

The Rise of Fringe Competitors In the Wake of an Emerging Middle Class: An Empirical Analysis*

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

Alberto Salvo[‡]
Department of Economics
National University of Singapore

August 2014

Abstract

The “emerging middle class” has become a force of economic importance in many consumer markets around the globe. A striking phenomenon in some of these markets is the growth of “generic,” low-price brands that compete with established premium brands over the expanding consumer purchasing power. In this paper we examine the impact of a substantial rise in Brazil’s living standards on the development of the country’s large soft-drink market, during a six-year period that saw unprecedented growth in the share of generic soda brands. Our demand study draws on data sources that capture both social mobility and market outcomes. Our analysis suggests that the emergence of a price-sensitive, new middle class aided the growth of the fringe. Our estimated demand model rationalizes a drastic price cut by Coca-Cola that allowed it to contain the fringe’s growth.

Keywords: emerging middle class, social mobility, differentiated-product demand, habit formation, generics, competitive fringe, premium brands

JEL Classification: L10, D12, O12

*A previous draft was circulated under the title “Grab Them Before They Go Generic: Habit Formation and the Emerging Middle Class.” 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 Itai Ater, Michael Beenstock, David Besanko, Eric Bradlow, Sofronis Clerides, Allan Collard-Wexler, Pierre Dubois, David Genesove, Matt Gentzkow, Paul Grieco, Noam Gruber, Phil Haile, Günter Hitsch, Kosin Isariyawongse, Dan Keniston, Saul Lach, Katja Seim, Olexander Scherbakov, Scott Stern, Yuya Takashi, and Maria Ana Vitorino, as well as participants at numerous conferences and seminars for valuable comments. All errors are our own.

[†]aloneiz@mssc.huji.ac.il

[‡]albertosalvo@nus.edu.sg

“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

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, experience a socioeconomic transformation: 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 Bolstering the demand for consumer goods, these “new consumers” provide a potential engine of growth for the global economy. This motivates studying both the nature of this new demand, and its implications for competition in emerging consumer markets.

One might expect established brands to be the primary beneficiaries from the increased purchasing power brought about by this socioeconomic transition. Nonetheless, in many relevant markets, we observe the rapid growth of a fringe of small producers, sometimes referred to as “B-brands.” To provide a few examples, 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). The Economist (2012) counts 100 “domestic carmakers” in China. Abbott India’s brands Digene, Eptoin and Cremaffin face competition from 211, 327 and 242 “regional generics,” respectively.² Gemma Ware (2013) documents that in sub-Saharan Africa’s markets, and in food and beverage in particular, B-brands have gained substantial ground, arguing that “...demand is driven by the rise of Africa’s middle classes.”

A unifying feature of the B-brands’ competitive strategy is the focus on low prices rather than on extensive advertising and branding efforts. Such brands also tend to operate locally, tailoring their marketing to the idiosyncratic characteristics of a relatively small geographic area. This stands in contrast to established brands (or “A-brands”) that invest heavily in advertising and brand-building, often on a national or even a global scale. B-brands pose a challenge to

¹The Economist (2011a) 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.” Nomura Bank states that by 2014 Indonesia should boast almost 150m “newly affluent Indonesians (who) are certainly spending” (The Economist, 2011b). Ferreira et al (2012) calculate that “at least 43% of all Latin Americans changed social classes between the mid 1990s and the end of the 2000s.” The emerging middle class is also referred to as the “lower-middle class,” to emphasize its vulnerability to income shocks, or to the risk of “slipping back,” as recently warned by World Bank chief economist Kaushik Basu—see Financial Times (2014).

²As shared by the company during a corporate presentation in late 2011.

established brands in emerging markets, as argued recently by executives of a global “fast-moving consumer goods” firm, stating 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.”³

Our paper explores potential explanations for the success of B-brands in emerging markets. We do so by empirically analyzing an important test case: the Brazilian market for carbonated soft drinks (or “soda”). We study the evolution of this market from December 1996 through March 2003, a six-year period over which two striking phenomena were evident: a substantial expansion in demand fueled by rising living standards, and the rapid growth of a competitive fringe of soda producers.

Brazil’s large soda market trails only the United States and Mexico by volume. Following a successful economic stabilization plan in 1994, aggregate soda consumption almost doubled by 1997, and continued to grow at an annual rate of about 10% through 1999. This growth was driven, at least in part, by pronounced upward mobility among lower income households, who were no longer subject to “inflation tax.”⁴ In 1999, the Financial Times reported that “the increased purchasing power that came with stable prices... allowed about 25m new consumers into the (soft-drink) market,” among a total population of 170m at the time. Other markets, ranging from fresh meat to refrigerators and television sets, saw similar expansions in demand.

The Brazilian market features established (or “premium”) soda brands marketed by two leading firms, the Coca-Cola Company (hereafter Coca-Cola) and Ambev, who in 1996 jointly accounted for almost 90% of Brazilian soda expenditure.⁵ As they attempted to tap into the new and expanding demand for soda, Coca-Cola and Ambev faced increased competitive pressure from hundreds of regionally focused B-brands, which we also label “generics.” Between 1996 and 1999, the combined volume share of these B-brands doubled from 20% to 40%.⁶ The generics’ marketing efforts focused on securing shelf space via low prices, in contrast to the dominant duopoly’s heavy advertising. 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).

Having kept prices broadly constant during the preceding years of the fringe’s expansion, in 1999 Coca-Cola abruptly cut prices across its brands by over 20%, a move that was soon matched by Ambev. Following this price cut, the growth in the market share of the generic fringe was

³The statement was made while meeting one of the authors in Delhi in late 2011. To be clear, all words—including the terms in italics—are the executives’ own, though in slightly rearranged order without modifying context.

⁴A substantial mass of households with no access to inflation-indexed bank accounts were the main beneficiaries of the taming of chronically high inflation: “...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.

⁵Ambev distributed the Pepsi brand, and is now part of the AB Inbev group.

⁶This pertains to the dominant market segment of family-size bottles sold through the “self-service outlets” distribution channel (supermarkets with checkouts) in urban areas.

halted. Nonetheless, the fringe continued to command a substantial market share.

Several factors might have aided the fringe including, on the supply side, changes in technology and distribution. Our primary focus in this paper, however, is on the demand side, and, specifically, the potential link between the demand expansion brought about by the emerging middle class, and the B-brands' success. We empirically investigate this relationship. Exploiting cross-sectional as well as temporal data variation, we first present empirical evidence that the generics' volume share is positively associated with the fraction of a regional market's population that belongs in the emerging middle class. We then go beyond the finding that the two phenomena are linked, and attempt to provide insights regarding the *nature of the relationship*. What is it about an emerging middle class that fosters the growth of a fringe of generic producers?

Our analysis emphasizes the possibility that emerging middle class consumers' demand was more elastic than that of the established middle class. To stay with the Financial Times' analysis, "(t)he new (soda) customers...had different priorities...(t)hey were less concerned about expensive TV ads and more interested in value." A price-sensitive, expanding consumer segment may help explain both generic brands' growth, and the premium sellers' price cut. Whether such differences in price sensitivity are present is an empirical question, which we address by estimating a structural model of demand for soft drinks in the Brazilian market.

In this model, consumers are segmented by their exogenous socioeconomic standing. They are classified into three groups: "Poor," "Established Affluent," and "Newly Affluent." Established affluent households are those who were already affluent before the process of upward mobility began, whereas newly affluent households represent the new middle class. The model allows these groups to be characterized by different price-sensitivity parameters and, therefore, allows us to explore whether the newly affluent were more price sensitive than the established affluent.

The model is estimated using panel data that tracks both market outcomes (brand prices and market shares), and aggregate socioeconomic mobility in a cross-section of regional markets over the sample period. The model captures key aspects of social mobility: it allows poor households to move up to newly affluent status, while also allowing downward mobility of newly affluents to poor status. Such downward mobility is observed during a recession in the later years of the sample. The model also accounts for urbanization, an important demographic shift.

Another ingredient we incorporate into the model is persistence in demand. We present features of the raw data that are consistent with such persistence. In particular, during the recessionary period, households fell back from newly affluent status to the ranks of the poor, yet soda consumption did not fall (nor did prices). The mostly aggregate data we have access to make it difficult to identify the exact persistence mechanism separately from other plausible sources of serial correlation. We do find, however, that controlling for a specific form of persistence in demand is helpful in terms of identifying demand elasticities.

Our baseline model postulates that the utility from consuming a generic (resp., premium) brand is affected by previous-period consumption of a generic (resp., premium) brand. We label this a “Brand Type Persistence” (BTP) model, and use it to explore the dichotomous nature of premium versus generic consumption. Such a mechanism may also play a role in driving the generics’ growth: if such persistence is important, then the emergence of a new middle class can be interpreted as the arrival into the market of millions of new customers who have not yet formed persistent consumption patterns. This may have aided generic firms in making inroads into this new consumer segment, and may have also provided a strong incentive for Coca-Cola to cut prices with the goal of defending its future market position.

Findings. Estimates from our baseline structural model indicate that newly affluent households are more price sensitive than established affluent households yet less price sensitive than the poor. For example, conditional on recent consumption of premium soda, raising the price of Coke by 1% lowers the demand for the brand among established affluent, newly affluent and poor households by 1.4%, 3.3% and 5.7%, respectively. Newly affluents are 1.4 times more likely to choose a generic brand over a premium brand, whereas for established affluents this ratio is only 0.5. The estimates also provide support for persistence in preferences.

A counterfactual analysis using our estimated baseline model predicts that, had premium brands failed to cut prices in mid 1999, their market shares would have declined substantially through 2003, resulting in diminished variable profits in the medium run. Our model, therefore, provides strong justification for the strategic price cut.

We further examine how variations in the magnitude of the price cut affect market outcomes. We find that a deeper price cut, by a magnitude of 30% rather than the observed 20%, would have been inferior, as it would have decreased variable profits. A less radical price cut of 10% would have been slightly more profitable, yet it would still have resulted in substantial market share losses, allowing the fringe to achieve market share parity with premium brands by 2003. In other words, compared to the 10% alternative, the 20% price cut slightly lowers profit but prevents the generic fringe from matching the premium market share, a situation Coca-Cola might have deemed undesirable.

It is conceivable that Coca-Cola’s revealed preference for the deeper 20% price cut reflects an objective function that considers not only the medium-run profits we calculate, but also market shares, as strategic variables that the firm wishes to protect. It is in this sense that our estimated model rationalizes the magnitude of Coca-Cola’s price cut. Our ability to draw such conclusions stems from the fact that we do not impose profit maximization, or assume any other forms of strategic behavior on the supply side, in estimating the model. Finally, we also find that changing the *timing* of the price cut would not substantially impact its effectiveness.

Related literature. Our study relates to several lines of research. Competition between

branded products and lower cost generics, particularly in pharmaceuticals, is examined by Chaudhuri, Goldberg and Jia (2006) in India, and Hurwitz and Caves (1988) and Scott Morton (2000) in the US. Another literature examines the relationship between the demographic composition of demand and prices, or inflation moderation, including Frankel and Gould (2001), Bils and Klenow (2004), Nevo and Hatzitaskos (2006), Lach (2007), and Calzolari, Ichino and Manaresi (2012). A recent literature studies demand by emerging-market consumers (e.g., Sancheti and Sudhir 2009). The empirical literature in economics and marketing has introduced habits or persistence into models of consumer choice (e.g., Eichenbaum, Hansen and Singleton 1988, Erdem 1996, Keane 1997, Shum 2004, Dubé, Hitsch, Rossi and Vitorino 2008). The role of persistent preferences in food and beverage has been emphasized in the literature (see Atkin 2013 and Bronnenberg, Dubé and Gentzkow 2012 for recent contributions).

The rest of the paper is organized as follows: Section 2 describes the data sources and the Brazilian market. Section 3 discusses several potential drivers of the fringe’s growth, and presents descriptive evidence regarding the relationship between this growth and the emergence of a new middle class. Section 4 develops our structural demand model, and explains how this model is identified and estimated. Section 5 reports our results, and Section 6 concludes.

2 Market and data

This study brings together data from three main sources. We now describe these sources and provide additional background on the industry.

2.1 Market-level data: the growth of the fringe and Coca-Cola’s response

We observe a panel of market-level data from Nielsen, consisting of total quantities and prices for soft-drink brands. There are $g = 1, \dots, 7$ regions and $t = 1, \dots, 57$ time periods, ranging from the December 1996-January 1997 bimonth to the March 2003 month (Nielsen raised the frequency of its bimonthly point-of-sale audits to a monthly basis in 2000). We therefore observe $7 \times 57 = 399$ region-period markets.

The seven geographic markets are urban and, as in Salvo (2009), we consider soft drinks sold through the “self-service” channel (supermarkets with checkouts) in the 2-liter family-size bottle. Our focus on this market segment is justified on several counts. First, the focus on urban areas is natural since more than 80% of Brazil’s population was urbanized by 1996 and, importantly, our framework allows for rural-to-urban migration. Second, urban households in Brazil perform most of their grocery shopping in supermarkets with checkouts, rather than in behind-the-counter “traditional” retail stores. Finally, sales of family-size bottles dominate those of “single-serve” (300ml) bottles or cans (mostly sold in bars and restaurants). Moreover, the competitive fringe,

whose success we wish to explain, was mostly present in the family-size bottle segment.

Also following Salvo (2009), we aggregate flavors and brands into $j = 1, \dots, 9$ brand-groups. These groups include eight premium brands: five brands of the Coca-Cola Company (Coke, Fanta, the guaraná-flavored Kuat, Diet Coke, and “Other Coca-Cola”), and three brands marketed by Ambev (Guaraná Antarctica, Pepsi, and “Other Ambev”). The ninth brand category is an aggregate of discount brands that form the generic fringe.⁷ Both premium and generic brands enjoyed substantial volume growth during the sample period, as Figure 1 indicates. The generic fringe grew much faster than the premium brands until Coca-Cola’s abrupt mid-1999 price cut. The figure also reveals seasonality for which we control in our demand model.

Table 1 describes volume shares (of the soda category) for each brand in each region at the first and last sample periods. Averaged arithmetically across regions, Coca-Cola’s brands accounted for a 50% share in the first period, with Coke being dominant, whereas Ambev enjoyed a 31% share, with Guaraná Antarctica and Pepsi as its flagship brands. From a mean 19% at the start of the sample to 40% at the end, growth in the generic share was stark and apparent in all regions, occurring at differential rates across them (e.g., being most pronounced in the relatively poor northeast). This variation is helpful in identifying our demand model. The table also reflects region-specific tendencies to consume particular brands, motivating the inclusion in the model of region-brand fixed effects.

The left panel of Figure 2 illustrates the evolution of (mean share-weighted) prices for premium brands and for generics in R\$ per liter (prices reported at constant Brazil CPI March 2003 terms, divide by 2 for rough US\$ values). Premium brands initially held prices broadly flat, at R\$ 1.15. In mid 1999, Coca-Cola cut its prices by more than 20%, a move that was soon matched by Ambev. The figure clearly indicates the sudden nature of this price cut, which we will exploit for identification purposes. Fringe prices, for their part, declined gradually but relentlessly, from R\$ 0.90 in late 1996 to R\$ 0.60 in late 2000. Falling generic prices are consistent with substantial entry and capacity expansion in the fringe, as competitive firms passed efficiency gains through to consumers. Fringe prices did not respond to the premium price cut in the sense that they did not deviate from their trend, consistent with competitive behavior.

