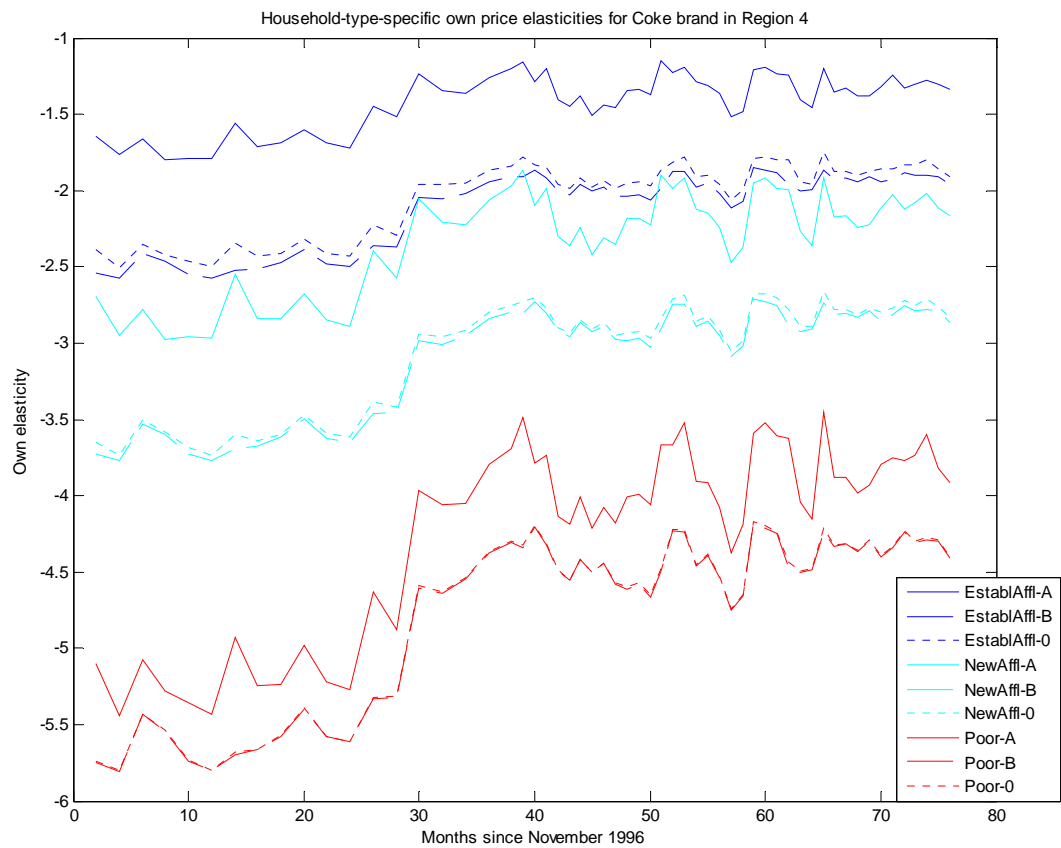


Figure 4: Evolution of own-price elasticities for Coke brand, by household type, in region 4 (São Paulo Metro). Source: Baseline model (Brand Type Persistence).

