American Economic Review 2012, 102(6): 2472–2508
http://dx.doi.org/10.1257/aer.102.6.2472

The Evolution of Brand Preferences: Evidence from Consumer Migration[†]

By Bart J. Bronnenberg, Jean-Pierre H. Dubé, and Matthew Gentzkow*

We study the long-run evolution of brand preferences, using new data on consumers' life histories and purchases of consumer packaged goods. Variation in where consumers have lived in the past allows us to isolate the causal effect of past experiences on current purchases, holding constant contemporaneous supply-side factors. We show that brand preferences form endogenously, are highly persistent, and explain 40 percent of geographic variation in market shares. Counterfactuals suggest that brand preferences create large entry barriers and durable advantages for incumbent firms and can explain the persistence of early-mover advantage over long periods. (*JEL* D12, L11, M31, M37)

Consumers appear to have high willingness to pay for particular brands, even when the alternatives are objectively similar. The majority of consumers typically buy a single brand of beer, cola, or margarine (Dekimpe et al. 1997), even though relative prices vary significantly over time, and consumers often cannot distinguish their preferred brand in blind "taste tests" (Thumin 1962; Allison and Uhl 1964). Consumers pay large premia to buy homogeneous goods like books and CDs from branded online retailers, even when they are using a "shopbot" that eliminates search costs (Smith and Brynjolfsson 2001). A large fraction of consumers buy branded medications, even though chemically equivalent generic substitutes are available at the same stores for much lower prices (Ling, Berndt, and Kyle 2002).

Theorists have long speculated that willingness to pay for brands today could depend on consumers' experiences in the past. Willingness to pay could be a

* Bronnenberg: Tilburg School of Economics and Management, PO Box 90153, 5000 LE Tilburg, the Netherlands (e-mail: bart.bronnenberg@uvt.nl); Dubé: University of Chicago Booth School of Business, 5807 South Woodlawn Avenue, Office 361, Chicago, IL 60637 (e-mail: jdube@chicagobooth.edu); Gentzkow: University of Chicago Booth School of Business, 5807 South Woodlawn Avenue, Office 514, Chicago, IL 60637 (e-mail: matthew.gentzkow@chicagobooth.edu). We thank Pradeep Chintagunta, Aimee Drolet, Jon Guryan, Emir Kamenica, Kevin Murphy, Fiona Scott Morton, Jesse Shapiro, Chad Syverson, and participants at the INFORMS Marketing Science Conference in Ann Arbor, Michigan, the second Workshop on the Economics of Advertising and Marketing in Paris, France, the NBER Summer Institute (IO) and the 2010 QME Conference for helpful comments. We gratefully acknowledge feedback from seminar participants at Boston College, the Einaudi Institute of Economics and Finance, Erasmus University Rotterdam, Goethe University Frankfurt, Hong Kong University of Science and Technology, London Business School, London School of Economics, Stanford University, Tel-Aviv University, University of California, Los Angeles, the University of Chicago, Universidade Nova de Lisboa, and University of Western Ontario. We thank Grace Hyatt and Todd Kaiser at Nielsen for their assistance with the collection of the data, and the Marketing Science Institute, the Neubauer Family Foundation, the Netherlands Organization for Scientific Research (NWO Vici Grant), and the Initiative on Global Markets at the University of Chicago Booth School of Business for financial support.

[†] To view additional materials, visit the article page at http://dx.doi.org/10.1257/aer.102.6.2472.

function of past consumption, which could enter expected utility directly (Becker and Murphy 1988), through switching costs (Klemperer 1987), or through beliefs about quality (Schmalensee 1982). It could depend on past exposure to advertising (Schmalensee 1983; Doraszelski and Markovich 2007), or on past observations of the behavior of others, as in Ellison and Fudenberg (1995). At the extreme, brand preferences could be entirely determined by experiences in childhood (Berkman, Lindquist, and Sirgy 1997). Under these assumptions, consumers' accumulated stock of "preference capital" could be a valuable asset for incumbent firms and a source of long-term economic rents.[1] In Bain's (1956) view, "the advantage to established sellers accruing from buyer preferences for their products as opposed to potential entrant products is on average larger and more frequent in occurrence at large values than any other barrier to entry" (p. 216).

Existing empirical evidence provides little support for the view that past experiences have a long-lasting impact on brand preferences. Large literatures have measured the effects of advertising, but these studies often find no effects (e.g., Lodish et al. 1995), and the effects they do measure are estimated to dissipate over a horizon ranging from a few weeks to at most five or six months (Assmus, Farley, and Lehmann 1984; Bagwell 2007). Empirical studies of habit formation and consumer switching costs have been limited to estimating short-run effects over horizons of at most one or two years (e.g., Erdem 1996; Keane 1997; Dubé, Hitsch, and Rossi 2010).

In this article, we study the long-run evolution of brand preferences, using a new dataset that combines Nielsen Homescan data on purchases of consumer packaged goods with details of consumers' life histories. Building on Bronnenberg, Dhar, and Dubé's (2007) finding that market shares of these goods vary significantly across regions of the United States, we ask how consumers' current purchases depend on both where they live currently, and where they lived in the past. This approach allows us to hold constant contemporaneous supply-side factors such as quality, availability, and advertising, and to isolate the causal effect of past experience on current purchases.

Our data include current and past states of residence for more than 38,000 households, which we match to 2006–2008 purchases in 238 consumer packaged goods product categories. Our primary dependent variable consists of the purchases of the top brand as a share of purchases of either of the top two brands in a category. Consistent with Bronnenberg, Dhar, and Dubé (2007), we show that this share varies significantly across space, with a mean of 0.63 and a cross-state standard deviation of 0.15 in the average product category.

We find strong evidence that past experiences are an important driver of current consumption. We first examine the way consumption patterns change when consumers move across state lines. Both cross-sectional and panel evidence suggest that approximately 60 percent of the gap in purchases between the origin and destination state closes immediately when a consumer moves. So, for example, a consumer who moves from a state where the market share of the top brand among lifetime residents is X percent to one where the market share is Y percent jumps from consuming X percent to consuming $(0.4X + 0.6Y)$ percent. Since the stock of past experiences has remained constant across the move, while the supply-side environment has changed,

[1] Throughout the article, we use "brand preferences" as a shorthand for willingness to pay. We intend this term to encompass channels such as learning that do not work through the utility function *per se*.

we infer that approximately 40 percent of the geographic variation in market shares is attributable to persistent brand preferences, with the rest driven by contemporaneous supply-side variables. We next look at how consumption evolves over time following a move. The remaining 40 percent gap between recent migrants and lifetime residents closes steadily, but slowly. It takes more than 20 years for half of the gap to close, and even 50 years after moving the gap remains statistically significant. Finally, we show that our data also strongly reject the hypothesis that *all* that matters is where consumers lived in childhood: consumers who move after age 25 still eventually converge to the consumption patterns of their new state of residence.

As a lens through which to interpret these results, we introduce a simple model of consumer demand with habit formation (Pollak 1970; Becker and Murphy 1988). Consumers in the model are myopic. Their choices in each period depend on the contemporaneous prices, availability, and other characteristics of the brands in their market, and on their stock of past consumption experiences, or "brand capital." The model has two key parameters: the weight on current product characteristics relative to the stock of past consumption (α), and the year-to-year persistence of brand capital (δ).

We next present evidence for two key identifying assumptions. The first is that a consumer's migration status is orthogonal to stable determinants of brand preferences. Panel evidence shows directly that migrants look similar to nonmigrants in their birth state before moving, and that age at migration is uncorrelated with purchases prior to moving. As additional evidence, we consider a subset of brands that were introduced late in our sample, and show that where a consumer lived before a brand pair was available does not predict her current consumption. The second assumption is that a brand's past market share in a given market is equal in expectation to the share today. We introduce historical data on market shares and show that, despite large changes over time in shares, the identifying assumption is approximately satisfied.

Under these two assumptions, we estimate that the weight on current characteristics in utility is $\alpha = 0.626$ and that the effect of a given year's consumption experiences depreciates at a rate of $1 - \delta = 0.025$ per year.

To shed more light on the economic implications of our findings, we simulate two counterfactual scenarios. First, we imagine that two brands enter a market sequentially, and we ask how difficult it will be for the second brand to equalize the market share advantage of the first. We show that a head start of even a few years creates a formidable barrier, with a second entrant needing to maintain a large advantage in supply-side variables (lower prices, more promotions, etc.) to catch up in subsequent years. Second, we introduce a simple model of endogenous firm choices and use it to study the persistence of brand advantages in the face of idiosyncratic shocks. We show that even with significant noise in the environment, our estimates can easily rationalize persistence of market shares over many decades, as observed in Bronnenberg, Dhar, and Dubé (2009).