It is worth noting that the brands we assign to Ambev were marketed under separate firms, Antarctica and Brahma, until their merger was announced in July 1999. One might ask whether Coca-Cola’s price cut, beginning in May 1999, was in part also a response to the soon-to-be announced Ambev merger. There is good reason to believe that this is not the case. First, Coca-Cola’s sudden price cut was directed at the family-size bottles sold through self-service outlets, precisely the business where the fringe had a growing presence. In contrast, Coca-

⁷The data provide limited information on the breakdown of this group into individual discount brands, as they are so numerous. Since Coca-Cola and Ambev are the prime users of these data, this aggregate structure is indicative of their perception of the fringe as a collection of small firms that do not offer substantial differentiation.

Cola’s single-serve prices in the bars channel, where the fringe had a much smaller presence, remained remarkably stable at around R\$ 3.1 throughout the sample. Indeed, had Coca-Cola been responding to the Ambev merger, one would expect prices in the bars channel to also change, since Antarctica and Brahma’s joint soft-drink share was relatively higher in bars than in supermarkets. Antarctica and Brahma were the country’s leading brewers and their merger was likely driven by developments in the fast-consolidating global beer industry, coupled with a soda business which—like Coca-Cola’s—faced growing competition from generics.⁸

As the right panel of Figure 2 shows, after 30 months of generics gaining share at the expense of premium brands, the premium price cut had a clear and immediate impact. It put an end to the staggering generic share growth, leading to stable premium and generic shares through the end of the sample. An interesting aspect of this process is that a sudden, permanent cut in premium prices led to a dramatic change in the volume share *trend*. In Section 3, we return to this pattern and discuss its likely relationship with an emerging middle class.

2.2 Data on aggregate social mobility

To track the undercurrent of social mobility in the Brazilian economy, we rely on the proprietary LatinPanel survey from IBOPE, a leading provider of data on consumer demographics.⁹ The survey, widely used by marketing practitioners, profiles urban households in Brazil’s different regions based on their expenditure on durable goods and services (e.g., ownership of a refrigerator, numbers 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 group.” The data we accessed provides the proportion of urban households pertaining to each of three groups, **AB**, **C** or **DE**, respectively with “high,” “intermediate” or “low” levels of affluence, in each of seven geographic regions over the period 1994-2006.¹⁰

The IBOPE data indicate that the demographic composition of urban households: (i) was stable between 1994 and 1996; (ii) displayed strong upward mobility from **DE** to **ABC** (i.e., {AB,C}) status between 1996 and 2000; and (iii) experienced a partial reversal of this upward mobility thereafter, consistent with a recessionary period. In aggregate, the proportion of **DE** households fell from 50% in 1996 to 33% in 2000, then rose to 44% by 2003 (conversely, the **AB** proportion rose from 19% in 1996 to 33% in 2000, then fell to 23% by 2003).

⁸Further, the policy concern with mergers is that merging firms and rivals may raise prices, rather than cut them. Coca-Cola’s response to increased competition from another premium brand (Pepsi) typically operates via aggressive advertising campaigns, not via price cuts. Contemporaneous accounts indeed describe the price cut as a response to the fringe’s growth. For example, Gertner et al (2004) report that “Coca-Cola cut its prices... attempting to stop tubaínas growth” (B-brands are colloquially referred to as tubaínas) while the Financial Times (1999) states “(t)he large companies are fighting back... ha(ving) cut the price of a Coke by 20 per cent.”

⁹Coca-Cola kindly shared the data with us for the purpose of this study.

¹⁰The points scale 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 2012, Economist 2007). 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, using nominal 2000 R\$/US\$).

These demographic patterns are consistent with media and market research reports. The 1996-1999 upward mobility was fueled by successful economic reforms in the early 1990s, including trade liberalization and, most notably, the taming of very high inflation by the 1994 *Real* stabilization plan. These reforms were followed by strong consumption growth across the Brazilian economy, particularly among lower-income households. Figure 3 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.¹¹

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). As for the downward 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 currency crisis, and the 2000-01 Argentine crisis.

To analyze the impact of the changing socioeconomic composition, we define three socioeconomic groups: “Established Affluent” (EA), “Newly Affluent” (NA) and “Poor” (P). Using the IBOPE proportions together with urban household counts, we track the number of households who belong in each of these groups, in each region and over time. Our “Established Affluent” group consists of urban households who were already in **ABC** status in 1996, i.e., before the process of upward mobility took off. The number of households in this group, in each of the seven regions, is thus fixed over time at the initial number of **ABC** households in that region. We define the size of the “Poor” group in each region-period market gt by the number of urban households that IBOPE assigns in that market to **DE**.

Finally, we define the size of the “Newly Affluent” group in market gt as the difference between the contemporaneous number of **ABC** households and region g ’s initial (i.e., 1996) number of **ABC** households. The urban household population in each region gradually grows over time, and we interpret this process as (net) rural-to-urban migration, assuming that migrants join the “Poor” group among city dwellers.¹²

Applying these definitions to the data, the number of newly affluent households is strictly

¹¹See Carvalho et al. (2012), Salvo (2009, 2010) and Neri (1995) for further discussion of the impact of the reforms.

¹²To illustrate our computations using data, in the South region there were: (i) in $t = 1$, 3149 (thousand urban) **ABC** households and 2116 **DE** households, and (ii) in $t = 2$, 3238 **ABC** households and 2045 **DE** households. Between these periods, $3238 - 3149 = 89$ poor households moved up to newly affluent status, and the number of migrants was $3238 + 2045 - (3149 + 2116) = 18$. Thus the numbers of established affluents, newly affluents and poor in this region, respectively, are (3149,0,2116) in period 1 and (3149,89,2045) in period 2. The online appendix provides consistency checks between different data sources.

positive for all regions and all time periods $t > 1$, and is equal to zero, by definition, for $t = 1$, the initial period of our Nielsen soda market data. The zero number of newly affluent 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 begins in late 1996. This assumption is also consistent with press and trade articles from the time. For example, our measure of the number of newly affluent households in 1999, summed across the seven Nielsen regions, translates into 20m consumers, well in line with a Financial Times’ (June 1999) count of “25m new consumers” across Brazil, noting that our study does not cover rural areas or the northern states.

Figure 4 plots the evolution of the 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 affluent reflects the joint effects of the recession and the urbanization process. There are large regional disparities, with states in the Northeast being the least affluent and São Paulo Metro being the most affluent (respectively, 65% and 36% of urban households in these regions 1 and 4 are initially poor).

2.3 Data on household-level brand choices

Our third main data source allows us to relate household characteristics to soda consumption choices at the start of our period of study. We use an urban household expenditure survey—hereafter HEX—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). Since the survey was conducted shortly before the beginning of our Nielsen market data, we use it to construct an initial condition for the evolving relationship between socioeconomic standing and consumption choices, as detailed in Section 4.

This survey reports the soda brand purchased, and the amount spent, for consumption inside the home. Households are not classified according to the ABCDE system, but we use the detailed information available (e.g., ownership of a refrigerator, numbers of TVs and bathrooms in the residence) to assign, like IBOPE does, each household to an affluence group from A to E. Table 2 reports the relationship between inside-the-home soda consumption and socioeconomic status. For example, 34.5% of São Paulo Metro’s **ABC** households purchased soda whereas only 19.8% of **DE** households did so. Table 2 also shows, for each region, the shares of **ABC** and **DE** households who consume premium and generic brands. The ratio of generic-to-premium consumption tended to be higher for **DE** soda-consuming households compared to **ABC** soda-consuming households, e.g., such ratios in São Paulo Metro were 2.6 : 17.3 for **DE** and 1.4 : 33.1 for **ABC**. As we note in the online appendix, our modeling of households at each point in time as either premium or generic shoppers, but not “hybrids,” is largely consistent with the data.

Additional data. Our analysis also draws on (see online appendix for sources): (i) urban household population by region over time; (ii) advertising intensity at the brand-market level, (iii) temperature, and (iv) input prices: sugar, electricity and fuel.

3 Possible drivers of growth in the fringe

Who are the generic fringe producers, and what explains their staggering growth? In contrast to the established Coca-Cola/Ambev duopoly, with their heavily advertised brands and nationwide distribution, fringe players ran small-scale operations, in most cases individually covering a fraction of a state, or even just a dozen or two municipalities (Brazil has 27 states, each comprising 310 municipalities on average). Often referred to as “Tubaínas” (originally, a trademark of a small São Paulo soda seller as early as the 1940s), fringe firms have been present in Brazil’s soft-drink market for over half a century. Such firms’ brands sell at substantially lower prices compared to premium brands. While some generics may invest in branding or advertising—if so, at a very local level—for the most part these are single-product firms offering little product differentiation. A typical retailer carries two to three generic brands, offered at similar prices.

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. Much of this growth was fueled by entry, rather than by the expansion of existing firms. An individual fringe firm’s distribution typically remained limited to a small geographic area, and no individual fringe firm accounted for more than 1.5% of the Brazilian market (Gertner et al 2004). No census of fringe operators exists, but industry sources suggest that following three years of net entry, the number of firms selling generic soda may have surpassed 500 by 1999. For example, Ambev’s (2003) financial reports referred to “700 low-price brands” while the Financial Times (1999) cited over 900 firms.

While premium brands enjoyed market power via product differentiation, fringe prices were likely to stay close to marginal cost. The evolution of fringe prices shown in Figure 2 is consistent with marginal costs declining as a consequence of entry by increasingly efficient generic producers (and possibly the exit of less efficient ones). The all-but-nonexistent response of fringe prices to the mid-1999 premium price cut is also consistent with competitive behavior in the fringe.

The rise of the fringe: supply-side factors. A likely supply-side factor contributing to fringe growth was the early-1990s shift from the 1-liter proprietary glass bottle (returned to the bottler for reuse, requiring a certain level of sophistication and scale) to the non-returnable 2-liter PET bottle. This may have lowered barriers to entry for small, unsophisticated firms (Ambev 2003). Consistent with the hypothesis of disposable containers aiding fringe growth, family-size PET bottles accounted for 92% of the fringe’s volume, compared to 78% of the premium volume, in supermarkets in 2001/02.

Another supply-side contributor is distribution. Our Nielsen data allow us to observe brands' retail presence and, importantly, our demand model controls for it. Figure A1 (right panel) indicates that the proportion of stores that stocked at least one generic brand exceeded 80% already in the first sample period, increasing to almost 100% over the first half of the sample. This rise in the share of retail locations stocking generics is consistent with the pronounced entry of generic brands into local markets. The already high initial generic penetration at retailers suggests that distribution, important as it may be, cannot fully explain the generic growth.

Demand-side factors. As in many other consumer markets in developing countries (discussed in the introduction), it is very likely that the rising Brazilian middle class acted as a catalyst for the growth of the fringe of soda producers. Before considering mechanisms that may drive this relationship, we provide descriptive evidence that such a link exists.

Table 3 reports OLS regressions of the generic fringe's volume share of the soda category on the share of the region-time market's population that is newly affluent (as defined in Section 2), where an observation is a region-time market. The newly affluent household share is positively and significantly associated with the generic volume share. This continues to hold as we add various controls: region intercepts, time trends, seasonal effects, premium brands' media advertising intensity (the sum of Gross Rating Points), and generic and (mean) premium brands' in-store presence. As in our subsequent structural demand model, we flexibly interact time trends, advertising and in-store presence with region dummies. Controlling for all these effects (column V), we find that a 1% increase in the fraction of the market's population that is newly affluent is associated with a 0.18% increase in the volume share of the generic fringe.

To get a sense of the quantitative implication, we refer to Figure 2 which shows that, between 1996 and 1999, the share of newly affluents went from 0 to roughly 20-30 percent across the seven regions. This suggests that the growth of the new middle class contributed, approximately, between 3.6 and 5.4 percentage points to the generic volume share. Noting that the generic volume share increased by 20 percentage points over this period, we conclude that the emergence of the new middle class can explain about 20 to 25 percent of the generics' growth. Further, this magnitude can plausibly be viewed as a lower bound: note the large drop in the newly-affluent-share coefficient as we move to column V and control for in-store presence. If some of the expanded distribution was prompted by the growth of the "new demand" for soda, then part of this effect should also be attributed to the new middle class. The contribution of the new middle class to generic growth, therefore, appears to be substantial.

While the results in Table 3 suggest a causal effect from the emergence of a new middle class to the growth of the fringe, they do not tell us much about the mechanisms underlying this link. One possible mechanism is *price sensitivity*: if newly affluent customers were more price sensitive than their established affluent counterparts, they may have been more prone to favor cheap generic

brands. This view is consistent with a survey by the Boston Consulting Group (2002) on the shopping preferences of Brazil's "Class C" (new middle class) consumers, claiming that brand was the least important factor in their food purchase decisions. This view is also consistent with that espoused by the Financial Times (1999), stated above, on the "new customers...(being) less concerned about expensive TV ads," as well as an article by the local business newspaper Exame (2002) stating that 63% of the market leaders in 157 product categories in Brazil lost market share between 1998 and 2000 (with leaders gaining share in only 12% of these categories).

Examining the right panel of Figure 2, one may interpret the 1996-2003 market share evolution as follows: until the premium price cut in 1999, generics consistently gained market share at the expense of premium brands thanks to the growing mass of newly affluent consumers, whose price sensitivity drove them towards low-cost substitutes. Following the price cut, it is possible that the incoming "new consumers" were more evenly split among premium and generic brands, explaining why volume shares immediately stabilized. This can be one explanation for why the large, permanent price cut led to a change in market share trends.

A second potentially relevant mechanism is *persistence in demand*. In our context, such persistence would imply that the emergence of a new middle class introduces into the market a substantial mass of customers who have not yet adopted any persistent shopping patterns or "habits" in the soda category. This provides an opportunity for generic entrants to win market share. If these new customers then develop the habit of consuming generic soda, convincing them to pay premium prices down the road may be costly. This mechanism may provide premium brands with clear incentives to cut prices to induce at least some of the new consumers to substitute from the outside option to premium sodas. The raw data offers some clues regarding the potential presence of a persistent component in demand. The recessionary years, post 2000, saw the proportion of newly affluent households shrink considerably (Section 2). The recession was not accompanied by declining soda prices (Figure 2). One would expect soda consumption to fall, yet it remained stable. Figure 3 contrasts the stability of soda consumption in the recessionary period with the substantial decline in the sales of cement, after years of growth for soda and cement alike. The differential pattern is suggestive of persistent preferences for soda, relative to other product categories, noting that cement can be considered a consumer good in Brazil.

In the next section we develop a structural demand model that accommodates both socioeconomic mobility and class-specific price sensitivities while controlling for persistence in preferences. The estimated model is consistent with the emerging middle class being relatively price-sensitive, and also helps rationalize Coca-Cola's large price cut.

4 A structural demand model

This section develops a model of household demand for soft drinks. Subsection 4.1 develops the utility framework, while subsection 4.2 describes the dynamic process that characterizes households' mobility across socioeconomic and consumption states. Intuitive arguments for identification and a description of the estimation algorithm follow in subsections 4.3 and 4.4.

4.1 Household types and the utility framework

We denote the eight premium brands (or A brands) as elements of the set \mathcal{A} , and the ninth brand category as the only element of the set of generics (or B brands) \mathcal{B} . We allow households' current preferences over substitute soft-drink brands to depend on both their current socioeconomic state, and on their previous-period consumption. In each period t , a household belongs in one of three *socioeconomic states* (EA, NA, P), standing for established affluent, newly affluent and poor, respectively (we subsequently explain how the model accommodates region-specific aggregate mobility across these states over time). A household is also characterized by its previous-period consumption choice, which we refer to as the household's *habit state*. This terminology is adopted for concreteness only—as we further discuss below, habits are just one form of persistent preferences, and it may be difficult to tell them apart from other forms of serial correlation.