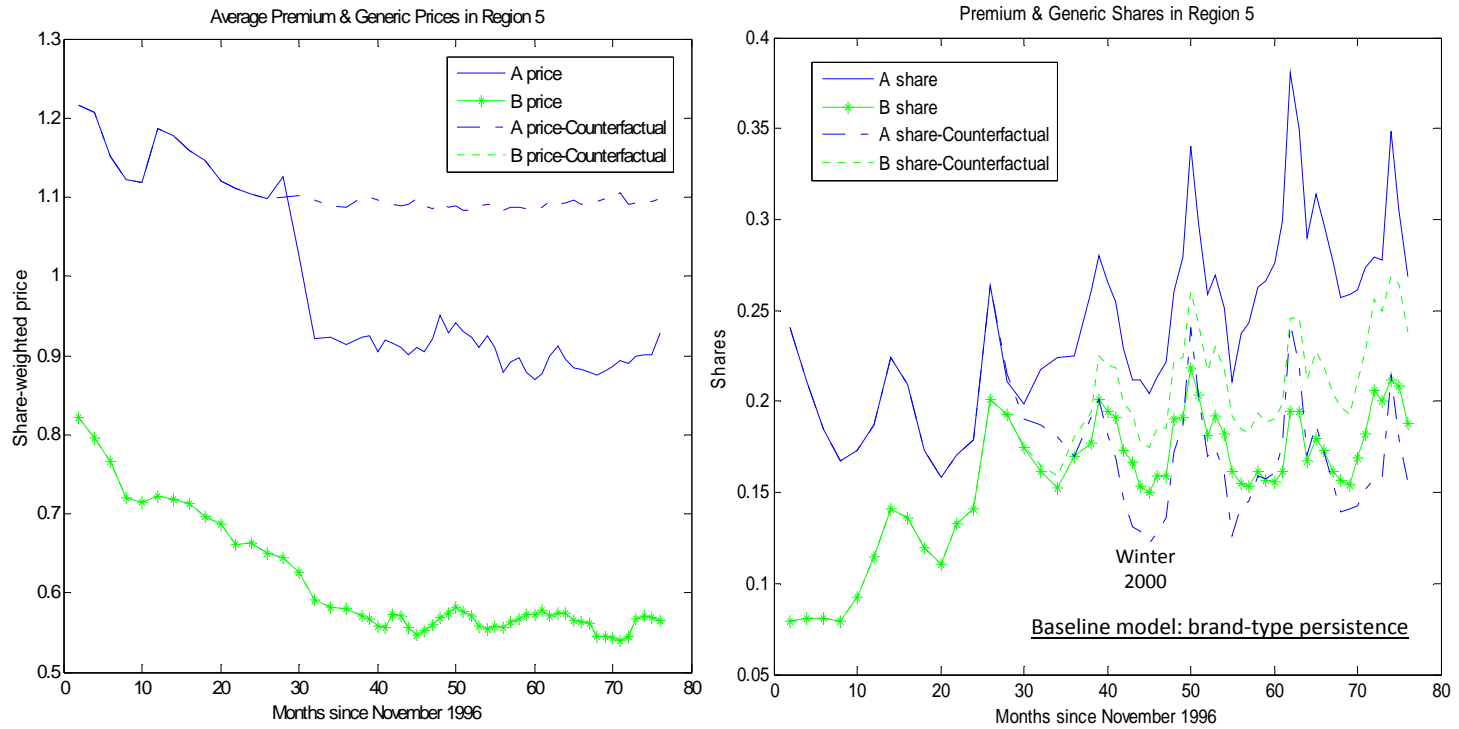


Figure 5: Actual against counterfactual price and share paths for premium (“A”) brands and generic (“B”) brands in region 5 (São Paulo Interior). Prices in the left panel and shares in the right panel. The counterfactual scenario considers premium brands not cutting prices in mid 1999. Source: Baseline model (Brand Type Persistence).