In the final section, we present evidence on the specific mechanisms that underlie our results. We show that the relative importance of brand capital is higher in categories with high levels of advertising and high levels of social visibility. Although we cannot interpret these relationships as causal, they are consistent with a model in which both advertising and observed consumption of peers make the stock of brand capital more important. At the same time, we observe substantial persistence even in categories where advertising and visibility are low, suggesting that some element of habit formation is likely necessary to rationalize the data. We also assess how much of the geographic variation in shares not explained by brand capital can be attributed to variation in prices, display advertising, feature advertising, and availability.

Our empirical strategy is closely related to work that uses migration patterns to study the formation of culture and preferences. Logan and Rhode (2010) show that nineteenth-century immigrants' expenditure shares for different types of food are predicted by past relative prices in their countries of origin. Luttmer and Singhal (2011) link immigrants' preferences for redistribution of wealth to the average preference for redistribution in their birth countries. Atkin (2010) shows that migrants within India are willing to pay higher prices to consume foods that are common in their state of origin. Our results also relate to the literature on the formation of preferences more broadly (Bowles 1998). Our work further relates to the broader literature on sources of entry barriers and incumbent advantages (e.g., Bain 1950; Williamson 1963). In particular, Foster, Haltiwanger, and Syverson (2010) show that the demand curves of manufacturing plants shift out over time, and that a model of endogenous demand-side capital formation similar to the one we develop herein can explain a significant share of older plants' size advantage relative to newer plants. Finally, our work relates to the conceptual literature on the long term effects of brand equity in marketing (e.g., Aaker 1991; Keller 1993).

Section I introduces our data. Section II presents descriptive evidence on the evolution of brand preferences. Section III introduces our model and estimation strategy. Section IV presents evidence supporting our key identifying assumptions. Section V presents estimates of the model parameters and derives implications for first-mover advantage and share stability. Section VI presents evidence on mechanisms. Section VII concludes.

I. Data

A. *Purchases and Demographics*

We use data from the Nielsen Homescan Panel on the purchases and demographic characteristics of 48,501 households. The panel is drawn from 50 states and covers purchases made between October 2006 and October 2008, inclusive. Each household receives an optical scanner and is directed to scan the barcodes of all consumer packaged goods they purchase, regardless of outlet. The data thus include purchases not only from supermarkets, but also from convenience stores, drug stores, and so on. The data cover food, beverages, and many nonfood items commonly found in supermarkets. See Einav, Leibtag, and Nevo (2010) for a recent validation study of the Homescan Panel.

The most granular notion of a product in the data is a UPC code. Nielsen groups UPCs into categories called modules. Examples include "canned soup," "regular cola," "cough drops," and "bar soap." Nielsen also groups UPCs by brand, with Coca-Cola 12-ounce cans and Coca-Cola 2-liter bottles both grouped under the brand "Coca-Cola." A single brand may span multiple modules. Our raw data include 382 modules and 51,316 brands.

We define the total number of purchases by a household of a particular module-brand combination to be the number of observed shopping trips on which the household purchased at least one UPC in that module-brand. A trip counts as a single purchase regardless of the size, number of units bought, or price paid. In Appendix B we show that our results are robust to alternative quantity measures.

We rank brands within each module by the total number of purchases across all households in the sample. Our main analysis focuses on the top two brands in each module. We refer to the best-selling brand in a module as brand 1 and to the second-best-selling brand as brand 2, respectively.[2]

For each household, we observe a vector of demographics that includes household income, whether the household's residence is rented or owned, and the household head's race and Hispanic status.

B. *Consumer Life Histories*

We supplement the purchase and demographic data with a survey of Homescan panelists' life histories, which we administered in cooperation with AC Nielsen. The survey was sent electronically to households in the panel, and we requested that each adult in the household complete the survey separately. The questionnaire asked individuals their country and state of birth, and their current state of residence. For those not currently living in their state of birth, we asked the age at which they left their state of birth, and the number of years that they have lived in their current state. Respondents also reported their gender, their date of birth, their highest level of educational attainment, whether they are currently employed, whether they personally make the majority of the household's purchase decisions (whether they are the "primary shopper"), and whether they are the "head of household."

The survey was sent to 75,221 households. From these, 80,077 individuals in 48,951 households responded for a response rate of 65 percent. The surveys were completed between September 13, 2008, and October 1, 2008.

From each household, we select a single individual whose characteristics we match to the purchase data. We first focus on individuals born in the United States. For the set of households with multiple respondents, we then apply the following criteria in order, stopping at the point when only a single individual is left: (i) keep only primary shopper(s) if at least one exists; (ii) keep only household head(s) if at least one exists; (iii) keep only the female household head if both a female and a male head exist; (iv) keep the oldest individual; (v) drop responses that appear to be duplicate responses by the same individual; (vi) select one respondent randomly.

We define a household to be a *nonmigrant* if the selected individual's current and birth state are the same and a *migrant* otherwise.

We use the reported birth date to define a respondent's age, assuming all surveys were completed on September 22, 2008. We define the "gap" in a consumer's reported history to be the difference between her age and the sum of the number of years she lived in her birth state and the number of years she has lived in her current

[2] In our data, we observe sales for both top brands in a typical state-module combination. However, for some state-modules, we observe sales for only one of the two brands. In Appendix B we show that our results are robust to focusing only on those state-module combinations where we observe sales for both top brands.

state. In cases where the sum of a respondent's reported years living in her birth state and current state exceeds her age (i.e., the gap is negative), we either recode the number of years lived in her birth state to be the difference between her age and the reported years in her current state (if the difference is only one or two years), or drop the household from the data (if the difference is more than two years).

C. *Additional Data Sources*

We supplement our core dataset with data on the historical market shares of a subset of the brands in our data from Consolidated Consumer Analysis (CCA). These volumes are published jointly by a group of participating newspapers from 1948 to 1968.[3] They aggregate results from consumer surveys conducted by the newspapers in their respective markets. For each product category and market, the surveys give the share of consumers who report purchasing each brand.[4] We match these brand-category pairs to brand-module pairs in the Nielsen data. We collapse to the state level, averaging each brand's share purchasing across years from 1948 to 1968 and across markets within states. We then define each brand's average share to be the share of consumers purchasing divided by the sum of this share across brands within the category.

To interpret our counterfactuals in terms of equivalent price changes, we use aggregate store-level data on 2001–2005 purchases and prices, spanning 30 product categories from the Information Resources, Inc. (IRI) Marketing Data Set (Bronnenberg, Kruger, and Mela 2008).

To measure module-level advertising intensity, we use data on 2008 advertising expenditures for each module from the TNS Media Intelligence Ad$pender database. We download total expenditures for each top-two Homescan brand in our sample, treating cases where no TNS data exist for the brand in question as zeros. We then sum expenditures by module and code the top 25 percent of modules by advertising expenditure as "high advertising."

D. *Final Sample Definition and Sample Characteristics*

We exclude modules from the main analysis in which we do not observe at least 5,000 households making purchases. We also exclude a small number of modules in which the top two brands as defined by Nielsen are in fact two varieties of a single brand (e.g., "Philadelphia" and "Philadelphia Light" in the Cream Cheese module). We exclude migrant households for which the gap as defined above is greater than five years. We also exclude individuals with a reported age less than 18 or greater than 99. Our final sample consists of 38,098 households and 238 modules. See Table A3 for a list of these modules.

[3] From 1948–1950, the *Milwaukee Journal* is listed as publisher. In 1948, the title is *Thirteen Market Comparison of Consumer Preferences.* In 1949 and 1950, the title is *Fourteen Market Comparison of Consumer Preferences.* From 1951 to 1968, all of the participating newspapers are listed as publisher (the exact set of newspapers varies by year). In 1951 and 1952 the title is *Consolidated Consumer Analysis Information*, and from 1953 to 1968 the title is *Consolidated Consumer Analysis.*

[4] Until 1958, consumers were asked to report the brand they "usually buy" in each category. From 1959 on, they were asked to report the brand they "bought last."

Table 1—Migration Patterns

	Region of residence			
Region of birth	Northeast	Midwest	South	West
Northeast	6,765	269	1,539	448
Midwest	165	10,654	1,377	885
South	193	435	9,725	292
West	56	214	341	4,740

Note: Table shows the number of households in the Nielsen Homescan sample by census region of birth and current residence.