We allow for three such habit states, differentiating households who, in the preceding period, consumed: (i) a premium brand $j \in \mathcal{A}$, (ii) a generic brand $j \in \mathcal{B}$, or (iii) did not consume any soda, i.e., chose the outside option (\mathcal{O}). Crossing together the three socioeconomic groups and the three prior-consumption states, we obtain nine discrete household types, indexed by r :

$$r \in \mathcal{R} := \{EA^{\mathcal{A}}, EA^{\mathcal{B}}, EA^{\mathcal{O}}, NA^{\mathcal{A}}, NA^{\mathcal{B}}, NA^{\mathcal{O}}, P^{\mathcal{A}}, P^{\mathcal{B}}, P^{\mathcal{O}}\} \quad (1)$$

Thus, for example, a time- t newly-affluent household who consumed a generic brand in period $t - 1$ is of type $r = NA^{\mathcal{B}}$, whereas an established affluent household who consumed a premium brand in the preceding period is of type $r = EA^{\mathcal{A}}$. Fixing a region-period market gt , let $F_{r,gt}$ denote the type- r fraction of that market's household population. We collect these fractions for the nine types in a 9-dimensional vector denoted \mathcal{F}_{gt} , to which we refer as market gt 's *type-distribution vector*. The indirect utility of household i of type r in market gt from consuming brand j is given by:

$$u_{i \in r, j, gt} = \delta_{jgt} + \alpha_r \cdot p_{jgt} + \lambda \cdot h_{jr} + \epsilon_{ijgt} \quad (2)$$

The term δ_{jgt} denotes a market-specific, household-invariant base utility from brand j : $\delta_{jgt} = x'_{jgt}\beta + \alpha \cdot p_{jgt} + \xi_{jgt}$. Vector x_{jgt} contains brand-region fixed effects, seasonal effects, market temperature, brand-level advertising and in-store presence, and region-specific time trends. These trends allow for region-specific temporal evolution in the utility from the outside option, such as differential rates of expansion in markets for soda alternatives (e.g., juices). The brand’s price is p_{jgt} , and ξ_{jgt} denotes a (brand-market specific) utility shock observed by firms and consumers, but unobserved to the econometrician.

The second and third terms in (2) introduce household-type heterogeneity. The parameter α_r shifts the base price sensitivity α in accordance with the household type r :

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

Whereas 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. Note that we allow price sensitivity to vary with a household’s socioeconomic standing, but not with its “habit.” Below we provide intuition for the role played by this restriction in identification.

The variable h_{jr} in (2) captures the persistence, or habit, feature. In our baseline specification, which we refer to as *Brand Type Persistence* (BTP), this is given by:

$$h_{jr}^{BTP} := \begin{cases} 1 & \text{if } r \in \{EA^A, NA^A, P^A\} \text{ and } j \in \mathcal{A} \\ 1 & \text{if } r \in \{EA^B, NA^B, P^B\} \text{ and } j \in \mathcal{B} \\ 0 & \text{otherwise} \end{cases} \quad (3)$$

This specification implies that consuming *any premium brand* in the previous period shifts one’s utility from consuming *any premium brand* in the current period by a magnitude of λ . Such a household is characterized by a “premium” habit in the current period. Similarly, consuming *any generic brand* in the previous period shifts one’s utility from consuming *any generic brand* in the current period by λ , a situation we refer to as a “generic” habit.

Our modeling of state dependence is, thus, parsimonious: persistence applies to a class of brands—premium or generic—rather than operating at the individual brand level. This choice is driven by our motivation: to capture a potentially important mechanism in an emerging-market setting characterized by rapid growth in discount brands with minimal advertising. We provide more motivation for this modeling choice below, in the context of observed data variation. The online appendix reports a robustness check in which brand-loyalty effects were allowed, producing

similar qualitative findings.¹³ We consider additional alternative specifications below.

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 specification by defining the utility from the outside option, $u_{i \in r, j=0, gt} = \epsilon_{i, 0, gt}$. We denote the model parameters to be estimated by $\theta = \{\beta, \alpha, \alpha_{EA}, \alpha_{NA}, \lambda\}$.

Two aspects of the model are worth noting. First, letting only the preceding period’s choice—rather than a longer history—determine the state is consistent with the literature. Second, again consistent with the literature, we do not allow consumers to be forward-looking: they maximize current-period utility and do not internalize the effect of current choices on future utility. We view this approach as appropriate given the nondurable nature of the product.

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

$$s_{j,r,gt}(\theta) = \frac{\exp(\delta_{jgt} + \alpha_r \cdot p_{jgt} + \lambda \cdot h_{jr})}{1 + \sum_{\ell=1}^J \exp(\delta_{\ell gt} + \alpha_r \cdot p_{\ell gt} + \lambda \cdot h_{\ell r})}, \quad (4)$$

where $J = 9$ is the number of brands sold in each market. The notation $s_{j,r,gt}(\theta)$ reflects the fact that these shares are model predictions that depend on parameter values. The predicted aggregate share for brand j is the weighted sum of brand j ’s share across each of the nine household types, where the weights are the fractions of the population of each type:

$$s_{jgt}(\theta) = \sum_{r \in \mathcal{R}} F_{r,gt} \cdot s_{j,r,gt}(\theta) \quad (5)$$

We, therefore, introduce a random-coefficient logit model where unobserved, discrete heterogeneity is present via the nine consumer types, noting that we only observe aggregate-level market shares in each region-period market, rather than type-specific shares.

Variants on the persistence mechanism. The above Brand Type Persistence (BTP) model offers a specific way of modeling persistence in utility and demand. It emphasizes the notion that customers may have persistent shopping patterns: buying expensive, highly-advertised brands as distinct from cheap generic substitutes. The raw data lends some support for such a mechanism: for instance, while the large premium price cut in 1999 slowed down generics’ growth, it did not steal the volume share they had acquired by 1999 (Figure 2). A persistent pattern to consume generic soda may, therefore, have shielded generic producers from losing existing customers following the premium price cut (while the price cut could have still been very effective in inducing consumers to switch from the outside option to premium consumption).

¹³Also reported in the appendix is a specification that allows the effects of previous consumption to differ in magnitude across brand types rather than modelling a uniform effect of λ . While that analysis suggests that a history of premium consumption could have a larger effect than generic consumption, another specification showed the effects to be similar. In the interest of parsimony, we maintain a uniform persistence parameter in our baseline specification.

While we view this mechanism as likely, alternative persistence mechanisms may have been present. The fact that total soda consumption did not fall during the recession (see Figure 3) could be consistent with there being a persistent tendency to consume soda *of any kind*. To capture this possibility, we introduce a variant of our baseline model dubbed *Soda Category Persistence* (SCP). In this model variant, the persistence term is defined as:

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

Here, 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. Alternatively, we consider a *No Persistence* (NP) model variant that turns off state dependence in preferences by constraining λ to equal zero. In Section 5 we report results for all three model variants—namely, our baseline BTP model, SCP and NP—in order to investigate the role played by different assumptions regarding the persistence mechanism.

4.2 Dynamic type evolution

We now describe the dynamic process that shifts the distribution of households across types over time. Social mobility (and, to a lesser extent, rural-to-urban migration) in a particular urban region g changes the number of households in each socioeconomic group. In addition, in each period t , households make consumption choices, contingent on type and on model parameters, that affect the state with which they enter period $t + 1$. These processes determine how the type-distribution vector \mathcal{F}_{gt} evolves over time.

We begin by computing, for each region, the first period’s type-distribution vector, \mathcal{F}_{g1} . This follows directly from the HEX survey that was conducted right before our Nielsen data begins, linking a household’s socioeconomic class to its choice among premium soda, generic soda, or “no soda.” Fractions of the population at $t = 1$ belonging in each of the three established affluent (resp., poor) types are set in proportion to the HEX shares for **ABC** (resp., **DE**) households across premium brands, generics, and no soda. Since the number of newly affluent households was set to zero at $t = 1$ (Section 2), corresponding population fractions are zero by construction.

Given a particular value for the model’s parameters θ , these fractions are updated forward for periods $t = 2, \dots, 57$. Fixing region g and period t , we explain how to obtain $\mathcal{F}_{g(t+1)}$ given \mathcal{F}_{gt} and a value for θ . Repeating this updating process for $t = 1, \dots, 56$, starting from a known \mathcal{F}_{g1} , yields the full trajectory of the distribution of household types over time.

Each generic value of θ yields a prediction, via (4), of the shares (and masses) of type- r households who consume premium and generic soda in period t . The social mobility and rural-urban migration that we observe in the aggregate data requires that we make assumptions on how

these demographic shifts interact with previous consumption. For instance, whenever aggregate upward mobility is detected in a given region between periods t and $(t + 1)$, it follows that some of period $(t + 1)$'s newly affluent households were poor in period t . Determining these households' $(t + 1)$ habit requires, therefore, an assumption regarding their soda choices in period t , when they were still "Poor." Aggregate downward mobility and urbanization similarly require assumptions. We specify two such assumptions. Appendix A.1 reports several checks that reassure us that our findings are robust to modifications in these assumptions (and also provides some examples).

Assumption 1 (Socioeconomic Mobility). *Previous-period choice probabilities for premium and generic brands among households moving up (down) from Poor to Newly Affluent (Newly Affluent to Poor) status are equal to the previous-period choice probabilities among all Poor (Newly Affluent) households.*

Assumption 1 implies that social mobility between periods t and $(t + 1)$ is independent of consumption choices at time t . A household who "moved up" from being poor at t to newly affluent at $(t + 1)$ is as likely to have consumed each type of soda at time t as any member in the wider population of poor households at time t . We similarly incorporate an assumption regarding rural-urban mobility:

Assumption 2 (Migration). *Households moving to urban areas join the Poor group and have a no-soda habit. Households moving out of urban areas leave the Poor group, and have premium, generic and no-soda habits in proportion to the shares of those habits among the Poor that remain.*

4.3 Sources of data variation and parameter identification

While a formal proof of identification lies beyond this paper's scope, we now intuitively discuss the ways in which the available data variation recovers the model parameters. The data offer two unique sources of variation. First, we observe socioeconomic transitions of considerable magnitude and different directions. Specifically, we observe both upward and downward mobility that occur at differential rates across the seven regions. Such variation is not typically encountered in a mature market, but it is available in the rapidly growing emerging market we study.

A second unique source of variation is the substantial variation in prices. The large and abrupt premium price cut midway through the sample is one such source of price variation (noting that, in mature markets, Coca-Cola and Pepsi typically charge rather stable prices). Further, fringe prices gradually decline over most of the sample period, a process we view as stemming largely from declining marginal cost.

In Appendix A.3 we illustrate the identification arguments by way of Monte Carlo studies that show how the socioeconomic transitions and price variation, of the magnitude that we observe

in the aggregate data, help identify the parameters associated with the modeled persistence mechanism and heterogeneity in price sensitivity.

Let us first describe data variation that should help pin down the *price sensitivity parameters*. As noted, we observe both the growth of a new middle class from 1996/97 on and its partial reversion during a recession beginning in 2000/01. Importantly, these demographic shifts occurred at differential rates across regions. While brand-region fixed effects control for *fixed* differences in preferences across regions (e.g., for cultural or historical reasons), the intra-region temporal shifts in the socioeconomic distribution of households provide a key source of variation. To illustrate, consider two regions which dynamic evolution varies substantially, the Northeast and São Paulo Metro. The following table depicts socioeconomic and product market data for both regions at several points in time:

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

Shares: No Soda & Generic Soda	Dec/Jan 1997		Dec/Jan 2000		Dec/Jan 2003	
	s_0	s_{gen}	s_0	s_{gen}	s_0	s_{gen}
Northeast (Region 1)	87%	0.3%	82%	5.0%	80%	8.4%
São Paulo Metro (Region 4)	61%	6.6%	62%	12.6%	59%	11.9%

Sources: Nielsen, IBOPE, IBGE (PNAD). Market shares s_0 and s_{gen} are for the outside option and for generics, respectively.

At the start of our sample, the Northeast is substantially poorer than São Paulo Metro, with 65% and 36% of urban households being poor, respectively. At the same time, the Northeast exhibits lower soda penetration, with 87% of that region’s households not consuming soda compared to 61% for its wealthier counterpart. This cross-sectional variation may be explained not only by the poor being more price sensitive than the established affluent, but also by fixed differences in the regions’ preferences, the latter being controlled for by brand-region fixed effects in the utility specification.

From 1997 to 2000, the Northeast boasted stronger upward mobility relative to SP Metro: by 2000, 24% of northeastern households were newly affluent compared with 16% of SP Metro’s households. Over these same years, soda penetration ($1 - s_0$) grew substantially in the Northeast, from 13% to 18%, while it was about flat in SP Metro. Generic brands enjoyed a huge gain in share (s_{gen}) in the Northeast, from 0.3% to 5.0%, while in SP Metro that share “only” doubled.

Importantly, this joint temporal variation in social mobility and soda consumption choices helps pin down the price sensitivity parameters ($\alpha_{EA}, \alpha_{NA}, \alpha$). Since our model allows price

sensitivity to vary by socioeconomic standing, intra-regional household mobility between poor and newly affluent status changes the aggregate price sensitivity in the region. Co-variation of these shifters of price sensitivity with aggregate market shares, controlling for prices, helps identify the heterogeneous price sensitivities. As expected, Monte Carlo experiments in which we restricted social mobility increased the variance of our estimates.

Next, let us consider data variation that provides information on the *persistence parameter* λ . The table above reports the differential effects of the recession in the two regions. The Northeast saw the proportion of newly affluents shrink considerably, from 24% in 2000 to 15% by 2003, yet soda penetration continued rising, albeit at a lower rate, from 18% to 20%—recalling that the recession was not accompanied by declining soda prices. This pattern is suggestive of persistence in preferences: since we identify the poor as being more price sensitive than the newly affluent, a model with completely static preferences would predict that soda consumption in the Northeast should fall during the recession, implying that the data reject such a model.

Another data variation that is indicative of persistent preferences is the premium price cut in mid 1999. The table shows that in both regions the generic share held up quite well following the price cut. This provides further evidence of persistence and, specifically, of the baseline BTP mechanism that induces households to stay within brand type.¹⁴

Some patterns in the data are, therefore, consistent with persistent preferences. We should emphasize two assumptions that are key to identifying the persistence parameter λ . First, we assume that the demand errors ξ are not serially correlated (further discussing this below). Second, we place a specific structure on households' systematic individual heterogeneity, allowing it to depend only on the socioeconomic status and on previous-period consumption. Violation of these assumptions would likely suggest that part of the effect that we capture via λ stems from other sources, such as serial correlation in the demand errors, or systematic individual heterogeneity that we do not model.

Absent more detailed micro-level data, it is difficult to relax these assumptions and to separately identify λ from these other potential sources of serial correlation. That being said, we believe that the unique variation available to us reveals that persistence in preferences *is present in the data*. Further, we show below that the “No Persistence” model variant does not perform well in our application. This failure is quite intuitive: such a model interprets the data variation during the recession as evidence that the poor are “not that price sensitive,” since all one would see through the lens of the model is an increasingly poor population consuming stable amounts of soda. It might then be more difficult to elicit greater price sensitivity among the less affluent.

¹⁴As pointed out by a referee, the premium brands' falling share in the early part of the sample could be viewed as evidence against there being a “habit” of premium soda consumption. We note that this loss pertains to share of actual soda volume (summing to one, see Figure 2), not to loss of share of potential market size (Figure 5 below) or to volume in levels (Figure 1), which did not fall. The large body of evidence surveyed in Section 3 pointed to a tendency of the new customers to flow toward generic brands, rather than to a massive switch of existing customers from premium to generic soda consumption.

To sum, we acknowledge the limitations of our approach in terms of pinning down the exact persistence mechanism. Still, we incorporate persistence into our estimated model for two reasons: first, we believe that this mechanism is present in the data and may be important in its own right for understanding the dynamics of competition between premium and generic brands. Second, we find that failing to control for this mechanism hampers our ability to identify the price sensitivity parameters.