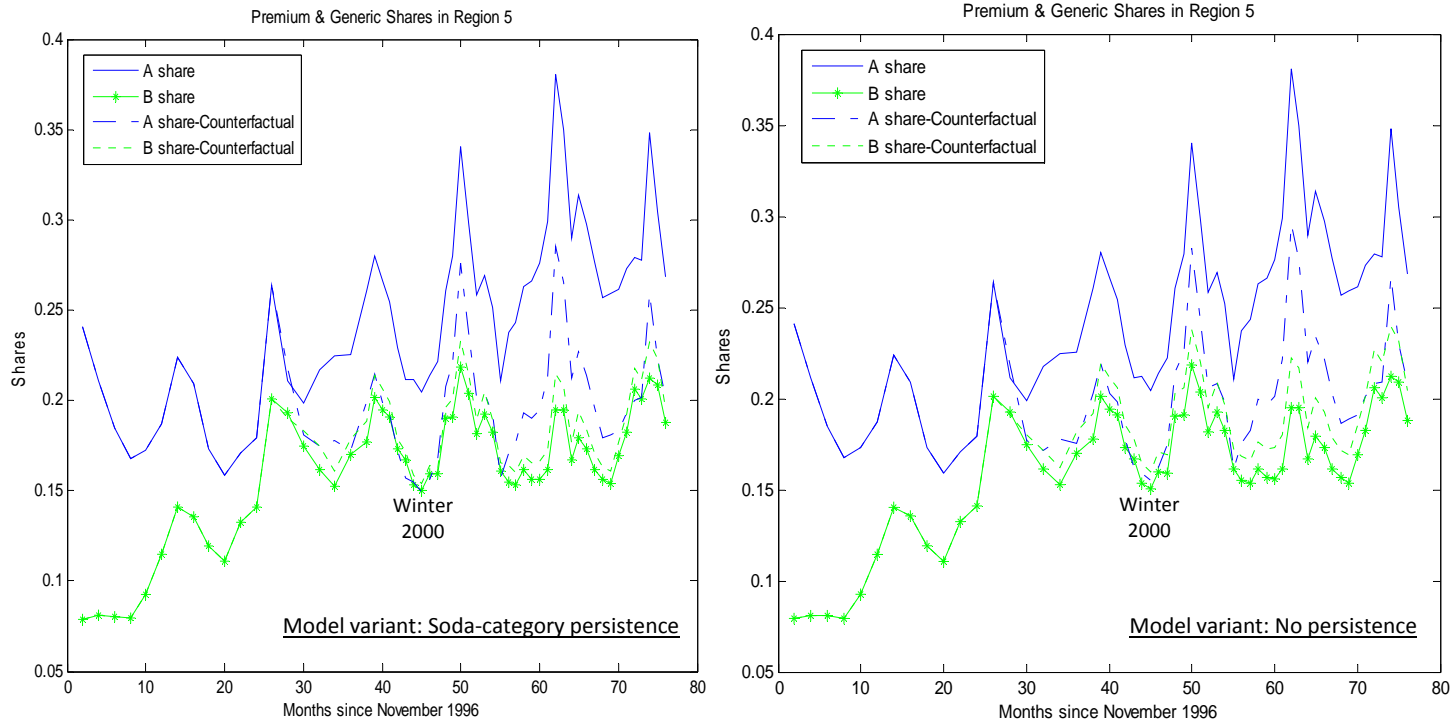


Figure 6: Share paths for premium (“A”) brands and generic (“B”) brands, in region 5, for the same counterfactual experiment of the earlier figure (premium brands not cutting prices in 1999), but employing each of the two model variants. Sources: Soda Category Persistence model variant in the left panel, No Persistence model variant in the right panel.

Region of survey. cities	Socioeconomic group	Households \times 1000		Soda purchasing	By brand type		No soda
		Universe	%		Premium	Generic	
1 (Northeast)	ABC (“Rich”)	696	36	28.0%	27.0%	0.9%	72.0%
	DE (“Poor”)	1230	64	9.1%	8.3%	0.8%	90.9%
2 (MG, ES, RJ interior)	ABC	529	57	39.9%	37.9%	2.0%	60.1%
	DE	404	43	23.2%	22.1%	1.2%	76.8%
3 (RJ Metro)	ABC	1625	55	31.9%	31.6%	0.3%	68.1%
	DE	1331	45	18.3%	18.3%	0.0%	81.7%
4 (SP Metro)	ABC	2586	60	34.5%	33.1%	1.4%	65.5%
	DE	1689	40	19.8%	17.3%	2.6%	80.2%
6 (South)	ABC	955	63	43.2%	42.5%	0.7%	56.8%
	DE	559	37	20.4%	20.1%	0.3%	79.6%
7 (DF, GO MS)	ABC	428	61	36.5%	34.4%	2.1%	63.5%
	DE	270	39	23.6%	21.1%	2.5%	76.4%
Total above	ABC (“Rich”)	6819	55	35.0%	33.9%	1.1%	65.0%
	DE (“Poor”)	5482	45	17.5%	16.3%	1.2%	82.5%