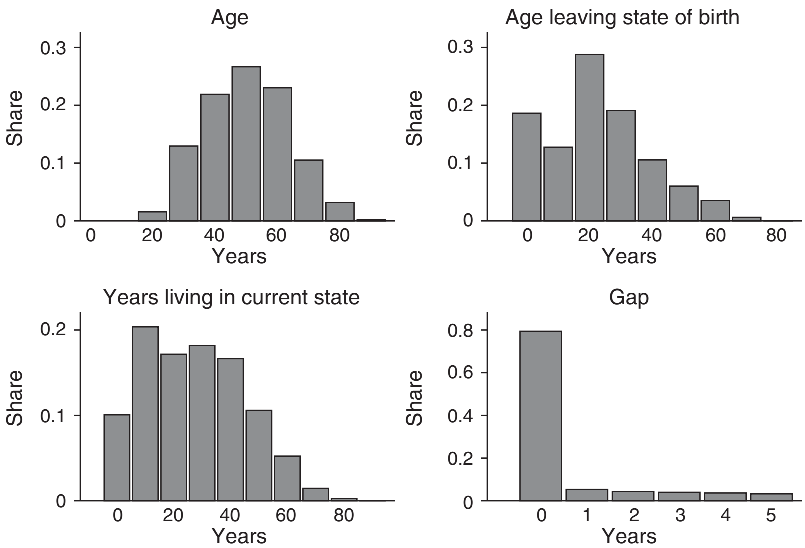

Figure 1. Summary Statistics for Migrants

Table 1 summarizes the migration patterns in our final sample. Approximately 16 percent of respondents are born in a different census region than the one in which they currently live. The most common moves have been out of the Northeast and Midwest and into the South and West regions of the United States.

Figure 1 shows the distribution of age of respondents in our final sample, along with the distributions of the age at which respondents moved out of their state of birth, the number of years respondents have lived in their current state of residence, and the gap between the year when they moved out of their state of birth and the year when they moved into their current state of residence. The figure shows that there is substantial variation in all of these measures, and that the majority of sample households have no gap between leaving their state of birth and arriving in their state of residence.

For simplicity, we treat the small set of households with a gap greater than zero and less than five years as if the gap were zero. That is, we assume the age at which the shopper left her birth state was her current age minus the number of years she reports living in her current state. We show in Appendix B that our results are robust to dropping these households.

II. Descriptive Evidence

A. *Measurement Approach*

Index consumers by i, modules by j, and states by s. We focus on the top two brands in each category as defined above. Let i's observed *purchase share* in category j, $\hat{y}_{ij}$, be the number of purchases of brand 1 in category j divided by the total purchases of brands 1 and 2. Let $\hat{\mu}_{sj}$ be the mean of $\hat{y}_{ij}$ across all nonmigrant households in state s.

For each migrant consumer i, we define the *relative share* in category j to be i's purchase share, scaled relative to the average purchase share of nonmigrants in her current and birth states:

$$\beta_{ij} = \frac{\hat{y}_{ij} - \hat{\mu}_{sj}}{\hat{\mu}_{s'j} - \hat{\mu}_{sj}}, \tag{1}$$

where s' is i's current state, and s is i's birth state.

We take β_{ij} as a summary of the way migrants' purchases compare to those of nonmigrants. If purchases depend only on contemporaneous supply-side variables like prices, availability, and advertising, migrants should behave identically to nonmigrants in their current state, and β_{ij} should equal one on average. If purchases depend only on experiences early in life, migrants should behave identically to nonmigrants in their birth state, and β_{ij} should equal zero on average. If preferences evolve endogenously throughout the life cycle, β_{ij} should fall between zero and one, on average, and should depend on the age at which a migrant moved and the number of years she has lived in her current state.

To look at these patterns in the data, we parameterize $\hat{y}_{ij}$ as

$$\hat{y}_{ij} = f(a_i, t_i)\,\hat{\mu}_{s'j} + [1 - f(a_i, t_i)]\,\hat{\mu}_{sj} + \eta_{ij},$$

where a_i is the age at which i moved, t_i is the number of years i has lived in her current state, and η_{ij} is an i.i.d. error term. This in turn implies

$$\beta_{ij} = f(a_i, t_i) + \frac{\eta_{ij}}{\hat{\mu}_{s'j} - \hat{\mu}_{sj}}. \tag{2}$$

We estimate equation (2) by weighted least squares.[5] The exact form of $f(\ )$ varies depending on the specification.

B. *Cross-Section*

Table 2 summarizes variation in purchase shares. The average of the purchase share $\hat{y}_{ij}$ across all consumers and modules in our sample is 0.63. Conditional on purchasing at least one of the top two brands, consumers in the typical category make 3.0 purchases of the top brand and 1.7 purchases of the second-place brand. The cross-state standard deviation of the purchase share is 0.15. The absolute value of the gap

[5] In the online Appendix we show that the results are unchanged if we allow the variance of η_{ij} to depend on the number of purchases made by consumer i in module j.

Table 2—Summary Statistics for Final Sample

Number of categories	238
Number of households	
Nonmigrant	27,686
Migrant	10,412
Mean across categories:	
Average purchases of number 1 brand	3.0
Average purchases of number 2 brand	1.7
Average purchase share ($\hat{y}_{ij}$)	0.63
Cross-state standard deviation of average purchase share for nonmigrants	0.15
Average absolute difference between purchase share in birth and current state for migrants	0.11

Notes: Number 1 and number 2 brand in each module defined by total purchases. Purchase share $\hat{y}_{ij}$ is purchases of number 1 brand/(purchases of number 1 brand + purchases of number 2 brand). Cross-state standard deviation is computed by averaging $\hat{y}_{ij}$ within each state-module pair, taking the standard deviation across states within each module, and then taking the mean of this standard deviation across modules.

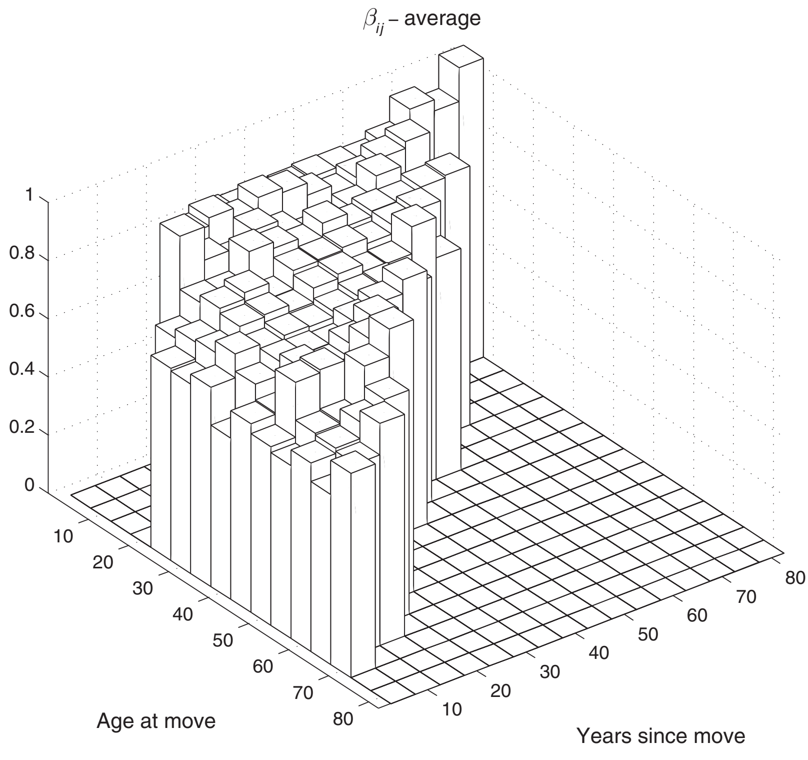

Figure 2. Relative Shares by Age at Move and Years Since Move

between the purchase share in a migrant's current state and in her birth state is 0.11 on average. These geographic differences are broadly consistent with the patterns reported in Bronnenberg, Dhar, and Dubé (2007). Table A3 reports the average purchase share and cross-state standard deviation for each module individually.