4.4 The estimation procedure

Since the type-specific shares $s_{j,r,gt}(\theta)$ from (4) are not observed in our data, estimation is based on matching the aggregate shares $s_{jgt}(\theta)$ predicted from (5) with shares s_{jgt} computed from the Nielsen data. This follows the spirit of the literature that estimates random-coefficient demand models with discrete-type unobserved heterogeneity (e.g., Berry, Carnall and Spiller 1996, Kalouptsi 2010). In this literature, types are often abstract groupings of “similar” consumers, and the population fractions of these types are treated as parameters to be estimated. For example, Nair (2007) models video-game consumers as either “high valuation” or “low valuation” types, and estimates their relative population fractions.

In contrast, our approach places a specific structure on the unobserved heterogeneity, tying it down to both the household’s current socioeconomic standing and to its previous-period consumption choice. Population fractions of the unobserved types are, therefore, computed by combining data on aggregate social mobility and model predictions of household choices. Because time- t choices determine time- $(t + 1)$ states, we incorporate a dynamic updating routine into each evaluation of the GMM objective function. We briefly describe the logic of the estimation algorithm and leave further details to Appendix A.2.

As in any discrete-choice framework, to compute the empirical analog of predicted shares s_{jgt} (5) we need to define the relevant market size, that is, the maximum amount of soft drinks that can potentially be consumed in each market gt . We define this maximum quantity, denoted \mathcal{M}_{gt} , as six liters per week over the duration of period t multiplied by the number of urban households residing in market gt . One may 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). We then compute brand j ’s share as $s_{jgt} = q_{jgt}/\mathcal{M}_{gt}$, where q_{jgt} denotes brand sales volume observed in the Nielsen data. The share of the outside option is then $s_{0gt} = 1 - \sum_j s_{jgt}$. The online appendix reports robustness checks with regard to the assumption on market size.

The following steps allow us to construct a GMM objective function and evaluate it at some generic value of the model’s parameters θ . Conditional on θ and the first-period type-distribution vector \mathcal{F}_{g1} (per subsection 4.2), (5) yields predicted aggregate brand shares in period $t = 1$. Using the contraction mapping from Berry, Levinsohn and Pakes (1995), we invert the market share

equation that equates predicted aggregate shares to observed shares, solving for the unique vector of brand-specific base utilities δ ; we “match” s_{jgt} rather than *type-specific* $s_{j,r,gt}$ as the latter are unobserved. We then use these δ in (4) to predict type-specific choices, from which we obtain the type-specific fractions of households who in $t = 1$ consume premium brands, generic brands, or no soda. These predicted choices (which generate second-period habit states) are then combined with the aggregate socioeconomic mobility observed between $t = 1$ and $t = 2$ to obtain the second-period type-distribution vector \mathcal{F}_{g2} , as explained in subsection 4.2.

Repeating this process for periods $t = 2, \dots, 56$ (and solving for δ in the last sample period $t = 57$), and then repeating for each region g , we obtain the base utilities for every brand in every region-period market. From $\xi = \delta - x'\beta - \alpha p$, we compute the demand unobservable $\xi_{jgt}(\theta)$ for each brand j in each market gt . We follow the demand estimation literature and make the identifying assumption that these unobservables are mean-independent of a set of instrumental variables Z . This assumption gives rise to a GMM objective function that attempts to set the co-variance between the instruments and the computed demand unobservables as close to zero as possible.

We noted above our assumption that ξ are not serially correlated. Our inclusion of brand-region fixed effects and region-specific trends alleviates this concern to some extent. We perform an ex-post check by calculating the simple correlation between current-period and previous-period estimated values $\hat{\xi}$ for each brand-region combination. This calculation yields correlations that are mostly negative (52 out of 63) and mostly small in absolute value. Of the six correlations that exceed 0.2 (with a maximum correlation of 0.5), five are in the Northeast (region 1), and they do not pertain to the leading brands Coke, Guaraná Antarctica and Fanta. These patterns are consistent with lack of positive serial correlation in ξ .

Choice of instruments. Following the literature, we seek instruments that would address the likely correlation of the demand errors ξ with prices. Importantly, in our application, pricing may be non-optimal. Coca-Cola set a certain pricing level up to mid 1999, and then changed it abruptly. The fact that Coca-Cola replaced the management team at its Brazilian subsidiary around the time of the price cut (Salvo 2009) suggests that the firm may not have applied a consistent pricing rule throughout the sample period. Rather, the firm may have been searching for an adequate response to the generics’ growth. Further, that Coca-Cola “shocked Wall Street... by announcing a rare drop in quarterly sales” (Financial Times 1999) suggests that agents might not have enjoyed perfect foresight with respect to growth in competition.

In the absence of a well-specified, formal rule for optimal pricing, how can one justify the choice of instruments for price? We believe that three sets of demand instruments can be justified.¹⁵ The first class of instruments are input prices, namely, sugar, electricity and fuel prices. These cost

¹⁵Salvo (2009) uses similar instruments to estimate an AIDS demand model, a different approach compared to the discrete-choice model we offer in this paper. Instrumenting for price endogeneity is similarly relevant to both frameworks.

shifters should affect prices, regardless of the exact pricing rule. Specifically, we interact these three input prices with two brand type dummies (premium or generic), forming six instruments.

The second class of instruments borrows from Hausman, Leonard and Zona (1994). We instrument for a brand’s price in a given region with the contemporaneous mean price for this brand in the other six regions, forming one instrument. The identifying assumption is that prices are correlated across regions through a common cost structure or through common shifts in the way firms strategically interact (for instance, the mid-1999 premium price cut). This approach can be challenged if common demand unobservables are present (see Bresnahan 1997a, 1997b). However, such issues are of a lesser concern in our setting, for two reasons. First, we control for region-specific, brand-level advertising intensity and in-store presence, often absent from demand studies. Second, the Brazilian market features considerable regional variation in demand, distribution and promotion practices. We also note that this second set of instruments provides valuable variation across brands and regions. Such variation is not available from the cost shifters, which only vary over time.

A third set of instruments is afforded by the premium brands’ mid-1999 price cut. Consistent with the change of management and related issues discussed above, we argue that this sharp price cut was exogenous to the brand-region-time specific demand unobservables ξ_{jgt} . It was a response to the demographic shifts and fringe expansion that we observe in 1996-1999, rather than a response to some sudden unobserved mid-1999 demand shock (noting that we also control for advertising intensity, in-store presence, weather shocks and region-specific drifts). In practice, we generate a post-July 1999 dummy variable and interact it with brand-region fixed effects, forming $9 \times 7 = 63$ instruments that allow the effects to vary by brand within each region.

5 Results from the structural demand model

5.1 Demand estimates and their economic significance

Table 4 reports estimates from the baseline demand model, Brand Type Persistence. The price sensitivity of poor households has the expected negative sign and is precisely estimated at -5.7. The parameters α_{EA} and α_{NA} are positively signed, with $\hat{\alpha}_{EA} > \hat{\alpha}_{NA}$. This implies that the established affluent are the least price sensitive group, the poor are the most sensitive group, and the newly affluents display an intermediate level of sensitivity. Namely, $\hat{\alpha} + \hat{\alpha}_{EA} = -2.1$ (with a standard error of 0.2) for the established affluent, and $\hat{\alpha} + \hat{\alpha}_{NA} = -3.9$ (with a standard error of 0.6) for the newly affluent. This ranking is consistent with the media and analysts’ reports that were referenced in Section 3. The newly affluent are marginally less price sensitive than the poor (the p-value of a test of equality is 0.10) but significantly more price sensitive than the

established affluent (the p-value of a test of equality is 0.00).¹⁶

The persistence effect λ comes out strongly significant, both statistically and economically. As we note in the online appendix, this finding is robust across different specifications. In particular, we estimated alternative specifications in which the base utilities from premium and generic brands are allowed to depend on social class, by adding class-specific intercepts to the utility function. We generally found that the large and precisely estimated λ survived, suggesting that our baseline λ is not picking up omitted differences in utility across types. The added intercepts in such specifications seemed to be confounded with price sensitivity parameters (and carried counterintuitive signs).¹⁷

Our utility framework provides a measure of the increased willingness to pay for generic (premium) soda that results from previous consumption of generic (premium) soda. In particular, the implied increase for a newly affluent household is $-\lambda/(\alpha + \alpha_{NA}) = 1.1$ R\$/liter.¹⁸ Further, the implied increase in WTP for a generic over a premium brand when the household previously consumed a generic *rather than a premium brand* is twice this amount.

Table 4 further reports the effects of several shifters of δ_{jgt} , the base utility of consuming brand j in market gt . Our specification includes brands' media advertising and in-store presence, both of which are allowed to have region-specific effects. In the case of advertising, for instance, this might reflect cross-sectional variation in households' exposure to media, and the fact that these measures pertain only to the main cities within each region, and are thus not comparable across regions. We generally obtain positive effects for media advertising and for in-store presence. To gain a sense of variation in advertising over time and across premium brands, note that intensity for the Coke brand in São Paulo Metro was measured at 2199 Gross Rating Points in December 2000, rising to 3587 GRP in December 2001, whereas Pepsi's GRP were 351 and 598 in these respective periods (fringe advertising is measured at zero GRP throughout). The online appendix further discusses these covariates and reports robustness tests in which they are dropped.

Coefficients on (similarly region-specific) time trends tend to be negative, consistent with continued improvement in the value of the outside option, which includes beverages other than soda such as juices and (tap or bottled) water, concomitant with the overall trend of economic growth.¹⁹ Finally, note that our specification controls for brand-region fixed effects, and for bi-monthly seasonal effects interacted with brand type, to allow these effects to differ across premium and generic brands. Over and above seasonality, market gt 's mean temperature is

¹⁶The online appendix reports robustness checks. Across the bulk of specifications we tried, the exact point estimate for α_{NA} is less robust when compared to other parameter estimates, but the finding that it does not exceed α_{EA} (i.e., newly affluent are no less price sensitive than established affluent) holds rather consistently.

¹⁷Details are available from the authors upon request. We thank one reviewer for suggesting these specifications.

¹⁸Corresponding amounts for established affluent and poor households are 2.0 and 0.7 R\$/liter, respectively.

¹⁹Forbes (2004) reports that the "the juice category grew twenty times over the past decade, albeit from a low base." IBGE's annual household surveys (PNAD) also indicate a sustained increase in access to tap water and piped sewerage in urban Brazil. Note that we rescale variables to vary between 0 and 1, so the trend coefficients reflect variation from the start to the end of the sample: these effects are economically small.

estimated to have a significantly positive effect on demand.

Table 5, Panel A provides predictions for the share of households purchasing soda, by type of household and brand. Overall soda penetration is 69%, 22% and 5% for the established affluent, the newly affluent and the poor, respectively (these are means computed across all regions over the last 12 months of the sample). Compared to the established affluent, newly affluent households are *more* likely to favor generic over premium brands: generic-to-premium consumption ratios are 0.5, 1.4 and 2.7 for the established affluent, newly affluent and poor groups, respectively.

The transition matrices in Table 5, Panel B illustrate the combined effects of the estimated heterogeneous price sensitivities and state dependence. The first three columns report mean transition rates over the first 12 months of the sample (and across all regions), prior to the premium price cut. For this subsample, a newly affluent household who in the previous period consumed a premium (generic) brand will consume the same brand type in the current period with a probability of 0.47 (0.56)—and these “stay with brand type” probabilities are higher for the established affluent, namely 0.85 (0.79) for premium (generic). Notice that such probabilities are smaller early in the sample compared with those for the last 12 months of the sample, after the price declines, shown in last three columns. For example, the stay-with-brand-type probabilities for a newly affluent household who in the previous period consumed premium (generic) is now 0.69 (0.75), compared to 0.47 (0.56) early on. The estimates suggest that the premium price cut was effective in raising these probabilities, thus strengthening market segmentation.

These transitions also reveal that while persistence (as we model it) is “large”—see the probability mass along each “diagonal”—consumers *do* transition across brand types, particularly by not consuming soda for at least a period. For example, in the early part of the sample, a newly affluent who in the preceding period consumed premium soda still has a substantial 0.52 probability of not consuming soda this period, in which case the probabilities that in the next period the then NA^O household chooses premium or generic are 0.01 and 0.02 respectively.²⁰

Table 6 reports estimated own-price elasticities (means over all markets gt), listing both aggregate elasticities for the leading brands, and elasticities by household type for Coke and for generics. A 1% increase in Coke’s price lowers its market share by 1.7%, compared with somewhat larger (in magnitude) elasticities of -2.0 to -2.1 for the other premium brands, Guaraná Antarctica, Fanta and Pepsi. The own-price elasticity for generics is -0.7 . While this value may seem low, note that this is the elasticity of demand for the aggregation of generic brands. The demand for each individual generic brand should be much more elastic, given the limited differentiation and fierce price competition within the fringe.

²⁰Our habit formation interpretation of the estimated state dependence is thus less extreme than the view espoused in the Economist (2009): “Across the developing world millions—perhaps billions—of people are currently forming tastes that will endure for the rest of their lives. Put one of Kraft’s Oreos or Cadbury’s Flakes in their hands and they may become loyal customers for decades to come.”

Examining the nine type-specific elasticities, and fixing the type of brand chosen in the previous period, we see that demand is more elastic the lower is the socioeconomic standing. For instance, considering households who previously chose a premium brand, Coke elasticities are (-1.4,-3.3,-5.7) for established affluents, newly affluents, and poor, respectively. Fixing the socioeconomic standing, demand elasticities are substantially associated with previous consumption. Considering, for example, the newly affluent group, demand for Coke is less elastic for households who previously chose premium soda (-3.3) compared to those who did not (-4.2).

The evolution of these type-specific own-price elasticities is shown in Figure 5 (for region 4). The left panel considers the Coke brand. Demand by all types becomes less elastic halfway through the sample, when premium brands cut prices. The least elastic demand for Coke is by the established affluent who previously consumed a premium brand; the most elastic demand is by the three poor types, particularly those who did not consume premium soda previously. Coke elasticities among newly affluent types lie between the elasticities displayed by the established affluent and the poor. The right panel depicts own-price elasticities for generics. The demand for generics is less elastic among newly affluent types who previously consumed a generic brand than among established affluent types who did not previously consume a generic brand. Such patterns illustrate the role played by heterogeneous price sensitivities and persistent preferences.

Demand Model variants. Table 7 presents estimates under our two alternative modeling choices for the persistence mechanism. Under the No Persistence (NP) specification (panel B), the price sensitivity of the poor is statistically indistinguishable from that of the established affluent, and *weaker* than that of the newly affluent. Our intuition for the failure of this model to deliver a reasonable ordering of price sensitivities was discussed above: by shutting off state dependence, this restricted model appears to interpret stable consumption during the recession as evidence that the poor are not very price sensitive.²¹

Similarly to our baseline BTP model, the Soda Category Persistence (SCP) model (panel A) also delivers a positive and significant estimate for the persistence parameter λ . This model, however, cannot distinguish the price sensitivity of the poor from that of the affluent groups, given that both α_{EA} and α_{NA} are statistically insignificant. This does not serve as clear evidence that the SCP model is misspecified, or that it is inferior compared to our baseline model. As discussed in Section 4, we incorporate persistence in the model since the raw data provides some evidence for it, and also since controlling for it helps us avoid the unreasonable price sensitivity ordering delivered by the NP model. Our ability to choose the best modeling choice for persistence among alternative specifications such as BTP and SCP is more limited. Ultimately, since our baseline BTP model appears to do a better job in capturing price sensitivity differences across

²¹Guided by a referee’s suggestion, we verified that incorporating class-specific intercepts into the utility from premium and generic brands does not improve the performance of this model, and does not “substitute” for controlling for persistence. Details are available from the authors upon request.

groups, we choose to carry this model forward to the counterfactual analyses reported next.