Table 1: The extensive margin of soda consumption inside the home by different socioeconomic groups in 1995/96. Socioeconomic groups are defined per the points scale used by IBOPE. Cities surveyed were: (Region 1) Recife, Fortaleza and Salvador; (region 2) Belo Horizonte; (region 3) Rio de Janeiro Metro; (region 4) São Paulo Metro; (region 6) Curitiba and Porto Alegre; (region 7) Brasília and Goiânia. No city was surveyed in region 5 (state of São Paulo excluding São Paulo Metro). We do not consider the northern city of Belém as it is located outside the area covered by Nielsen. Source: IBGE HEX (POF) 1995/96.

Baseline model: Brand Type Persistence (BTP)	[1] coeff	(s.e.)	[2] coeff	(s.e.)	[3] coeff	(s.e.)
A. Parameter Estimates						
Estimates: θ_2						
α_{EA}	1.882***	(0.441)	2.401***	(0.913)	2.291***	(0.806)
α_{NA}	0.671	(0.557)	1.515**	(0.684)	1.663***	(0.571)
λ	3.526***	(0.236)	3.842***	(0.358)	2.574***	(0.395)
Estimates: θ_1						
α	-3.466***	(0.494)	-4.344***	(0.960)	-3.323***	(0.827)
Dec/Jan seasonal (summer begins)	0.278***	(0.032)	0.302***	(0.040)	0.234***	(0.028)
...Feb/Mar, Apr/May, Jun/Jul...	Yes		Yes		Yes	
Aug/Sep seasonal	-0.002	(0.020)	0.020	(0.020)	0.010	(0.015)
Temperature	0.058***	(0.006)	0.076***	(0.007)	0.056***	(0.005)
Region 1-specific advertising GRPs	0.547**	(0.245)	0.522**	(0.234)	0.011	(0.075)
...Regions 2,3,4,5,6-specific GRPs...	All +ve		All +ve		Most +ve	
Region 7-specific advertising GRPs	0.321**	(0.140)	0.291**	(0.142)	0.114	(0.080)
Brand-and-region fixed effects	Yes		Yes		Yes	
Region time trends	No		Yes		No	
Brand-and-region time trends	No		No		Yes	
B. Economic Implications of Estimates: Elasticities (columns correspond to above specifications)						
Aggregate own-price elasticities, $\eta_{jj,gt}$:						
Coke (premium, \mathcal{A})	-1.440		-1.638		-1.058	
Guaraná Antarctica (premium, \mathcal{A})	-1.685		-1.958		-1.173	
Fanta (premium, \mathcal{A})	-1.640		-1.905		-1.141	
Pepsi (premium, \mathcal{A})	-1.635		-1.900		-1.136	
B brands (fringe of generics, \mathcal{B})	-0.743		-0.751		-0.659	
HH-type specific own elasticit., $\eta_{jj,r,gt}$:						
Coke, EA^A	-1.094		-1.325		-0.763	
Coke, EA^B	-1.684		-2.074		-1.077	
Coke, EA^O	-1.611		-1.990		-1.030	
Coke, NA^A	-2.309		-2.191		-1.354	
Coke, NA^B	-2.995		-3.033		-1.753	
Coke, NA^O	-2.960		-2.985		-1.713	
Coke, P^A	-3.128		-4.085		-3.282	
Coke, P^B	-3.724		-4.678		-3.565	
Coke, P^O	-3.705		-4.666		-3.555	
B brands, EA^A	-0.981		-1.208		-0.627	
B brands, EA^B	-0.266		-0.278		-0.265	
B brands, EA^O	-0.914		-1.130		-0.580	
B brands, NA^A	-1.728		-1.755		-1.007	
B brands, NA^B	-0.718		-0.572		-0.502	
B brands, NA^O	-1.678		-1.693		-0.959	
B brands, P^A	-2.145		-2.700		-2.037	
B brands, P^B	-1.095		-1.444		-1.451	
B brands, P^O	-2.112		-2.676		-2.018	