Figure 2 plots the key information in our data: how the relative share, β_{ij}, varies with a migrant's age at move (a_i) and years since move (t_i). We plot estimates of

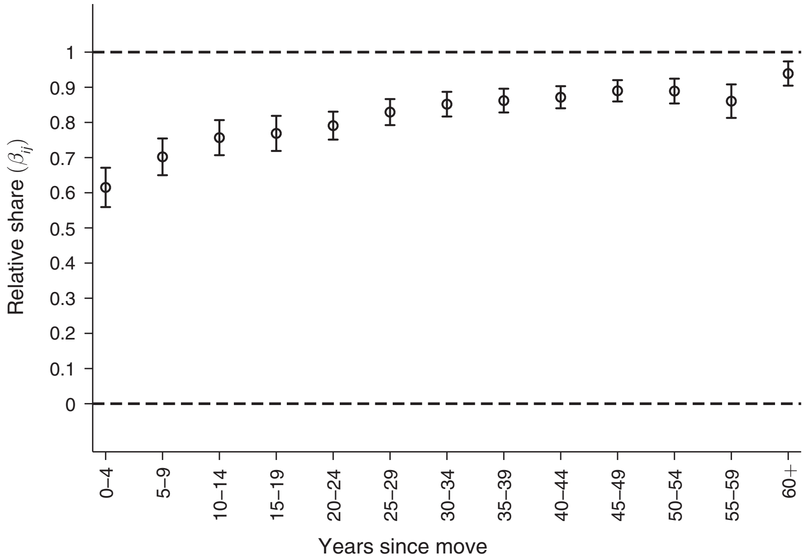

Figure 3. Relative Shares by Years since Move

Notes: Whiskers indicate 95 percent confidence intervals. Standard errors clustered by module.

equation (2), parameterizing $f(a_i, t_i)$ with dummies for each combination of a_i and t_i, pooled in five-year bins. The figure shows that β_{ij} is clearly less than one on average, rejecting the view that purchases are entirely driven by contemporaneous supply-side variables. It shows that β_{ij} is clearly greater than zero, rejecting the view that purchases are entirely driven by childhood experiences. The figure also suggests that the purchases of migrants converge gradually toward those of nonmigrants in their destination states.

To illustrate the patterns of convergence more clearly, Figures 3 and 4 show the same information as Figure 2 collapsed to two dimensions. Figure 3 shows variation with respect to years since move, pooling across the age-at-move categories. Notice, first, that even very recent movers have relative shares far from zero. This fact suggests that there is a discrete "on-impact" change in purchases at the time an individual moves, equal to approximately 60 percent of the gap between the two states. Referring back to Figure 2, we see that this jump is of similar magnitude regardless of the age at which a consumer moves. Second, note that migrant purchases converge slowly toward those of nonmigrants in the years following a move. It takes more than 20 years for half of the remaining gap in relative shares to close (reaching $\beta_{ij} = 0.8$), and even after 50 years the difference between migrants and nonmigrants remains statistically significant.[6]

Figure 4 shows variation with respect to age at move, pooling across the years-since-move categories. Migrants who moved during childhood have relative shares close to those of nonmigrants in their current states, while those who move later look closer to nonmigrants in their birth states. This pattern is consistent with the brand

[6] In the online Appendix we show versions of Figure 3 for subgroups defined by education, income, and gender.

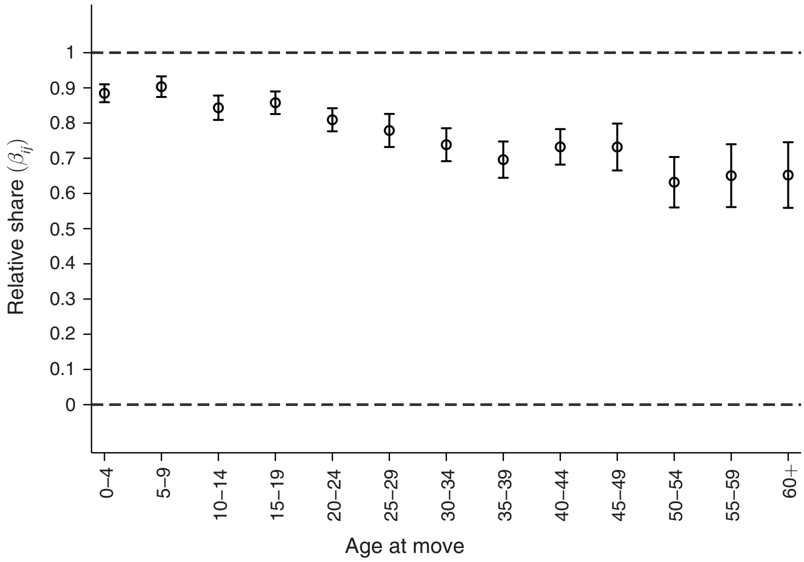

Figure 4. Relative Shares by Age at Move

Notes: Whiskers indicate 95 percent confidence intervals. Standard errors clustered by module.

capital model we introduce below, which predicts that the preferences of a consumer who has spent more time in her birth state will converge less quickly following a move. It is also consistent with results in marketing that show older consumers consider fewer brands when making a choice and are less likely to switch brands (Lambert-Pandraud and Laurent 2010; Drolet, Suppes, and Bodapati 2008). Interestingly, even consumers who moved before age 5 have relative shares slightly below 1, possibly reflecting the influence of parental preferences on childhood consumption.[7]

Note that the mechanical correlation between age at move and years since move means that Figures 3 and 4 partly repeat the same information. To separate the effect of age and years, Table 3 presents estimates of equation (2) where we include linear terms in a_i, t_i, and t_i squared. To make the coefficients easier to read, we divide both a_i and t_i by ten. For reference, the first column shows the regression analog of Figure 3 where we condition only on years since move. The constant in this regression gives the "on-impact" effect of moving, which we estimate to be 0.62. Relative shares start out converging at a rate of 10 percentage points per decade. The quadratic term is significantly negative, suggesting the rate of convergence slows over time. The second column adds age at move, a_i, which we find is significantly negative, showing that the preferences of older migrants indeed converge less quickly to those of their new state even after controlling for time since moving. The third and fourth columns control flexibly for time since move and age at move respectively. The linear and quadratic terms remain strongly significant and similar in magnitude in these regressions, confirming that time since move and age at move have independent effects.

[7] In the online Appendix, we present an extension of our model which allows parents to influence the consumption choices of young children. This extension rationalizes relative shares less than one for young movers but does not change any of our qualitative conclusions.

Table 3—The Evolution of Brand Preferences for Migrants

Dependent variable: Relative share (β_{ij})	(1)	(2)	(3)	(4)	(5)
Decades since move	0.098	0.079	0.075	—	0.092
	(0.009)	(0.009)	(0.010)	—	(0.016)
Decades since move squared	−0.009	−0.008	−0.007	—	−0.010
	(0.001)	(0.001)	(0.001)	—	(0.004)
Age (in decades) when moved	—	−0.018	—	−0.019	−0.013
	—	(0.005)	—	(0.005)	(0.008)
Constant	0.624	0.705	—	—	0.668
	(0.029)	(0.026)	—	—	(0.037)
Decades since move fixed effects	no	no	no	yes	no
Age when moved fixed effects	no	no	yes	no	no
Sample	all	all	all	all	age moved $\geq$ 25
Number of modules	238	238	238	238	238
Number of HH-module observations	528,621	528,621	528,621	528,621	212,957

Notes: The dependent variable β_{ij} is the share of a migrant's top-two brand purchases going to the top brand, scaled relative to nonmigrants in her current and birth states. $\beta_{ij} = 1$ implies her purchase share matches nonmigrants in her current state. $\beta_{ij} = 0$ implies her purchase share matches nonmigrants in her birth state. See Section II for details.

The final column repeats the regression of column 2 with the sample restricted to those moving at age 25 or later. We present this regression as a further test of the hypothesis that childhood experiences are decisive in shaping preferences. Both the jump on moving and convergence over time remain similar in magnitude and highly significant. So preferences do change, even for those who move late. This result provides some evidence against the common assertion that parental influence is dominant in shaping children's preferences (e.g., Moore, Wilkie, and Lutz 2002).[8]

C. *Panel*

Under assumptions we discuss in more detail in Section IV below, the cross-sectional variation in relative shares shown in Figure 2 is informative about how a given migrant's purchases evolve over time. In this section, we look at within-consumer variation in purchases more directly. The panel dimension of our data is limited, but we do observe a small number of consumers who move during the two years of our sample. For these consumers, we can follow purchases before and after their move and ask whether the panel lines up with our inferences from the cross-section.

Restricting attention to those for whom the gap between leaving their state of birth and arriving in their current state is zero, we observe 115 consumers who report moving in the past year and 111 consumers who report moving between one and two years ago. Given that our survey was fielded in September 2008, we expect the first

[8] Consumer behavior textbooks cite examples of parental influence. For instance, Berkman, Lindquist, and Sirgy (1997) state that "[i]f Tide laundry detergent is the family favorite, this preference is easily passed on to the next generation. The same can be said for brands of toothpaste, running shoes, golf clubs, preferred restaurants, and favorite stores" (p. 422–23).