5.2 Counterfactual analysis of the premium price cut

A striking feature of the data is the premium brands' sharp price cut of about 20%, led by Coca-Cola and promptly adopted by Ambev as well, almost halfway through the sample period. Per-liter premium brand prices stayed broadly flat at about R\$ 1.15 until mid 1999, then dropped abruptly to R\$ 0.90 and stayed at this lower level. Fringe prices, in contrast, experienced a gradual decline from R\$ 0.80 to R\$ 0.55 between late 1996 and mid 2000.

We employ the estimated model to simulate the evolution of market shares under alternative pricing paths for the premium sellers. We avoid assuming that Coca-Cola's pricing was optimal (recalling, in particular, the change in management shortly before the price cut). Instead, we investigate what might have prompted Coca-Cola to engage in the observed price cut by exploring the outcomes of alternative actions. Our findings are consistent with there being some clear advantages for the observed price cut over the alternatives we consider.

Specifically, we consider counterfactual scenarios in which (i) premium brands avoid the mid-1999 price cut altogether, (ii) premium brands cut prices in mid 1999 by alternative magnitudes, and (iii) premium brands cut prices earlier than mid 1999. In each of these counterfactual analyses, we prescribe the same policy (e.g. a 30% price cut) to all premium brands, i.e., both Coca-Cola's and Ambev's brands, motivated by the observation that the duopoly's observed prices followed similar patterns (i.e., high prices early and a cut of a similar magnitude in mid 1999). In addition, in each of these counterfactual analyses we hold generic prices at their observed paths. This is motivated by the view espoused earlier that fringe prices largely tracked marginal cost. The fact that fringe prices did not deviate from their downward trend in response to the observed premium price cut supports this approach.

Counterfactual 1: eliminating the price cut altogether. Figure 6 depicts a counterfactual in which the price cut is avoided (we plot the predicted effect of this policy for one region, the variation being similar in other regions). The left panel displays pricing, with the solid lines corresponding to observed premium and fringe prices, and the dashed lines corresponding to counterfactual premium prices. The right panel of Figure 6 reports the simulated impact on aggregate premium and generic market shares. Observed shares are marked by solid lines, whereas counterfactual shares are marked by dashed lines (noting that shares in this figure are out of the total market size, which includes the outside option, so that the premium and generic shares do not sum to one).

A clear picture emerges from this analysis: had premium producers failed to cut prices, they would have suffered a deep loss of market share, hitting a rock bottom in the winter of 2000. At that point, the premium share would have been 12%, compared to a share of over 20% in

the observed sample. In this counterfactual scenario, generic brands grow relentlessly at the expense of their premium competitors. The figure reports that the generic market share would have surpassed the premium share early in 2000. By 2003, generics would have enjoyed a sizable market share advantage over premium brands, 25% against 16%.

The deep price cut, therefore, seems to have prevented a substantial drop in premium market share. A key insight from this analysis is that the price change was especially effective in attracting customers who otherwise would have chosen the “no soda” option, and less effective in converting consumers of generic brands into premium consumption. This can be inferred by comparing the actual versus counterfactual shift in the premium shares over time to the smaller corresponding shift in generic shares. For example, actual inside shares in 2001/02 averaged 46% (28% premium plus 18% generic) to be compared with counterfactual inside shares of 40% (17% premium plus 23% generic). This is suggestive of substantial market segmentation, consistent with state dependence that reduces “business stealing” across the types of brand offered. Still, in the absence of a premium price cut, the fringe’s share would have been almost 30% larger.

While our analysis indicates that the premium price cut succeeded in avoiding a deep loss in share, this was achieved at a cost: premium sellers sacrificed a non-negligible portion of their margins to protect their market shares. To assess the impact of the price cut on earnings, we perform a back-of-the-envelope calculation of *variable profit* in both the observed sample, and under the counterfactual no-price-cut scenario. Using information gleaned from Ambev’s local SEC filings, and from conversations with industry insiders among other sources, we estimate that the premium brands’ combined variable profits (excluding fixed costs) during the first three years after the price cut amounted to R\$ 861 million, to be compared to counterfactual profits of R\$ 807 million, had the price cut not occurred.²² That is, a 6% loss in variable profit, considering only the medium run, was avoided by the premium price cut.

Though this analysis offers some justification for the price cut, it provides little insight regarding the effectiveness of the action’s specific magnitude and timing. We consider this next.

Counterfactual 2: changing the magnitude of the price cut. Holding the mid-1999 timing of the price cut, but varying its magnitude, we first consider a 30% price cut, i.e., a more drastic action than the observed 20% price cut. We find that this would have proved excessive: estimated variable profits would have amounted to R\$ 764 million compared to R\$ 861 million under observed prices, over the same three years.

Second, we consider a smaller 10% price cut in mid 1999. Three-year variable profits would then have been R\$ 871 million, slightly higher than the R\$ 861 million estimated under the observed price path. One may interpret this as evidence that the observed 20% premium price cut was excessive, in that a smaller price cut would have yielded higher profits. Examining the

²²See the online appendix for details on how this calculation was performed.

impact on *market shares*, however, suggests a more nuanced view. Figure 7 displays the 10% price cut and the corresponding market share evolution in the left and right panels, respectively. Compared to observed share paths (solid lines), counterfactual paths (dashed lines) imply that the 10% price cut would have allowed the generic fringe to achieve market share parity with premium brands by 2003. To the extent that the premium brands sought to protect both profits *and* market shares, the observed 20% price cut may be viewed as superior to the 10% cut: while slightly lowering profits, it accomplishes a potentially important goal of maintaining high market shares and, specifically, dominance over the fringe.

Protecting market shares, even at a cost, may indeed be rational insofar as current market share is an “asset” in the sense that it is predictive of future profit. The persistence of brand market shares is studied, for instance, by Bronnenberg, Dhar and Dubé (2009).

Counterfactual 3: an earlier price cut. We use the model to predict share paths had premium brands cut prices by the same observed magnitude but at the start of our sample, in early 1997, rather than in mid 1999. This earlier price cut would have allowed premium sellers to slightly increase variable profit from R\$ 861 million to R\$ 873 million. The benefits to an early price cut appear, therefore, to be limited. Given the persistence in preferences embedded in the model, one might have expected an earlier price cut to be very effective by locking a larger number of consumers into the “premium habit.” However, recalling the transitions of Table 5, premium-habit customers still exhibit substantial switching rates despite the large estimate for λ , mostly by switching “out” to the outside option. The fact that an earlier price cut would not have created substantial gains does not mean that we should interpret the observed mid-1999 timing of the price cut as “optimal.” It is plausible that the timing of the price cut was associated with Coca-Cola’s management change that led the firm to reshape its pricing strategy.

Discussion. The counterfactual analysis lends justification to the premium sellers’ strategic price cut. Our estimates suggest that the growing market segment of newly affluent consumers were relatively sensitive to price. The price cut allowed a larger share of these upwardly mobile “new consumers” to flow into premium soda consumption, whereas absent the price cut they would have largely migrated toward generic soda consumption.

We further demonstrate that the magnitude of the price cut can be justified if a premium brands’ objective function is defined over two arguments: medium-run profits, and market share dominance. A deeper price cut would have eroded margins and profits, while a more modest one would have been marginally more profitable than the observed 20% cut, but would have allowed the generic fringe to continue its relentless gain in share at the expense of the premium brands. We also do not find evidence that changing the timing of the price cut would have been beneficial to premium brands. In this sense, our estimated model rationalizes the price cut observed in the data as a strategy that effectively curbed the expansion of the generic fringe.

6 Concluding remarks

This paper examines two salient features of the Brazilian soda market: the emergence of a new middle class, and the rapid growth of a fringe of discount brands at the expense of premium brands. We show that the two phenomena are statistically linked, suggesting that a rising middle class provided fertile ground for the generic fringe.

Using mainly aggregate data that exhibit very rich variation, we estimate a structural demand model that highlights the price sensitivity of the emerging middle classes. The model also controls for persistence in demand, a feature that finds support in the raw data patterns and which may, in itself, help explain the success of generic entrants. The estimated model allows us to rationalize an observed, drastic price cut by premium brands as a means of protecting both medium-run profits, and long-term market share dominance over the fringe.

While our application focuses on the Brazilian soft-drink market, we view the issues addressed here as pertinent to many consumer goods markets in the developing world, where a tension between advertised branded offerings and discounted generics exists or is developing, with premium brands being prompted to cut prices in the wake of an emerging middle class. Understanding the features of demand and the nature of competition in such markets should be of 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] Atkin, D. (2013). Trade, tastes and nutrition in India. *American Economic Review* 103, 1629-63
- [4] Berry, S., 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] Bills, M. and P. J. Klenow (2004). Some evidence on the importance of sticky prices. *Journal of Political Economy* 112, 947-985

- [7] Boston Consulting Group (2002). Mercados pouco explorados: Descobrimo a classe C. BCG Publication, São Paulo. (Under-explored markets: Discovering the “C” class, in Portuguese.)
- [8] 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
- [9] Bresnahan, T. F. (1997b). The Apple Cinamon Cheerios war: Valuing new goods, identifying market power, and economic measurement. Mimeo, Stanford University
- [10] 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
- [11] Bronnenberg. B., J.-P. Dubé and M. Gentzkow (2012). The evolution of brand preferences: Evidence from consumer migration. *American Economic Review* 102, 2472-2508
- [12] Calzolari, G., A. Ichino, F. Manaresi, and V. Nellas (2012). When the baby cries at night: Uninformed and hurried buyers in non-competitive markets. CEPR working paper DP8856
- [13] Carvalho Filho, I. and M. Chamon (2012). The myth of post-reform income stagnation: Evidence from Brazil and Mexico. *Journal of Development Economics* 97, 368-386
- [14] 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
- [15] 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
- [16] 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
- [17] Economist (2011a). Politics in emerging markets: The new middle classes rise up. Print edition of September 3, 2011
- [18] Economist (2011b). Indonesia’s middle class: Missing BRIC in the wall. Print edition of July 21, 2011
- [19] Economist (2012). China’s motor industry: Stepping on the gas. Edition of April 24, 2012
- [20] 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
- [21] Erdem, T. (1996). A dynamic analysis of market structure based on panel data. *Marketing Science* 15, 359-78
- [22] Exame (2002). Por que as marcas líderes estão perdendo participação de mercado para produtos populares e baratos. Print edition of September 1, 2002. (Why leading brands are

- losing market share to cheap and popular products, in Portuguese.)
- [23] Ferreira, F. H. G., J. Messina, J. Rigolini, L.-F. Lopez-Calva, M. A. Lugo and R. Vakis (2012). *Economic Mobility and the Rise of the Latin American Middle Class*. World Bank, Washington D.C.
- [24] 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
- [25] Financial Times (2014). Slowdown puts 1bn middle class at risk. Edition of April 13, 2014
- [26] Forbes (2004). Segmento de sucos cresceu 20 vezes nos últimos dez anos, mas brasileiro ainda consome pouco. (Juice category grew twenty times over the past decade, yet Brazilians still consume little, in Portuguese.) Edition of February, 2004
- [27] Frankel, D. M. and E. D. Gould (2001). The retail price of inequality. *Journal of Urban Economics* 49, 219-239
- [28] 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
- [29] Gemma Ware (2013). The rise of Africa's B-brands. Online edition of May 16, 2013
- [30] Gertner, D., R. Gertner and D. Guthery (2004). Coca-Cola's marketing challenges in Brazil: The Tubainas war. Thunderbird case study A12-04-0025
- [31] 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
- [32] 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
- [33] Kalouptsi, M. (2010). From market shares to consumer types: Duality in differentiated product demand estimation. *Journal of Applied Econometrics* 27, 333-342
- [34] Keane, M. P. (1997). Modeling heterogeneity and state dependence in consumer choice behavior. *Journal of Business & Economic Statistics* 15, 310-27
- [35] Lach, S. (2007). Immigration and prices. *Journal of Political Economy* 115, 548-587
- [36] 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
- [37] 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
- [38] Nevo, A. (2000). A practitioner's guide to estimation of random coefficients logit models of demand. *Journal of Economics & Management Strategy* 9, 513-548

- [39] Nevo, A. and Konstantinos Hatzitaskos (2006). Why does the average price paid fall during high demand periods? Mimeo, Northwestern University and University of California, Berkeley
- [40] 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
- [41] 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
- [42] Sancheti, S. and K. Sudhir (2009). Education consumption in an emerging market. Mimeo, Yale University
- [43] 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
- [44] SuperHiper (2003). Survey presents consumer purchasing habits at retail. Trade publication, edition of 03/2003
- [45] Wall Street Journal (2011). A rags-to-riches career highlights Latin resurgence. Front page article in the print edition of November 15, 2011

A Appendix

A.1 Dynamic type evolution

We provide examples, from the data, of the dynamic updating process. We consider two transitions, both for São Paulo Metro (region 4). The first transition, from $t = 1$ to $t = 2$, features upward mobility and, unusually in the data (yet we need to allow for this), a slight flow of residents out of the region (“net urban-to-rural migration”). The second transition, from $t = 10$ to $t = 11$, features upward mobility and the rural-to-urban migration that is prevalent in the data. We illustrate these transitions at the estimated model parameters θ^* . We also comment on the robustness of our estimates to the baseline mobility Assumptions 1 and 2.

Region 4, $t = 1$ to $t = 2$. The initial type-distribution vector is

$$\begin{aligned} \mathcal{F}_{g=4,t=1} &= \{F_{EA^A,4,1}, F_{EA^B,4,1}, F_{EA^O,4,1}, F_{NA^A,4,1}, F_{NA^B,4,1}, F_{NA^O,4,1}, F_{PA,4,1}, F_{PB,4,1}, F_{PO,4,1}\} \\ &= \{0.255, 0.029, 0.361, 0, 0, 0, 0.048, 0.009, 0.299\} \end{aligned}$$

As explained, the last element, for instance, is the product of (region 4’s) poor household count in $t = 1$ (observed from combining IBOPE and IBGE) and the share of the outside option among region 4’s **DE** households (calculated from the HEX 95/96), divided by the total household count (IBOPE/IBGE), i.e., $1346585 \times 0.84093/3789771 \simeq 0.299$. From $s_{j,r,g=4,t=1}(\theta^*)$ (see (4)), we obtain the mass of households for each of the nine types who choose to consume premium, generic, or no soda. For example, the premium soda share among established affluent households who in the previous period consumed premium soda is $\sum_{j \in \mathcal{A}} s_{j,EA^A,g=4,t=1}(\theta^*) \simeq 94\%$.

In contrast, the premium soda shares among established affluents who previously consumed generic soda and no soda are, respectively, 3% and 20%. Thus, since the established affluent population is constant over time (at 2443186), the number of established affluent households going into $t = 2$ in the premium-soda prior-consumption state is (in thousands, hereafter) $3790 (0.255 \times 0.94 + 0.029 \times 0.03 + 0.361 \times 0.20) \simeq 1192$.

As for mobility, according to IBOPE/IBGE, the socioeconomic distribution of households evolves from $(\mathbf{ABC,DE}) = (2443, 1347)$ in $t = 1$ to $(2511, 1269)$ in $t = 2$. It follows that, in $t = 2$: (i) $2511 - 2443 = 68$ households are newly affluent; (ii) $(2443 + 1347) - (2511 + 1269) = 10$ households migrated out of the urban area (again, this rarely happens in the data); and (iii) 1269 households are poor. Following Assumption 1, the 68 upwardly mobile households entering $t = 2$ are assigned to states in proportion to the choices of poor households in $t = 1$ among premium, generic and no soda (where these proportions are calculated as illustrated for established affluents, for which a proportion $1192/2443 \simeq 49\%$ chose premium rather than generic or no soda). These counts (summing 68) are deducted from the $t = 1$ poor population (1347) that is transitioning to $t = 2$ in proportion to the poor's choices across brand types. Similarly, following Assumption 2 (Migration), the 10 households leaving the city are dropped from the counts of the poor (totaling $1347 - 68$) in proportion to the poor's choices across brand types.