Table 2: Results in the baseline model: Brand Type Persistence. All specifications include Coke's price reduction, five sets of Hausman price means, and factor prices as instruments. Reported elasticities are means across region-and-time markets, and only a few elasticities are shown due to space constraints. * $p < .1$, ** $p < .05$, *** $p < .01$

Variante 1: Soda Category Persistence (SCP)	[1] coeff	(s.e.)	[2] coeff	(s.e.)	[3] coeff	(s.e.)
A. Parameter Estimates						
Estimates: θ_2						
α_{EA}	1.221**	(0.517)	2.116***	(0.749)	2.312***	(0.760)
α_{NA}	0.109	(0.510)	0.644	(0.616)	1.387***	(0.537)
λ	2.027***	(0.616)	2.302***	(0.499)	2.067***	(0.436)
Estimates: θ_1						
α	-2.536***	(0.562)	-3.691***	(0.808)	-3.218***	(0.806)
Dec/Jan seasonal (summer begins)	0.202***	(0.032)	0.235***	(0.037)	0.229***	(0.030)
...Feb/Mar, Apr/May, Jun/Jul...	Yes		Yes		Yes	
Aug/Sep seasonal	-0.003	(0.019)	0.012	(0.019)	0.012	(0.015)
Temperature	0.049***	(0.007)	0.066***	(0.007)	0.058***	(0.006)
Region 1-specific advertising GRPs	.00006**	(.00003)	.00006**	(.00003)	.00000	(.00001)
...Regions 2,3,4,5,6-specific GRPs...	All +ve		All +ve		Most +ve	
Region 7-specific advertising GRPs	.00003***	(.00001)	.00003**	(.00001)	.00001	(.00001)
Brand-and-region fixed effects	Yes		Yes		Yes	
Region time trends	No		Yes		No	
Brand-and-region time trends	No		No		Yes	
B. Economic Implications of Estimates: Elasticities (columns correspond to above specifications)						
Aggregate own-price elasticities, $\eta_{jj,gt}$:						
Coke (premium, \mathcal{A})	-1.507		-1.660		-1.059	
Guaraná Antarctica (premium, \mathcal{A})	-1.561		-1.770		-1.119	
Fanta (premium, \mathcal{A})	-1.507		-1.718		-1.088	
Pepsi (premium, \mathcal{A})	-1.493		-1.705		-1.082	
B brands (fringe of generics, \mathcal{B})	-0.918		-1.063		-0.705	
HH-type specific own elasticit., $\eta_{jj,r,gt}$:						
Coke, EA^A and Coke, EA^B	-1.118		-1.308		-0.764	
Coke, EA^O	-1.309		-1.561		-0.896	
Coke, NA^A and Coke, NA^B	-2.319		-2.937		-1.705	
Coke, NA^O	-2.546		-3.218		-1.903	
Coke, P^A and Coke, P^B	-2.445		-3.711		-3.293	
Coke, P^O	-2.668		-3.938		-3.440	
B brands, EA^A and B brands, EA^B	-0.635		-0.742		-0.432	
B brands, EA^O	-0.757		-0.902		-0.517	
B brands, NA^A and B brands, NA^B	-1.208		-1.480		-0.885	
B brands, NA^O	-1.449		-1.829		-1.080	
B brands, P^A and B brands, P^B	-1.270		-1.869		-1.691	
B brands, P^O	-1.518		-2.247		-1.963	

Table 3: Results in the first model variant: Soda Category Persistence. All specifications include Coke's price reduction, five sets of Hausman price means, and factor prices as instruments. Reported elasticities are means across region-and-time markets, and only a few elasticities are shown due to space constraints. * $p < .1$, ** $p < .05$, *** $p < .01$