Figure 5. Relative Shares by Month (Moved 10/07–9/08)

Notes: Whiskers indicate 95 percent confidence intervals. Standard errors clustered by module. The sample consists of migrants who report having lived in their current state less than one year. The dotted line at 0.62 indicates the relative share of recent migrants predicted from the cross-section.

group to have moved between October 2007 and September 2008, and the second group to have moved between October 2006 and September 2007.

Figure 5 shows relative shares by month for those who report moving in the past year. We plot estimates of equation (2), parameterizing f with dummies for t_i pooled in one-month bins. These consumers' relative shares for the months up to October 2007 are close to zero, indicating that their purchases before they move are similar to those of nonmigrants in their states of birth. If moves are distributed uniformly within the October 2007 to September 2008 period, and if an individual's relative share jumps to 0.62 on moving, we should expect the points to increase linearly from zero to 0.62 in the second half of the figure. This pattern is exactly what we observe.

Figure 6 shows relative shares by month for those who report moving between one and two years ago. As we would expect based on the cross-sectional evidence, relative shares increase roughly linearly from October 2006 to September 2007 and then are flat at 0.62 or slightly increasing thereafter.

III. Model and Estimation

As a lens through which to interpret these results, we introduce a simple model of consumer demand with habit formation (Becker and Murphy 1988). The model serves two purposes. First, it allows us to quantify the preference persistence we observe in terms of an economically meaningful structural parameter: the rate at which the stock of preference "capital" derived from past experience decays. Second, it lets us consider the implications of our results for firms' short-run and long-run demand curves, the importance of first-mover advantage, and the stability of market shares over time.

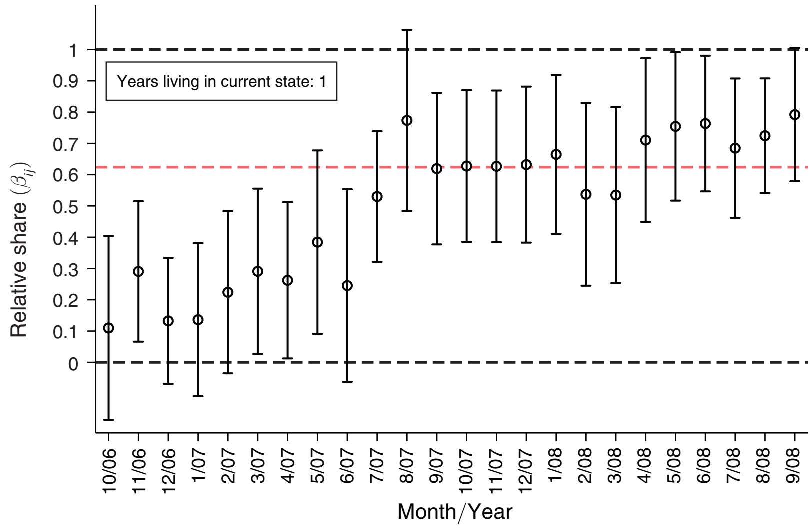

FIGURE 6. RELATIVE SHARES BY MONTH (MOVED 10/06–9/07)

Notes: Whiskers indicate 95 percent confidence intervals. Standard errors clustered by module. The sample consists of migrants who report having lived in their current state between one and two years. The dotted line at 0.62 indicates the relative share of recent migrants predicted from the cross-section.

A. *Setup*

We model a consumer deciding which of the top two brands to purchase in a particular module. We treat states as the relevant product market, assuming that supply-side characteristics of all brands are constant within state. We add subscripts for consumers, modules, and states when we turn to estimation in subsection C below.

The difference between the consumer's indirect utility from the top brand and the second brand is

$$U = \alpha\mu(\mathbf{X},\xi) + (1 - \alpha)k - \nu. \tag{3}$$

Here, $\mu(\mathbf{X},\xi) \in (0,1)$ is the consumer's *baseline utility*, $\mathbf{X}$ is an observed vector of consumer characteristics, ξ is an unobserved vector of product characteristics, $k \in [0,1]$ is the consumer's stock of *brand capital*, $\alpha \in (0,1]$ is a parameter governing the relative importance of past consumption in current preferences, and $\nu \sim \text{Uniform}(0,1)$ is a utility shock drawn independently across purchase occasions.

We assume the consumer prefers the top brand to the second brand if and only if $U \geq 0$. The probability that the consumer chooses the top brand (conditional on purchasing one of the top two) is therefore

$$y = \alpha\mu(\mathbf{X},\xi) + (1 - \alpha)k. \tag{4}$$

Equation (4) is a version of the standard linear probability model of demand (Heckman and Snyder 1997).

The baseline utility, $\mu(\mathbf{X}, \xi)$, captures the influence of all demand factors other than past consumption. $\mathbf{X}$ includes consumer characteristics such as cohort and income. ξ includes all relevant state-level characteristics of the top two brands, including their prices, availability, advertising levels, and qualities.

The stock of brand capital summarizes the consumer's past consumption experiences. We define the stock of brand capital to be the discounted average of past purchase shares:

$$k = \frac{\sum_{a=1}^{A-1} \delta^{A-a} \hat{y}_a}{\sum_{a=1}^{A-1} \delta^{A-a}}, \tag{5}$$

where $A \geq 1$ is the consumer's age and $\hat{y}_a$ is the consumer's actual purchase share across all purchase occasions at age a. The parameter $\delta \in [0,1]$ governs the persistence of capital over time.

We assume that equation (3) describes the consumer's purchases at all earlier ages. We also assume that α and $\mathbf{X}$ are constant, but that the capital stock, k, and the product characteristics, ξ, may have changed over time (for example, because the consumer moved from one state to another). When $A = 1$, and, thus, k is undefined, we assume $U = \mu(\mathbf{X}, \xi) - \nu$. We can thus think of $\mu(\mathbf{X}, \xi)$ as the expected utility of a consumer who has never before purchased either of the top brands in module j and so has acquired no brand capital.

It is straightforward to show that the linear recursive structure of equations (4) and (5) means we can write y as a weighted average of past $\mu(\mathbf{X}, \xi)$ plus a mean zero shock:

$$y_A = \sum_{a=1}^{A} w_a^A \mu(\mathbf{X}, \xi_a) + \varepsilon_A, \tag{6}$$

where ξ_a is the vector of product characteristics the consumer faced at age a, $E_\nu(\varepsilon_A) = 0$, $w_a \in [0,1]$, and $\sum_{a=1}^{A} w_a = 1$.

Consider, now, the special case in which product characteristics, ξ, vary across states but are constant over time. It is immediate that if the consumer has lived in the same state throughout her life, her expected purchase share is simply $y = \mu(\mathbf{X}, \xi) + \varepsilon$, where ξ are the product characteristics in her current state. Suppose instead that the consumer has moved exactly once: she lived in a state with characteristics ξ until age a^* and then moved to a state with characteristics ξ'. It is immediate from equation (6) that

$$y_A = \beta\mu(\mathbf{X}, \xi') + (1 - \beta)\mu(\mathbf{X}, \xi) + \varepsilon_A, \tag{7}$$

where $\beta = \sum_{a=a^*+1}^{A} w_a^A$ and, hence, $\beta \in (0,1)$.

It is straightforward to derive an explicit expression for β as a function of the age at which the consumer left her birth state (a^*) and the number of years she has lived in her current state ($t^* = A - a^*$):

$$\beta = 1 - (1 - \alpha)\left[\prod_{r=1}^{t^*-1}\left(1 - \frac{\alpha}{\sum_{\ell=0}^{a^*+r-1} \delta^\ell}\right)\right], \tag{8}$$

if $t^* > 1$, and $\beta = \alpha$ if $t^* = 1$. See Appendix A for the derivation of equation (8). Note that $\lim_{t^* \to \infty} \beta = 1$, and that β is increasing in t^*. Note also that β is decreasing in a^* for $t^* > 1$.

B. *Discussion*

The weight, β, in equation (7) is the model analog of the relative share defined in Section II: $\mu(\mathbf{X}, \xi)$ is the average purchase share among nonmigrants in a migrant's birth state, $\mu(\mathbf{X}, \xi')$ is the average purchase share among nonmigrants in her current state, and $\beta = \frac{y - \mu(\mathbf{X}, \xi)}{\mu(\mathbf{X}, \xi') - \mu(\mathbf{X}, \xi)}$.