Region 4, $t = 10$ to $t = 11$. This example highlights mobility. The type-distribution vector following choices made in $t = 9$ and mobility into $t = 10$ is

$$\mathcal{F}_{g=4,t=10} = \{0.250, 0.109, 0.279, 0.003, 0.007, 0.137, 0.000, 0.002, 0.214\}$$

Having updated from $t = 1$, the history of choices and mobility now determines the distribution of states across each socioeconomic group. At $t = 10$, the generic-to-premium ratio is $0.109 : 0.250 = 0.4$ among the established affluent (see Table 5). From IBOPE/IBGE data, the mass of households by socioeconomic group (in thousands) in $t = 10$ is computed as: 2443 established affluent (this stays constant), 560 newly affluent and 826 poor (see Figure 4; $t = 10$ is the Jun-98/Jul-98 bimonth). The evolution of $(\mathbf{ABC,DE})$ from $(3003, 826)$ in $t = 10$ to $(3060, 784)$ in $t = 11$ implies that: (i) the newly affluent count grows by 57 (to 617); (ii) 16 migrants arrive at the city and join the ranks of the poor; and (iii) the poor count drops by $57 - 16 = 42$ (to 784). The 57 upwardly mobile households making choices with newly affluent status in $t = 11$ belong to states in proportion to the $t = 10$ choices of the poor they left behind (Assumption 1). The 16 migrants who are new to the city are assigned to the no-soda state (Assumption 2).

Robustness to Assumptions 1 and 2. Our results are robust to alternative assumptions, namely: (i) modifying Assumption 1 to assign households moving up from poor to newly affluent status to prior-consumption states in proportion to the previous-period soda choices of the newly affluents they are joining, rather than the poor they are leaving behind²³ (and analogously with respect to households moving down from newly affluent to poor status, based on previous-period choices by the poor); and (ii) modifying Assumption 2 to assign households moving to urban areas to prior-consumption states in proportion to the previous-period choices of the city-dwelling poor they are joining. For example, under (ii), $(\alpha_{EA}, \alpha_{NA}, \alpha)$ and λ are estimated, respectively, at $(3.62, 1.87, -5.76)$ and 4.21 (with standard errors of $(1.42, 1.45, 1.47)$ and 0.26), very close to baseline estimates (Table 4). Estimates under variant (i) are also very close to baseline.

²³The exception is the first transition, from $t = 1$ to $t = 2$, in which the newly affluent are a random sample of the poor as, by definition, there are no newly affluents in $t = 1$.

A.2 The estimation algorithm

We classify the parameters θ into “linear” and “non-linear” parameters, $\theta_1 = \{\beta, \alpha\}$ and $\theta_2 = \{\alpha_{EA}, \alpha_{NA}, \lambda\}$, consistent with familiar terminology from the literature. Given any generic value for the non-linear parameters θ_2 , steps 1 to 5 of the algorithm below yield an $N \times 1$ vector $\delta(\theta_2)$, containing the base utilities for every brand in every region-period market ($N = 9 \cdot 7 \cdot 57$). As noted in Section 4, conditioning on the full parameter vector $\theta = (\theta_1, \theta_2)$, one obtains an $N \times 1$ vector of demand unobservables by subtracting the systematic portion of the base utility from δ_{jgt} , i.e., $\xi_{jgt} = \delta_{jgt} - x'_{jgt}\beta - \alpha p_{jgt}$. Stacking these unobservables, we obtain:

$$\xi(\theta) = \delta(\theta_2) - X\theta_1$$

where the $N \times K_1$ matrix X contains the K_1 base-utility covariates (including price), and let K_2 denote the dimension of θ_2 . Now let Z denote a $N \times L$ matrix of instruments containing all covariates in X but price, as well as excluded instruments, where $L > K_1 + K_2$. Writing $W = (Z'Z)^{-1}$, the GMM objective is defined by:

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

Computation time can be reduced substantially by noting (see BLP 1995, Nevo 2000) that, conditional on θ_2 , there is a closed-form solution for θ_1 that minimizes the objective:

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

This allows us to maximize the objective by searching only over values for θ_2 . At every guess $\tilde{\theta}_2$ for the non-linear parameters, the GMM objective is evaluated via the following steps:

1. For every region $g = 1, \dots, 7$, and period $t = 1$, given \mathcal{F}_{g1} and $\tilde{\theta}_2$, use the BLP contraction mapping to solve for the unique vector of base utilities that matches observed aggregate market shares with those predicted by the model.
2. For every region $g = 1, \dots, 7$ and household type $r = 1, \dots, 9$, use equation (4), the base utilities recovered in step 1, and $\tilde{\theta}_2$, to predict the shares of type- r households who consume premium brands, generic brands or no soda in period $t = 1$.
3. For every region $g = 1, \dots, 7$, use the shares obtained in step 2, data on aggregate social mobility and migration, and Assumptions 1 and 2, to obtain next period’s type-distribution vector, \mathcal{F}_{g2} (recall Section 4.2).
4. Repeat steps 1-3 for periods $t = 2, \dots, 57$.
5. Stack the base utilities for all brands in all regions and periods in the $N \times 1$ vector $\delta(\tilde{\theta}_2)$, and evaluate the GMM objective at $\tilde{\theta}_2$, as explained above.

A.3 Monte Carlo experiments

Data generating process. Demand for soda follows the state-dependent household-level choice model developed in Section 4, namely indirect utility (2) with baseline persistence feature (3). We take as true parameters the point estimates reported in Table 4—these are reproduced in column I of Table A1. We design each simulated dataset to have the same dimensions as our empirical dataset: 9 brands (8 of which are premium), 7 regions and 57 time periods. This enables us to take covariates x_{jgt} , the first-period distribution of types \mathcal{F}_{g1} , and the evolution of each region’s household population by socioeconomic group as observed in the data (see definitions in Section 4). Assumptions 1 and 2 (“orthogonality”) dictate how demographic shifts interact with previous-period consumption. Prices are simulated according to:

$$p_{jgt} = \begin{cases} \Lambda_t^{prem} c_{jt}^{prem} + \rho \xi_{jgt} & \text{if } j \in \mathcal{A} \\ c_{jt}^{gen} & \text{if } j \in \mathcal{B} \end{cases}$$

where c_{jt}^{prem} and c_{jt}^{gen} are marginal costs for premium and generic brands, respectively, that are flat in output and vary across brands and over time, but not across regions (one can relax this); Λ_t^{prem} is a time-varying price markup over marginal cost for premium brands (this can also be made to vary across brands and regions); unobserved taste shocks ξ_{jgt} are i.i.d. across brands (including generics), regions and time; and $0 \leq \rho < 1$ is a pass-through ratio from utility shocks for premium brands to prices. Define marginal costs as the time-varying price of inputs W (e.g., sugar) times a cost efficiency (inverse productivity) parameter τ that varies by brand type (premium or generic) and over time, plus a brand-and-time varying disturbance term u :

$$\begin{aligned} c_{jt}^{prem} &= W_t \tau_t^{prem} + u_{jt} \\ c_{jt}^{gen} &= W_t \tau_t^{gen} + u_{jt} \end{aligned}$$

Specifically, the simulations reported in Table A1—except column V—consider price variation that is inspired by (is “comparable” to) what we observe in the real data: (i) premium brands’ markup $\Lambda_t^{prem} \sim N(2, 0.02^2)$ until April-May 1999 and, following the 20% price cut, $\Lambda_t^{prem} \sim N(1.6, 0.02^2)$ thereafter; (ii) factor price $W_t \sim U(0.85, 1.15)$; (iii) premium brands’ cost efficiency τ_t^{prem} equal to 0.55 throughout the sample period; (iv) generics’ cost efficiency τ_t^{gen} equal to 0.9 in the first period, decreasing linearly to 0.6 in August 2000, and constant thereafter; and (v) brand cost disturbance $u_{jt} \sim N(0, 0.01^2)$.

By contrast, the simulation reported in column V, marked “less simulated data variation versus real”: (i) drops the premium price cut, simulating the premium brands’ markup according to $\Lambda_t^{prem} \sim N(2, 0.02^2)$ throughout; (ii) drops the fringe’s price decline, setting the generics’ cost efficiency τ_t^{gen} already at 0.6 from the first period; and (iii) freezes each region’s household population by socioeconomic group (EA, NA, P) from the second period on.

Completing the description of the simulated datasets, we consider a pass-through ratio $\rho = 0.3$ and model utility shocks $\xi_{jgt} \sim N(0, \sigma^2)$ with varying orders of magnitude of variance, as reported in each column of Table A1. Notice that our experiments do not rely on the established firms pricing optimally, and that premium prices are endogenous since the established firms pass through a proportion of brand-market specific taste shocks to prices. Because taste shocks are i.i.d. and prices correlate across regions through a common cost (and markup) structure, prices

in one region are a valid instrument for prices in another region.

Estimation. To estimate using simulated data, we follow the estimation procedure proposed in Section 4. This includes adopting the same three instrument classes, namely: (i) the price of inputs, (ii) the contemporaneous mean price for a brand in the other regions, and (iii) a dummy variable indicating periods after July 1999 interacted with brand-region fixed effects. Since we calculate household-type specific brand shares analytically, there is no sampling variation in the logit shock ϵ_{ijgt} . For every simulation $\omega = 1, \dots, \Omega$, we complete the experimental—i.e., aggregate—dataset and make inference using the true (baseline) demand model (see columns II to V), as well as the No Persistence model variant (column VI).

Results. With $\sigma = 0.001$, the estimated baseline model recovers the true parameters to at least 2 decimal places (column II). Precision falls but is still quite high, particularly for $\hat{\lambda}$, when $\sigma = 0.1$ (column IV). Estimates using simulated data that exhibits less rich variation than what we observe empirically—namely, where we shut down the premium price cut, the generic price decline and the socioeconomic transitions—are substantially more noisy (column V compared to column IV). Estimates using the No Persistence model variant (column VI compared to column IV) are noisier and, importantly, price sensitivity appears biased downward (for brevity, the bottom of the table reports aggregate own-price elasticities only).

B Tables and Figures

Table 1: Brand Volume Shares of the Soda Category

Brand	Region 1 Northeast		Region 2 MG, ES, RJ		Region 3 RJ Metro		Region 4 SP metro		Region 5 SP		Region 6 South		Region 7 FD, GO, MS	
	t=1	t=57	t=1	t=57	t=1	t=57	t=1	t=57	t=1	t=57	t=1	t=57	t=1	t=57
Coke	0.40	0.21	0.37	0.26	0.36	0.26	0.32	0.24	0.31	0.25	0.28	0.31	0.34	0.29
Fanta	0.08	0.06	0.05	0.08	0.05	0.09	0.08	0.09	0.06	0.08	0.06	0.08	0.03	0.05
Kuat	0.03	0.06	0.02	0.04	0.04	0.06	0.07	0.06	0.06	0.04	0.04	0.04	0.05	0.03
Diet Coke	0.02	0.03	0.02	0.03	0.03	0.05	0.02	0.04	0.03	0.03	0.03	0.04	0.02	0.03
Other Coca-Cola	0.02	0.03	0.02	0.02	0.02	0.02	0.05	0.06	0.04	0.05	0.05	0.04	0.05	0.02
Guarana Antartica	0.17	0.09	0.07	0.06	0.06	0.09	0.09	0.10	0.08	0.07	0.05	0.05	0.09	0.10
Other Ambev	0.19	0.02	0.11	0.01	0.15	0.01	0.16	0.02	0.12	0.01	0.17	0.01	0.06	0.02
Pepsi	0.06	0.03	0.08	0.03	0.19	0.07	0.05	0.05	0.06	0.04	0.11	0.09	0.09	0.05
Generics	0.03	0.47	0.26	0.47	0.12	0.36	0.17	0.34	0.25	0.41	0.21	0.33	0.27	0.41

Volume shares of the soda category, by brand in each region in the first and last time periods. All statistics pertain to family-size bottles sold in supermarkets with checkouts. Coke, Fanta, Kuat, Diet Coke, and “Other Coca-Cola” are premium brands marketed by the Coca-Cola Company. Guarana Antartica, Pepsi, and “Other Ambev” are premium brands marketed by Ambev. Region 1 covers states in the Northeast excluding MA and PI. Region 2 covers the states of MG, ES and RJ excluding RJ metro. Region 5 covers the state of SP excluding SP metro. Region 6 covers the Federal District, and the states of GO and MS. Source: Nielsen.

Table 2: Soda Consumption by Socioeconomic Group (HEX)

Region of survey. cities	Socioeconomic group	Households $\times 1000$		Soda purchasing	By brand type		No soda
		Universe	%		Premium	Generic	
1 (Northeast)	ABC	696	36	28.0%	27.0%	0.9%	72.0%
	DE	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%

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. Metropolitan areas surveyed were: (Region 1) Recife, Fortaleza and Salvador; (Region 2) Belo Horizonte; (Region 3) Rio de Janeiro Metro; (Region 4) Sao Paulo Metro; (Region 6) Curitiba and Porto Alegre; (Region 7) Brasilia and Goiania. No city was surveyed in Region 5 (state of Sao Paulo excluding Sao Paulo Metro). We do not consider the northern city of Belem as it is located outside the area covered by Nielsen. Source: IBGE HEX 1995/96.

Table 3: Linking the Fringe's Volume Share to the Emerging Middle Class

	I	II	III	IV	V
Newly Affluent Share of All Households	0.351 (0.051)	0.463 (0.053)	0.439 (0.053)	0.450 (0.053)	0.181 (0.068)
Region DVs	No	Yes	Yes	Yes	Yes
Region-specific Time Trends	No	Yes	Yes	Yes	Yes
Seasonal DVs	No	No	Yes	Yes	Yes
Advertising GRPs (premium brands) \times Region DVs	No	No	No	Yes	Yes
Generic Brand In-store Presence \times Region DVs	No	No	No	No	Yes
Premium Brand In-store Presence (mean) \times Region DVs	No	No	No	No	Yes
No. Observations	399	399	399	399	399
No. Regressors	2	15	20	27	41
R-squared	0.107	0.768	0.778	0.782	0.867

Dependent variable is the generic brands' volume share of the soda category. Standard errors in parentheses. All regressions include a constant term. DVs denotes a set of dummy variables. Seasonality effects are bimonthly (December/January, February/March, etc).

Table 4: Baseline Demand Estimation Results

	coeff	(s.e.)		coeff	(s.e.)
Price Sensitivity Parameters			Parameter Combinations:		
α_{EA}	3.56	(1.39)	$\alpha + \alpha_{EA}$	-2.14	(0.17)
α_{NA}	1.80	(1.42)	$\alpha + \alpha_{NA}$	-3.90	(0.61)
α	-5.70	(1.43)	$\alpha_{EA} - \alpha_{NA}$	1.76	(0.52)
Persistence Parameter					
λ	4.20	(0.26)			
Other Effects					
Constant	-3.84	(0.37)			
Temperature	2.91	(0.32)			
Advertising GRPs:			In-store Presence:		
Advertising×Region 1	0.44	(0.22)	Distribution×Region 1	3.10	(0.27)
Advertising×Region 2	0.25	(0.37)	Distribution×Region 2	3.70	(0.46)
Advertising×Region 3	0.00	(0.27)	Distribution×Region 3	4.09	(0.96)
Advertising×Region 4	0.40	(0.37)	Distribution×Region 4	2.72	(0.45)
Advertising×Region 5	0.32	(0.32)	Distribution×Region 5	3.33	(0.48)
Advertising×Region 6	0.66	(0.34)	Distribution×Region 6	1.24	(0.28)
Advertising×Region 7	0.49	(0.32)	Distribution×Region 7	1.11	(0.33)
Time Trends:					
Region 1	-0.37	(0.12)			
Region 2	0.13	(0.19)			
Region 3	-0.33	(0.12)			
Region 4	-0.66	(0.09)			
Region 5	0.02	(0.17)			
Region 6	-0.05	(0.17)			
Region 7	-0.07	(0.21)			
Seasonality×Brand Type Effects		Yes			
Brand-Region Fixed Effects		Yes			

Standard errors in parentheses. Source: estimated baseline model.