Variant 2: No Persistence (NP)	[1] coeff	(s.e.)	[2] coeff	(s.e.)	[3] coeff	(s.e.)
A. Parameter Estimates						
Estimates: θ_2						
α_{EA}	0.414	(0.444)	1.212*	(0.630)	0.681	(0.907)
α_{NA}	0.107	(0.316)	-0.223	(0.396)	0.364*	(0.207)
λ (constrained to zero)	-	-	-	-	-	-
Estimates: θ_1						
α	-1.975***	(0.353)	-2.652***	(0.571)	-1.603**	(0.650)
Dec/Jan seasonal (summer begins)	0.187***	(0.019)	0.189***	(0.022)	0.183***	(0.017)
...Feb/Mar, Apr/May, Jun/Jul...	Yes		Yes		Yes	
Aug/Sep seasonal	-0.014	(0.018)	-0.002	(0.018)	0.000	(0.014)
Temperature	0.046***	(0.004)	0.055***	(0.005)	0.048***	(0.003)
Region 1-specific advertising GRPs	0.00005**	(0.00003)	0.00005**	(0.00003)	-0.00000	(0.00001)
...Regions 2,3,4,5,6-specific GRPs...	All +ve		All +ve		Most +ve	
Region 7-specific advertising GRPs	0.00003***	(0.00001)	0.00003***	(0.00001)	0.00001*	(0.00001)
Brand-and-region fixed effects	Yes		Yes		Yes	
Region time trends	No		Yes		No	
Brand-and-region time trends	No		No		Yes	
B. Economic Implications of Estimates: Elasticities (columns correspond to above specifications)						
Aggregate own-price elasticit., $\eta_{jj,gt}$:						
Coke (premium, \mathcal{A})	-1.701		-1.764		-1.124	
Guaraná Antarctica (premium, \mathcal{A})	-1.682		-1.779		-1.115	
Fanta (premium, \mathcal{A})	-1.611		-1.716		-1.070	
Pepsi (premium, \mathcal{A})	-1.585		-1.698		-1.055	
B brands (fringe of generics, \mathcal{B})	-0.973		-1.099		-0.659	
HH-type specific own elasticit., $\eta_{jj,r,gt}$:						
Coke, $r \in \{EA^A, EA^B, EA^O\}$	-1.509		-1.343		-0.885	
Coke, $r \in \{NA^A, NA^B, NA^O\}$	-1.850		-2.967		-1.221	
Coke, $r \in \{P^A, P^B, P^O\}$	-1.971		-2.710		-1.618	
B brand, $r \in \{EA^A, EA^B, EA^O\}$	-0.860		-0.777		-0.506	
B brand, $r \in \{NA^A, NA^B, NA^O\}$	-1.042		-1.648		-0.689	
B brand, $r \in \{P^A, P^B, P^O\}$	-1.107		-1.506		-0.905	

Table 4: Results in the second model variant: No Persistence. All specifications include Coke's price reduction, five sets of Hausman price means, and factor prices as instruments. Reported elasticities are means across region-and-time markets, and only a few elasticities are shown due to space constraints. * $p < .1$, ** $p < .05$, *** $p < .01$

	Soda penetration	Premium share	Generic share	Premium:Generic ratio
Baseline (Brand Type Persistence)				
Established Affluent	54.1%	36.9%	17.2%	2.1
Newly Affluent	25.4%	13.5%	11.8%	1.1
Poor	6.8%	2.6%	4.1%	0.6
Variant 1 (Soda Category Persistence)				
Established Affluent	55.0%	36.9%	18.0%	2.0
Newly Affluent	17.8%	9.6%	8.2%	1.2
Poor	10.2%	5.1%	5.1%	1.0
Variant 2 (No Persistence)				
Established Affluent	45.0%	30.6%	14.4%	2.1
Newly Affluent	20.8%	11.6%	9.3%	1.2
Poor	23.1%	13.3%	9.8%	1.4

Table 5: Soda penetration, for the entire category and by type of brand, in each socioeconomic group, in the baseline model and in the two model variants. Means across region-and-time markets.