The predictions of the model are consistent with the facts documented in Section II. A migrant's expected purchase share falls between the share among nonmigrants in her market of current residence and nonmigrants in her market of birth ($0 < \beta < 1$). When an individual moves, a fraction α of the market share gap between the two markets is closed immediately, as the product characteristics the consumer faces change from ξ to ξ' ($\beta = \alpha$ at $t^* = 1$). The parameter α therefore captures the "on-impact" effect of moving. The on-impact effect is the same regardless of the age at which the consumer moved. The remaining $1 - \alpha$ portion of the share gap closes gradually over time as her stock of brand capital adjusts. The adjustment is slower if δ is close to one, and if the consumer was older when she moved (since in this case she has accumulated a larger stock of past brand experiences).

The model is restrictive in several important ways. First, we model only the relative utilities of the top two brands. We do not model the extensive margin of whether or not to make a purchase in a module at all, and we suppress substitution with other brands.

Second, we assume that the capital stock, k, and the current demand characteristics, $\mu(\mathbf{X}, \xi)$, are separable in the indirect utility function. The influence of prices or advertising on indirect utility and, hence, on demand, will be the same regardless of a consumer's past experiences. The separability assumption delivers the prediction that the jump in relative share on moving (or "on-impact" effect) is the same regardless of the age at which a consumer moves. We make this assumption for tractability, and because it is consistent with the observed data, as seen in Figure 2.

Third, consumers in our model are myopic. We assume the consumer prefers the top brand to the second brand if and only if $U \geq 0$. A sophisticated, forward-looking consumer would take account of the way purchases today will affect her capital stock and, thus, her expected utility tomorrow. Demand would therefore depend not only on current product characteristics, but also on expected future product characteristics.

Finally, we assume that the capital stock is a weighted average of past consumption. As discussed above, past experiences could affect present demand through other channels. Past consumption might matter because of learning, and so enter current demand through beliefs rather than preferences. Past exposure to advertising or past observation of peers might matter independently of the level of past consumption. We see our evidence as potentially consistent with all of these stories, and our data do not allow us to distinguish them completely. We specialize to a habit model mainly because it is a simple way to capture the key facts. We consider evidence for advertising and peer effects in Section VI below.

C. *Estimation*

Index consumers by i, modules by j, and states by s as in Section II. Index years by t. For each consumer i, we observe a vector of purchase shares with typical element $\hat{y}_{ij}$, a vector of observables $\mathbf{X}_i$, and a vector $\mathbf{M}_i$ which encodes i's history of migration—her current and birth state, the age at which she moved (a_i^*), and the number of years she has lived in her current state (t_i^*). We use $\hat{\mathbf{y}}$, $\mathbf{X}$, and $\mathbf{M}$ to denote the matrices which pool these vectors across i.

We parameterize baseline demand $\mu()$ as

$$\mu(\mathbf{X}_i, \xi_{jst}) = \gamma_{jst} + \mathbf{X}_i \lambda_j, \tag{9}$$

where λ is a vector of parameters and γ_{jst} is shorthand for the value $\gamma(\xi_{jst})$ of a function mapping the vector of product characteristics ξ_{jst} to a scalar. The vector $\mathbf{X}_i$ includes log income, as well as dummies for cohort, Hispanic identity, race, educational attainment, and employment status.

Our first identifying assumption is that there are no unobserved consumer characteristics correlated with both purchases and the exogenous variables $\mathbf{M}_i$ and $\mathbf{X}_i$: $E(\hat{y}_{ij} - y_{ij} | \mathbf{X}, \mathbf{M}) = 0$.

Our second identifying assumption is that, conditional on observables, the expectation of baseline demand in a given module-state pair in a past period is equal to the expectation in the current period. Denoting the value of γ_{jst} in the current period by γ_{js}, we assume: $E(\gamma_{jst} - \gamma_{js} | \mathbf{X}, \mathbf{M}) = 0 \ \forall t$.

For a consumer born in state s and currently living in s', we then have

$$E(\hat{y}_{ij} | \mathbf{X}, \mathbf{M}) = \begin{cases} \gamma_{js} + \mathbf{X}_i \lambda_j & \text{if} \quad s = s' \\ \beta(a_i^*, t_i^*; \alpha, \delta)[\gamma_{js'} + \mathbf{X}_i \lambda_j] & \\ + [1 - \beta(a_i^*, t_i^*; \alpha, \delta)][\gamma_{js} + \mathbf{X}_i \lambda_j] & \text{if} \quad s \neq s', \end{cases} \tag{10}$$

where γ_{js} denotes the current value of γ_{jst}, and $\beta(a_i^*, t_i^*; \alpha, \delta)$ is given by equation (8). Note that we now allow ξ_{jst} to vary over time within a market. It is straightforward to show that $\beta(a_i^*, t_i^*; \alpha, \delta)$ is the same as in equation (8), where we assumed that ξ was constant over time within a market.

We estimate the parameters of this model using a two-step, nonlinear least squares estimator. In the first step, we estimate the parameters $\{\gamma_{js}\}_{\forall s}$ and λ_j for each module j by running an OLS regression of $\hat{y}_{ij}$ on $\mathbf{X}_i$ and a vector of state dummies using only the nonmigrant consumers (for whom $s = s'$). In the second step, we estimate the remaining parameters, α and δ, by minimizing $[\hat{y}_{ij} - E(\hat{y}_{ij} | \mathbf{X}, \mathbf{M})]^2$, holding $\{\gamma_{js}\}_{\forall j,s}$ and $\{\lambda_j\}_{\forall j}$ constant at their estimated first-step values.[9]

We compute bootstrap standard errors over 25 bootstrap samples, clustered by module. That is, we sample J modules with replacement for each bootstrap sample and include all households in each selected module. Our standard error estimates

[9] In Appendix B we show that the results are essentially unchanged if we allow the variance of η_{ij} to depend on the number of purchases made by consumer i in module j.

are therefore robust to within-module correlation induced by, for example, variation over time in γ_{jst} or household-module-level unobservables.

IV. Evidence on Identifying Assumptions

A. *No Selection on Unobservables*

Our first identifying assumption is that there are no unobserved consumer characteristics correlated with both purchase shares, $\hat{y}_{ij}$, and the observables, $\mathbf{M}_i$ and $\mathbf{X}_i$. Of particular concern is the possibility that migrants are selected to have unobserved brand preferences intermediate between the typical nonmigrant in their state of birth and their current state of residence. It could also be the case that migrants who stay in a state for many years after moving have characteristics more similar to lifetime residents of that state than migrants who stay for only a few years.

The first test of our identifying assumption is the within-consumer analysis presented in Figures 5 and 6 and discussed in Section II above. We see that the migrants look similar to nonmigrants in their birth states in the months before they move. The mean relative share pooling months 10/06 to 9/07 for migrants living in their current state less than a year is 0.093, the 95 percent confidence interval is $(-0.025, 0.211)$, and we fail to reject $\beta = 0$ at the 10 percent level ($p = 0.12$). The data are also consistent with a discrete jump in migrant purchases on moving. Moreover, purchase shares for these consumers prior to moving are not significantly related to the age at which they moved ($p = 0.37$), providing no support for the hypothesis that the correlation between relative shares and age at move or years since moving in Figure 2 is primarily driven by selection on unobservables.

As a second test of our identifying assumption, we consider a subsample of brands that were introduced relatively recently. Under the assumptions of our model, a migrant who moved before either of two brands was introduced should have an expected purchase share no different from nonmigrants in her current state of residence. If the identifying assumption was violated, where a consumer lived before the brands were introduced would be predictive of her characteristics, and so migrants who moved before a brand pair was introduced would look significantly different from nonmigrants.

To execute this test, we select pairs of brands that we have confirmed were introduced in 1955 or later. To maximize the power of the test, we do not restrict attention to top-two brands but include any pair of brands for which we could find information indicating that both were introduced in 1955 or later, and in which 500 or more nonmigrant households in our sample purchased at least one of the two. When there is more than one such pair in a single module, we select the pair with the most total purchases. Our final sample includes 52 brand pairs. We compute relative shares, β_{iw}, for each pair w as in equation (1) and estimate the regression

$$\beta_{iw} = (\omega_0 + \omega_1 t_i^*) I(t_i^* \leq T_w) + [\omega_2 + \omega_3 t_i^*] I(t_i^* > T_w) + \varepsilon_{iw}, \tag{11}$$

where T_w is the number of years at least one brand in pair w has been available, t_i^* is the number of years since i moved, and $I()$ is the indicator function. We weight observations by $(\hat{\mu}_{s'j} - \hat{\mu}_{sj})^2$ as in equation (2) above. Under our identifying assumption, we expect $\omega_1 > 0$, $\omega_2 = 1$, and $\omega_3 = 0$.