Table 5: Predicted Consumption Patterns and Transitions by Type

A: Predicted Soda Penetration and Consumption Patterns by Socioeconomic Group*

	Soda Penetration	Premium Share	Generic Share	Generic:Premium Ratio
EA	0.69	0.47	0.23	0.5
NA	0.22	0.09	0.13	1.4
P	0.05	0.01	0.04	2.7

* Soda penetration, for the entire category and by type of brand, in each socioeconomic group.

B: Transitions by Household Type**

	First 12 months of the sample, prior to premium price cut:			Last 12 months of the sample, after price declines:		
	Premium Soda	Generic Soda	No Soda	Premium Soda	Generic Soda	No Soda
EA^A	0.85	0.01	0.14	0.90	0.01	0.09
EA^B	0.02	0.79	0.19	0.02	0.87	0.11
EA^O	0.08	0.06	0.86	0.14	0.09	0.78
NA^A	0.47	0.01	0.52	0.69	0.01	0.30
NA^B	0.01	0.56	0.44	0.01	0.75	0.24
NA^O	0.01	0.02	0.97	0.04	0.04	0.92
P^A	0.12	0.01	0.87	0.34	0.01	0.65
P^B	0.00	0.27	0.73	0.00	0.53	0.46
P^O	0.00	0.00	0.99	0.01	0.02	0.98

** Reported transition rates are means across all regions in the first 12 sample months vs. last 12 months

Table 6: Demand elasticities

Aggregate Own-Price Elasticities		Household-Type Specific Own-Price Elasticities			
Coke	-1.65	Coke, EA^A	-1.42	Generic, EA^A	-1.33
Guaraná Antarctica	-2.06	Coke, EA^B	-2.29	Generic, EA^B	-0.23
Fanta	-2.01	Coke, EA^O	-2.18	Generic, EA^O	-1.24
Pepsi	-2.01	Coke, NA^A	-3.29	Generic, NA^A	-2.42
Generic	-0.68	Coke, NA^B	-4.20	Generic, NA^B	-0.86
		Coke, NA^O	-4.16	Generic, NA^O	-2.37
		Coke, P^A	-5.70	Generic, P^A	-3.56
		Coke, P^B	-6.14	Generic, P^B	-2.14
		Coke, P^O	-6.14	Generic, P^O	-3.54

Elasticities are means across region-and-time markets, only a few elasticities shown due to space constraints.

Table 7: Estimates from Demand Model Variants

	coeff	(s.e.)		coeff	(s.e.)
A: "Soda Category Persistence" Demand Model Variant					
Price Sensitivity Parameters			Parameter Combinations:		
α_{EA}	-0.09	(0.61)	$\alpha + \alpha_{EA}$	-1.81	(0.31)
α_{NA}	-0.41	(0.33)	$\alpha + \alpha_{NA}$	-2.14	(0.53)
α	-1.72	(0.33)	$\alpha_{EA} - \alpha_{NA}$	0.33	(0.76)
Persistence Parameter					
λ	3.30	(0.50)			
Other Effects					
Constant	-4.24	(0.37)			
Temperature	1.97	(0.21)			
B : "No Persistence" Demand Model Variant					
Price Sensitivity Parameters			Parameter Combinations:		
α_{EA}	0.47	(0.58)	$\alpha + \alpha_{EA}$	-1.51	(0.19)
α_{NA}	-0.85	(0.35)	$\alpha + \alpha_{NA}$	-2.83	(0.65)
α	-1.98	(0.42)	$\alpha_{EA} - \alpha_{NA}$	1.32	(0.77)
Other Effects					
Constant	-2.72	(0.24)			
Temperature	1.64	(0.11)			

Following the baseline model, both model variants include: Advertising×Region, In-store Presence×Region, Time Trend×Region, Seasonality×Brand Type Effects, and Brand-Region Fixed Effects. Standard errors in parentheses. Source: estimated soda category persistence model variant, estimated no persistence model variant.

Table A1: Monte Carlo Experiments

	I	II		III		IV		V		VI	
Model		Baseline		Baseline		Baseline		Baseline		NP	
Simulated vs. real		'Comparable'		'Comparable'		'Comparable'		'Less'		'Comparable'	
σ	True	Median	(SD)	Median	(SD)	Median	(SD)	Median	(SD)	Median	(SD)
Price sensitivity											
α_{EA}	3.56	3.56	(0.00)	3.57	(0.04)	3.55	(0.38)	3.20	(0.71)	4.71	(0.90)
α_{NA}	1.80	1.80	(0.00)	1.80	(0.04)	1.80	(0.37)	0.67	(23.80)	3.15	(0.62)
α	-5.70	-5.70	(0.00)	-5.71	(0.04)	-5.66	(0.42)	-5.24	(0.65)	-5.55	(0.95)
Combinations:											
$\alpha + \alpha_{EA}$	-2.14	-2.14	(0.00)	-2.14	(0.01)	-2.09	(0.06)	-2.04	(0.14)	-0.76	(0.15)
$\alpha + \alpha_{NA}$	-3.90	-3.90	(0.00)	-3.90	(0.01)	-3.85	(0.14)	-4.34	(23.82)	-2.31	(0.59)
$\alpha_{EA} - \alpha_{NA}$	1.76	1.76	(0.00)	1.76	(0.01)	1.76	(0.10)	2.27	(23.83)	1.47	(0.56)
Persistence											
λ	4.20	4.20	(0.00)	4.20	(0.01)	4.16	(0.09)	4.18	(0.15)	-	-
Other Effects											
Constant	-3.84	-3.84	(0.00)	-3.84	(0.01)	-3.82	(0.09)	-3.87	(0.17)	-2.47	(0.20)
Temperature	2.91	2.91	(0.00)	2.91	(0.01)	2.85	(0.07)	2.82	(0.07)	1.66	(0.21)
Advertising×Region 1	0.44	0.44	(0.00)	0.44	(0.01)	0.43	(0.07)	0.45	(0.07)	0.22	(0.08)
Advertising×Region 2	0.25	0.25	(0.00)	0.25	(0.01)	0.25	(0.08)	0.25	(0.08)	0.21	(0.08)
Advertising×Region 3	0.00	0.00	(0.00)	0.00	(0.01)	0.01	(0.08)	0.02	(0.08)	0.02	(0.09)
Advertising×Region 4	0.40	0.40	(0.00)	0.40	(0.01)	0.40	(0.09)	0.39	(0.09)	0.48	(0.08)
Advertising×Region 5	0.32	0.32	(0.00)	0.32	(0.01)	0.33	(0.08)	0.32	(0.08)	0.36	(0.09)
Advertising×Region 6	0.66	0.66	(0.00)	0.66	(0.01)	0.67	(0.09)	0.67	(0.09)	0.82	(0.09)
Advertising×Region 7	0.49	0.49	(0.00)	0.49	(0.01)	0.47	(0.07)	0.48	(0.07)	0.48	(0.09)
Distribution×Region 1	3.10	3.10	(0.00)	3.10	(0.01)	3.11	(0.06)	3.12	(0.06)	4.12	(0.12)
Distribution×Region 2	3.70	3.70	(0.00)	3.70	(0.01)	3.69	(0.10)	3.69	(0.10)	3.62	(0.11)
Distribution×Region 3	4.09	4.09	(0.00)	4.09	(0.01)	4.08	(0.14)	4.09	(0.13)	4.59	(0.21)
Distribution×Region 4	2.72	2.72	(0.00)	2.72	(0.01)	2.73	(0.12)	2.72	(0.12)	2.72	(0.15)
Distribution×Region 5	3.33	3.33	(0.00)	3.33	(0.01)	3.34	(0.12)	3.33	(0.12)	3.31	(0.11)
Distribution×Region 6	1.24	1.24	(0.00)	1.24	(0.01)	1.25	(0.07)	1.24	(0.07)	1.41	(0.08)
Distribution×Region 7	1.11	1.11	(0.00)	1.11	(0.01)	1.11	(0.06)	1.11	(0.06)	0.97	(0.12)
Time Trend Region 1	-0.37	-0.37	(0.00)	-0.37	(0.00)	-0.36	(0.03)	-0.38	(0.02)	0.03	(0.15)
Time Trend Region 2	0.13	0.13	(0.00)	0.13	(0.00)	0.13	(0.03)	0.13	(0.02)	0.24	(0.21)
Time Trend Region 3	-0.33	-0.33	(0.00)	-0.33	(0.00)	-0.32	(0.02)	-0.32	(0.02)	0.29	(0.13)
Time Trend Region 4	-0.66	-0.66	(0.00)	-0.66	(0.00)	-0.65	(0.03)	-0.65	(0.02)	-0.20	(0.10)
Time Trend Region 5	0.02	0.02	(0.00)	0.02	(0.00)	0.03	(0.04)	0.02	(0.02)	0.52	(0.19)
Time Trend Region 6	-0.05	-0.05	(0.00)	-0.05	(0.00)	-0.05	(0.03)	-0.05	(0.02)	0.37	(0.16)
Time Trend Region 7	-0.07	-0.07	(0.00)	-0.07	(0.01)	-0.07	(0.06)	-0.08	(0.03)	0.47	(0.22)
Aggregate Own-Price Elasticities											
Coke		-1.33	(0.02)	-1.33	(0.02)	-1.31	(0.03)	-1.43	(0.04)	-0.83	(0.11)
Guaraná Antarctica		-1.98	(0.02)	-1.98	(0.02)	-1.94	(0.05)	-2.09	(0.07)	-0.96	(0.14)
Fanta		-2.02	(0.02)	-2.02	(0.02)	-1.99	(0.05)	-2.15	(0.07)	-0.97	(0.14)
Pepsi		-2.06	(0.02)	-2.06	(0.02)	-2.02	(0.05)	-2.19	(0.07)	-0.98	(0.14)
Generic		-0.75	(0.01)	-0.75	(0.01)	-0.75	(0.01)	-0.55	(0.02)	-0.77	(0.09)

Medians and standard deviations (in parentheses) of estimated parameters are taken over $\Omega = 50$ simulations. Reported elasticities are means across region-and-time markets.

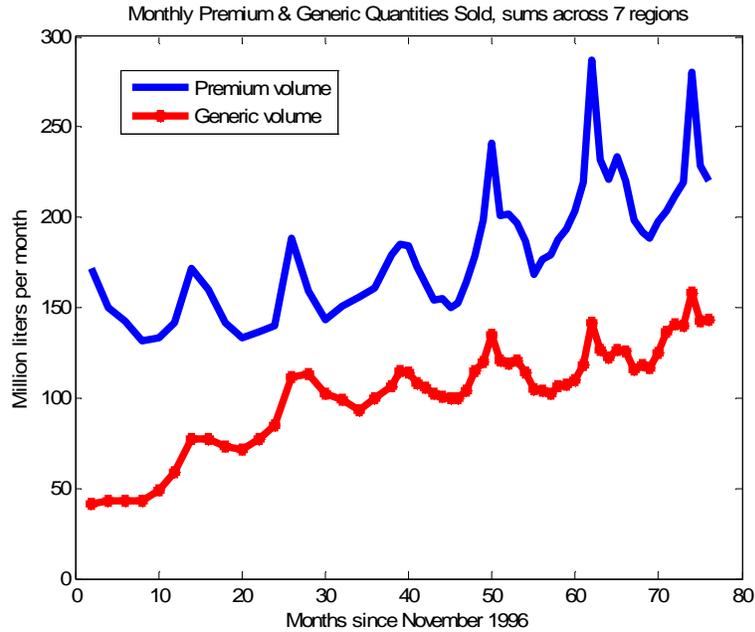


Figure 1: The evolution of quantities (in million liters/month) by type of brand (Premium versus Generic), for soda sold in family-size bottles through the self-service channel across the seven Nielsen regions, in the period Dec-96 to Mar-03. Source: Nielsen.

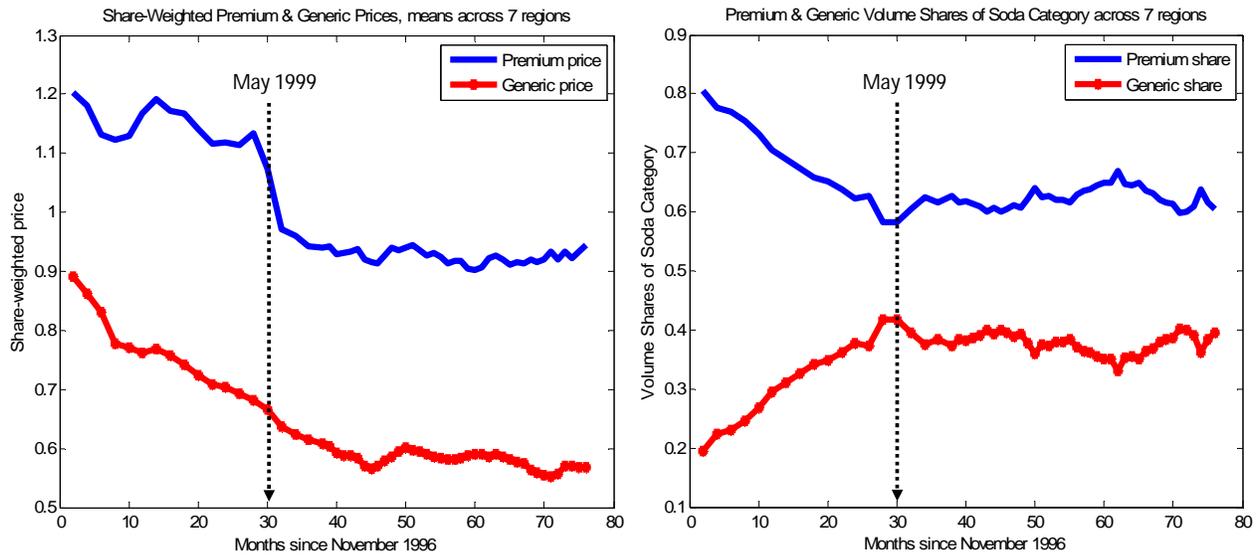


Figure 2: The evolution of prices (in constant Brazilian R\$/liter) and category volume shares (in percent, summing to one) by type of brand (Premium versus Generic), for soda sold in family-size bottles through the self-service channel across the seven Nielsen regions, in the period Dec-96 to Mar-03. Source: Nielsen.