Table 4—Brand Pairs Introduced after 1954

Dependent variable: Relative share (β_{ij})	(1)	(2)	(3)
Moved after brand introduced:			
Decades since move (ω_1)	0.007	0.007	0.018
	(0.002)	(0.003)	(0.005)
Constant (ω_0)	0.657	0.701	0.693
	(0.055)	(0.075)	(0.090)
Moved before brand introduced:			
Decades since move (ω_3)	0.001	0.001	0.001
	(0.002)	(0.002)	(0.002)
Constant (ω_2)	0.854	0.852	0.880
	(0.100)	(0.101)	(0.101)
Only brand pairs introduced after	1954	1975	1985
Number of brand pairs	52	24	11
Number of HH-pair observations	86,805	43,083	22,088

Notes: The dependent variable β_{ij} is the share of a migrant's top-two brand purchases going to the top brand, scaled relative to nonmigrants in her current and birth states. $\beta_{ij} = 1$ implies her purchase share matches nonmigrants in her current state. $\beta_{ij} = 0$ implies her purchase share matches nonmigrants in her birth state. The sample includes purchases of brand pairs introduced in 1955 or later. The coefficients in the first two rows apply to migrants who moved after the first brand in the pair in question was introduced. The coefficients in the following two rows apply to migrants who moved before the first brand in the pair was introduced. See Section IVA for details.

Table 4 presents the results. Consistent with our assumption, the coefficient on decades since moving is highly significant for those moving after the pair in question was introduced ($\omega_1 > 0$), but insignificant for those moving before the pair was introduced ($\omega_3 \approx 0$). We also cannot reject that the average shares of migrants who moved before the pair was introduced have the same average shares as nonmigrants in their current state of residence ($\omega_2 \approx 1$). We do, however, reject the joint hypothesis that $\omega_2 = 1$ and $\omega_3 = 0$; this could be evidence of a small amount of selection, or of measurement error in brand introduction dates. The results are robust to focusing on the complete set of pairs introduced since 1955, pairs introduced after 1975, and pairs introduced after 1985.[10]

B. *Expected Past Shares Equal Present Shares*

Our second identifying assumption is that, conditional on observables, the expectation of baseline demand in a given module-state pair in any past year is equal to the expectation in the current year.

To test this assumption, we study the 27 modules for which we observe purchases of both current top-two brands in the historical CCA data. For each module-state pair, we compute the current purchase share in the Homescan data across both migrants and nonmigrants. We then compare this share to the analogous share in the CCA data for the years 1948–1968, computed as described in Section C above. Under our identifying assumption, we expect that the regression of past shares on current

[10] In the online Appendix, we present a placebo version of this exercise where we replace observed brand introduction dates with randomly generated dates. In the placebo, decades since moving is highly significant for those who moved before a pair's introduction, and the constant term for these households is significantly less than one.

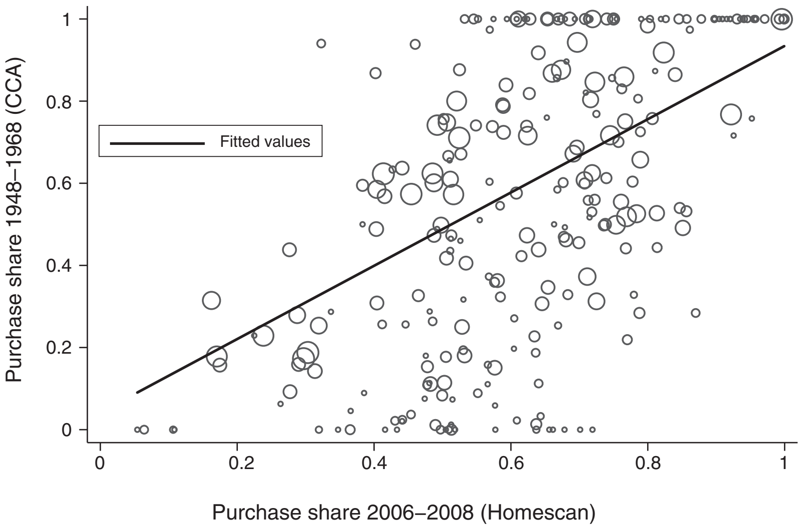

FIGURE 7. HISTORICAL AND CURRENT PURCHASE SHARES

Notes: Each observation is a state-module pair. The y-axis is average purchase share between 1948 and 1968, calculated using Consolidated Consumer Analysis. The x-axis is the average purchase share in the 2006–2008 Homescan sample. The size of the circles indicates the number of years of CCA data used to calculate the historical purchase share. See Section IVB for details.

shares should have an intercept of zero and a slope of one. Note that this is all that is required for consistency of our estimates; it is not necessary that past purchase shares *equal* current purchase shares, so long as they are the same in expectation.

Strictly speaking, a test of our identifying assumption requires that we compare past and current purchases of nonmigrants. We cannot perform this test, because the CCA data do not report shares by migration status. The regression of past on current shares will still be informative, however, so long as migrants are a relatively small share of the population and/or migration patterns have been relatively stable over time.

Figure 7 presents a scatterplot of current versus past purchase shares. Each observation is a state-module pair. The diameters of the circles are proportional to the number of years of CCA data we have for the observation. The current and past shares are clearly not equal, possibly reflecting real changes in market structure over time as well as sampling variability. However, the fitted values are very close to the 45-degree line.

Table 5 presents the corresponding regression of past shares on current shares, weighting by the number of years of CCA data and clustering by module. The estimated constant is 0.084, and the estimated slope is 0.822. We cannot reject the joint hypothesis that the constant equals zero and the slope equals one ($p = 0.30$).

A possible concern is that the coefficient in this regression may be attenuated by measurement error in the current shares. Consistent with this hypothesis, restricting the regression to state-module pairs where we observe at least 200 households making purchases in the Homescan data increases the estimated slope to 0.926 and reduces the estimated constant to 0.027. Restricting the sample to state-module pairs with at least 500 households increases the estimated slope to 1.039 and reduces the estimated constant to 0.001.

Table 5—Current and Historical Purchase Shares

Dependent variable: Purchase share 1948–1968	(1)	(2)	(3)
Current purchase share	0.822	0.926	1.039
	(0.119)	(0.105)	(0.089)
Constant	0.084	0.027	0.001
	(0.082)	(0.077)	(0.080)
Only include observations if number of Homescan HHs	≥ 0	≥ 200	≥ 500
p-value for (coeff. $=1$) & (cons. $=0$)	0.300	0.746	0.793
Number of modules	27	25	21
Number of state-module observations	325	188	115

Notes: Each observation is a state-module pair. The dependent variable is the estimated average purchase share in the state-module between 1948 and 1968, calculated using Consolidated Consumer Analysis. The right-hand side variable is the average purchase share in the 2006–2008 Homescan sample. All regressions weighted by the number of years of CCA data used to calculate the historical purchase share. The second column excludes observations where the number of observations used to compute the current purchase share is less than 200. The third column excludes observations where the number of observations is less than 500. See Section IVB for details.

Together, this evidence supports the assumption that the best predictor of a past purchase share given the data we observe is the present purchase share.

V. Results

A. *Parameter Estimates*

Table 6 presents estimates of the brand-stock model described by equations (4) and (5). The first parameter of interest is α, which represents the "on-impact" effect of moving to a different state. We estimate $\alpha = 0.626$, which is consistent with our descriptive analysis above and confirms that about 60 percent of the preference gap between territories is crossed on-impact when moving. Under the assumptions of our model, it also implies that 60 percent of the observed cross-state dispersion can be attributed to variation in supply-side factors ξ. The remainder, about 40 percent of regional share variation, can be attributed to consumers' stock of brand capital.

The estimate of the persistence parameter, δ, is 0.975. This magnitude is consistent with the earlier evidence that preferences appear highly persistent. The estimates suggest that it takes 27.2 years for half of a given year's contribution to the capital stock to decay.

In the online Appendix, we present fitted values and residuals from the model. The model successfully matches the qualitative features of the data.

B. *Demand Dynamics*

To see what these estimates imply for long-run and short-run price responses, consider a hypothetical market in which the top two brands, A and B, have equal market shares $(\mu(\mathbf{X}, \xi) = 0.5)$. Assume that the market has the same age distribution as the one observed in our Homescan sample, and that the current capital stock is $k = 0.5$ for all consumers.