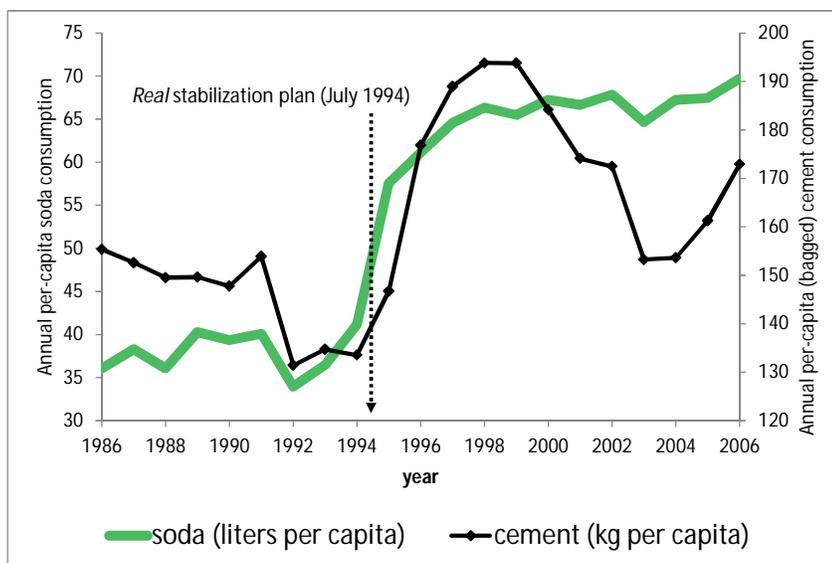


Figure 3: Annual aggregate per capita consumption of soft drinks (in liters per person) and of bagged cement (in kilograms per person). (Cement sold in bags, as opposed to sales in bulk, filter out any large-scale construction activity, such as government spending on infrastructure.) Sources: Brazilian trade associations for soft drink makers and for cement producers, ABIR and SNIC respectively.

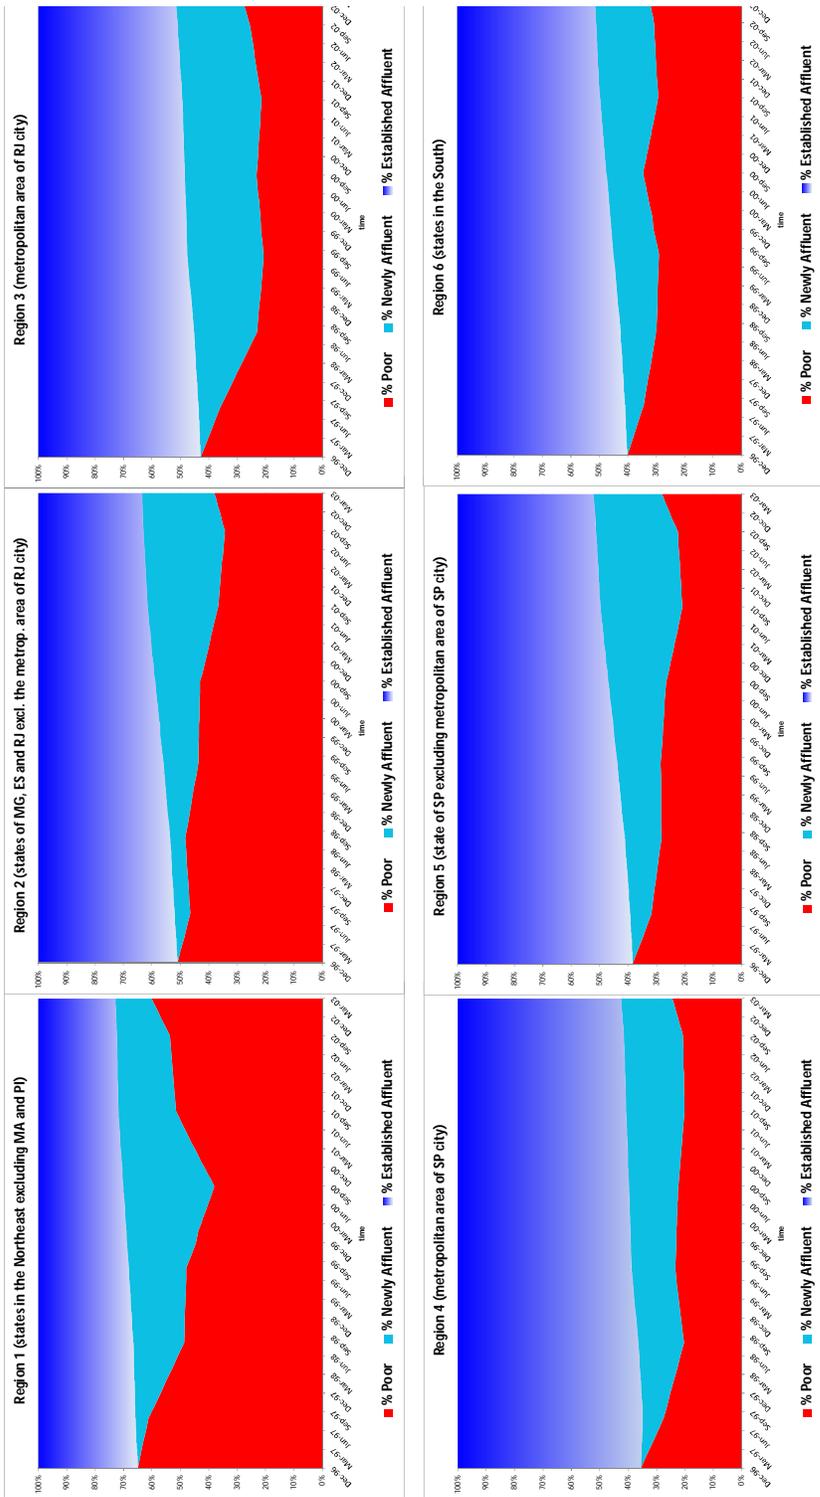


Figure 4: The rise of “newly affluent” households: Proportion of urban households in each of three constructed socioeconomic groups (“Established Affluent,” “Newly Affluent” and “Poor”), as defined in the text, by region in the period Dec-96 to Mar-03. The smallest region by number of households, region 7 (Federal District and states of GO and MS), is not shown for lack of space (the pattern is similar to region 2). Sources: IBOPE LatinPanel, IBGE PNAD.

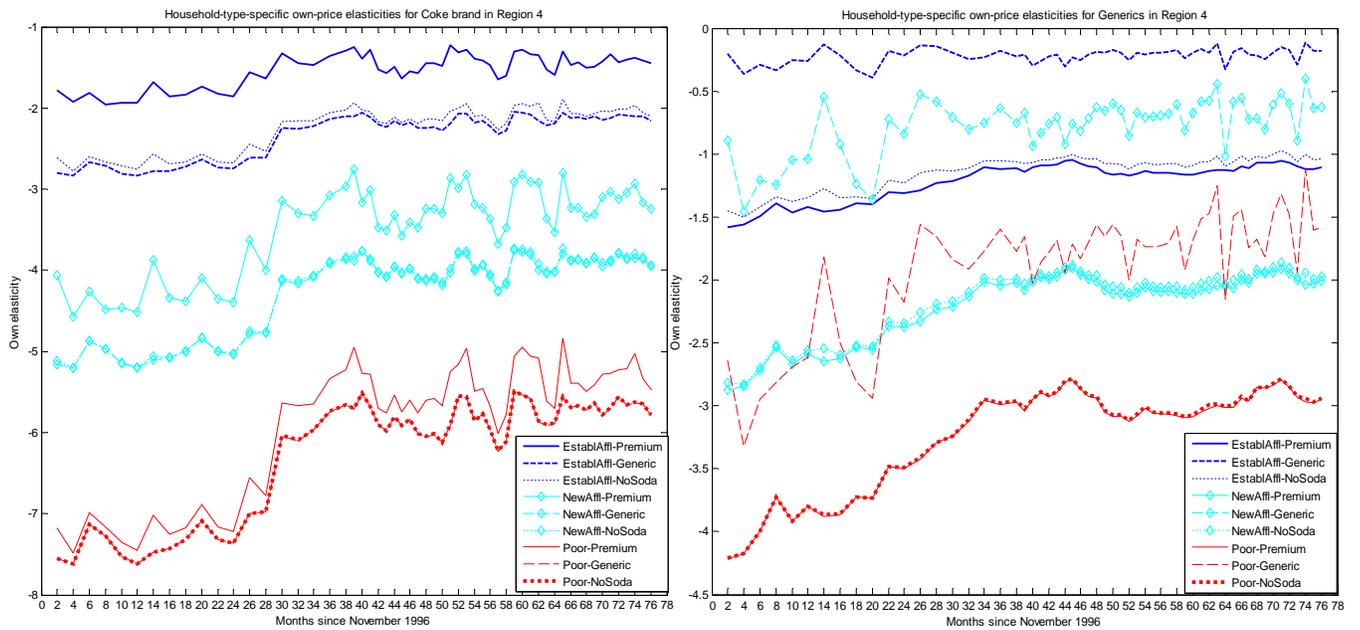


Figure 5: Evolution of own-price elasticities for Coke brand (left panel) and Generics (right panel), by household type, in region 4 (São Paulo Metro). Source: baseline model.

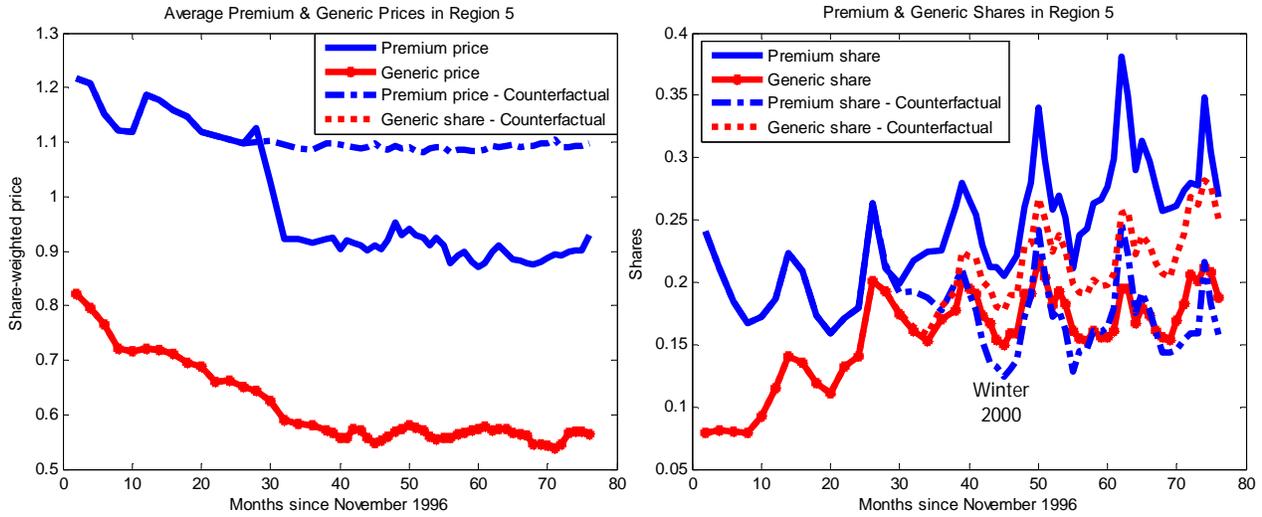


Figure 6: Actual against counterfactual price and share paths for premium brands and generic brands in region 5 (São Paulo Interior). Prices in the left panel and shares in the right panel. The counterfactual scenario considers premium brands not cutting prices in mid 1999. Source: baseline model.

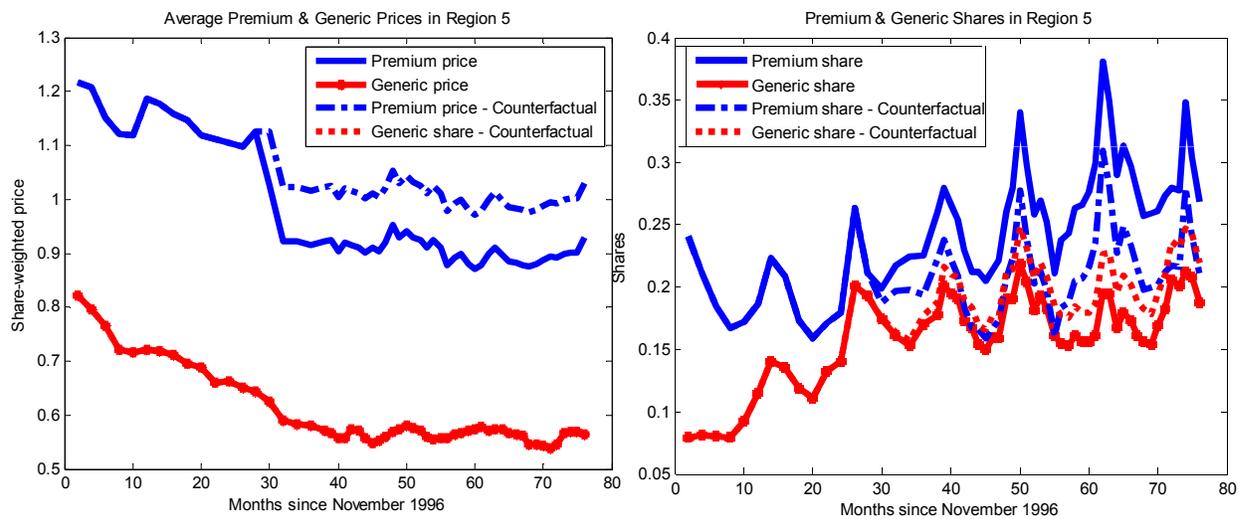


Figure 7: Share paths for premium brands and generic brands, in region 5, for a counterfactual in which premium brand prices are cut by only 10% in mid 1999. Source: baseline model.

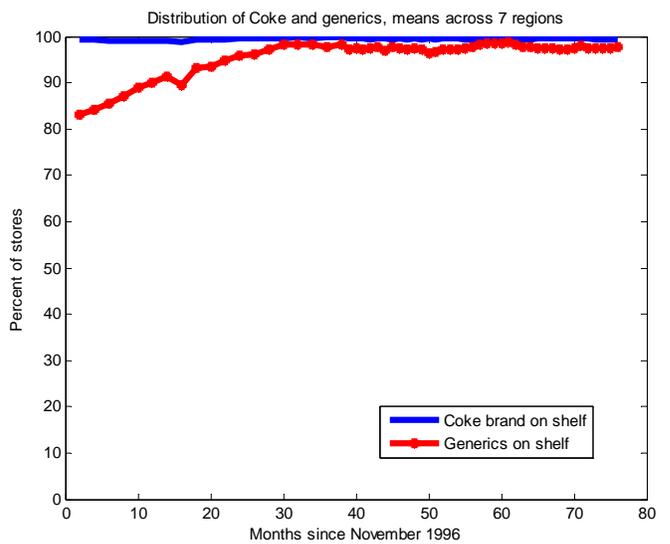
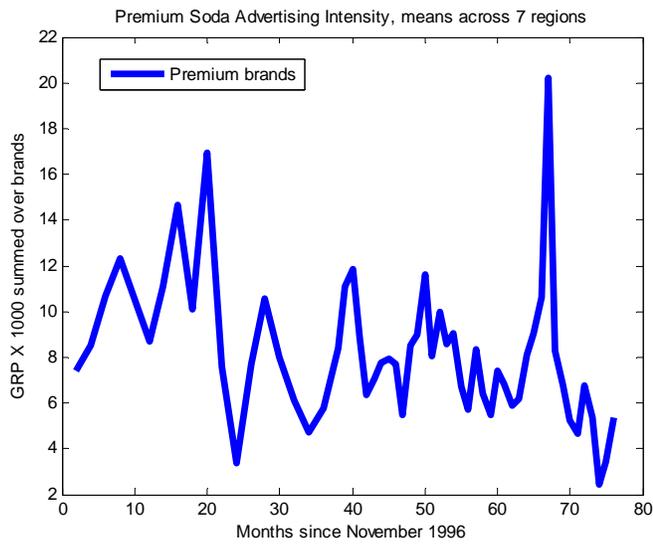


Figure A1: The evolution of premium soda media advertising intensity (in monthly Gross Rating Points, GRP, $\times 1000$) and the proportions of stores with specific brands in stock (in percent), in the period Dec-96 to Mar-03. GRP and percentages shown are means across the seven Nielsen regions. Sources: McCann-Erickson, Nielsen.