TABLE 6—STRUCTURAL PARAMETERS

α	0.626
	(0.025)
δ	0.975
	(0.006)
Mean of $\gamma_{js} + X_i \lambda_j$	0.636
	(0.013)
Half-life of brand capital (years)	27.2
fval ($\times e05$)	0.885

Note: Table reports two-stage NLLS extimates of model parameters as defined in Section VA.

Suppose, now, that brand A cuts its price to a level that increases baseline demand, $\mu(\mathbf{X}, \xi)$, from 0.5 to 0.6.[11] This change causes an immediate increase in brand A's purchase share from 0.5 to $0.6\alpha + 0.5(1 - \alpha) = 0.563$.

For a permanent price cut, the model implies that the purchase share will eventually rise to 0.6. These long-run payoffs will take many years to materialize, however. The dynamics of the purchase share following a permanent price cut will, by assumption, be the same as the dynamics of a migrant's share following a move, and so will have a path very similar to that shown in Figure 3.

Our model also implies that the price cut will have long-run effects even if it is temporary. Given the estimated parameters, however, these effects will typically be very small. If brand A reverts to its original price after one year, its purchase share falls from 0.563 to 0.502. The long-run effect of the price cut is, thus, 3.2 percent of the on-impact effect (although the slight increase will last for a long time). This observation may explain why studies of temporary changes in advertising intensity have generally failed to detect significant long-run effects beyond a horizon of a few months (Assmus, Farley, and Lehmann 1984; Bagwell 2007). It also suggests that the long-run preference formation we are studying here is a distinct phenomenon from the habit effects documented by Dubé, Hitsch, and Rossi (2010), where brief price cuts lasting days or weeks have large effects on subsequent purchase behavior.

C. *Early Entry and Catching up by the Later Entrant*

In this section, we consider the implications of our findings for first-mover advantage. We simulate a hypothetical market in which two brands, A and B, enter sequentially. Let equation (4) be stated in terms of relative demand for brand B, so that $y = 0$ corresponds to all consumers buying from A and $y = 1$ corresponds to all consumers buying from B. We assume that A and B are identical in every respect except their prices, so that at equal prices baseline demand would be $\mu(\mathbf{X}, \xi) = 0.5$. For a given head start by brand A, we ask how much and for how long brand B would have to discount its price to achieve parity in purchase shares.

For example, suppose that A has a head start of five years. During this period, $y = 0$ as all consumers buy brand A. The accumulated capital stock at the end of those five years is $k = 0$. Brand B then enters, and the two firms play a game

[11] In Appendix D, we show that for a typical category this would amount to a discount of approximately 18 percent.

that determines prices. Abstracting from the details of this game, we know that if prices are equal ($\mu = 0.5$), we will have $y < 0.5$, and y will converge toward 0.5 but never reach it. If B offers a lower price, so that $\mu > 0.5$, both y and k will reach 0.5 in some finite number of years. The larger is the price discount, the faster the convergence.

We estimate the relationship between baseline demand $\mu(\mathbf{X}, \xi)$ and price using store-level price and quantity data from IRI. The details of this exercise are given in Appendix D. The average sensitivity of baseline demand to relative prices is $\partial\mu/\left(\partial \log \frac{price_A}{price_B}\right) = -0.490$. From this, we estimate that baseline demand levels $\mu \in \{0.55, 0.60, 0.65, 0.70, 0.75\}$ correspond to price discounts of $1 - \frac{price_A}{price_B}$ $\in \{0.10, 0.18, 0.26, 0.34, 0.40\}$. For each of these price discounts, we consider head starts for brand A of $t \in \{1, 5, 10, 15, 25\}$ years.

Over sufficiently long horizons, it is important to account for the fact that some consumers will die (destroying some of A's capital) and others will be born (with much less of A's capital). We run the simulations assuming that the age distribution is stable over time and matches the empirical distribution we observe in our Homescan sample.

Table 7 shows the results. Equalizing shares in a reasonable amount of time requires significant investment. If A's head start is five years, B would need to discount its price by 18 percent to reach market share parity in just more than a decade. To catch up in only three years, B would need to discount its price by 34 percent. If A's head start is 15 years, B would require 23 years at 18 percent price discount, or seven years at 34 percent price discount, to reach market share parity.

D. *Persistence under Market Shocks*

Bronnenberg, Dhar, and Dubé (2009) show that regional share differences in consumer packaged goods industries persist over remarkably long periods of time. Current local shares are strongly predicted by who was the first entrant in a market, even when that entry happened a century ago, few consumers alive remember a time when both brands were not widely available, and the intervening years have seen large shocks to the economic environment such as the growth of supermarkets, changes in real income, wars, depression, and so on.

Our model does not predict how much persistence we should expect to see because it does not endogenize firm choices. The previous section showed that a second entrant would have to make large investments to catch up to the first entrant; it does not say anything about whether or not we will see those investments in equilibrium. In this section, we consider a specific assumption under which our model does have strong implications about persistence: complementarity between the stock of capital (k) and current investments in gaining market share (ξ).

Fix a hypothetical market with N consumers and focus on a single category. Assume that $\mu(\mathbf{X}, \xi) = \xi$ and interpret ξ as firm 1's share of "marketing expenditures." Extend the example of the previous section and suppose that purchase shares are subject to a shock κ_t each period:

$$(12) \qquad y_{it} = \alpha\, boldsymbolxi_t + (1 - \alpha) k_{it} + \kappa_t .$$

Table 7—First Mover Advantage

	Years to equate shares				
	Price discount by second entrant				
First entrant's head start (t)	10%	18%	26%	34%	40%
1 year	10	4	2	1	1
5 years	27	12	6	3	1
10 years	33	19	10	5	2
15 years	36	23	13	7	3
25 years	37	26	16	9	4
Baseline demand (μ) of second entrant implied by this price discount	0.55	0.60	0.65	0.70	0.75

Notes: An entry in the table is the number of years that a second entrant would need to maintain a certain price discount in order to achieve parity in brand shares. Rows indicate the assumed number of years that the first entrant was in the market alone. Columns indicate the size of the price discount. See Section C for details. The relationship between price discounts and baseline demand is estimated from aggregate IRI data on prices and quantities as described in Appendix D.

Assume that κ_t is distributed i.i.d. uniformly on $[-\overline{\kappa}, \overline{\kappa}]$. Assume that the age distribution is stable over time and matches the empirical distribution in the Homescan sample, as in Section C. Finally, assume that equilibrium marketing expenditures are proportional to expected market share, where the expectation is taken after κ_t is realized.[12] That is, marketing expenditures in period t satisfy $\xi_t = \frac{1}{N}\sum_i y_{it}$, which in turn implies

$$\xi_t = \frac{1}{N}\sum_i k_{it} + \frac{\kappa_t}{1 - \alpha}. \tag{13}$$

These assumptions are highly stylized. A proper treatment of either the impact of specific marketing investments on consumer behavior or the firms' optimal choice of these investments is beyond the scope of this paper. Nevertheless, this setup captures the intuition that a brand that has a lead in the capital stock of experienced consumers will receive more marketing investment and consequently be purchased more often by inexperienced consumers.

We assume an existing market share for the leading brand of 0.75, which has been in place for as long as consumers live. We fix $\alpha = 0.626$ (our empirical estimate) and simulate the evolution of market shares for different values of δ, from 0.975 (our empirical estimate) in steps of 0.25 down to 0.225. We assume that the parameter governing the shock process is $\overline{\kappa} = 0.05$, a number we choose because it is at the upper end of typical annual share movements in our data.[13] We then forward-simulate 100 years of evolution for our hypothetical market.

[12] According to Jones (1990) firms often set their advertising budget proportional to market share: "Most manufacturers use a case rate system […] which ties a brand's ad budget to its sales by allocating a certain number of advertising cents or dollars to each case sold" (p. 38). In the context of investments in shelf space, Bultez and Naert (1988) write that "Commercial models actually used by supermarkets determine space allocation following rules of proportionality to sales, revenue or profit" (p. 212).

[13] Under the allocation in equation (13), the shocks on market shares are uniformly distributed on $[-\overline{\kappa}/(1-\alpha), \overline{\kappa}/(1-\alpha)] \approx [-0.12, +0.12]$ at our estimated value for α. For comparison, focusing on modules with more than 5,000 purchases, we compute the year-to-year change in purchase shares for each market-category-year in our data. The average absolute change across these observations is 0.04, and 96 percent have changes lying in $[-0.12, +0.12]$.

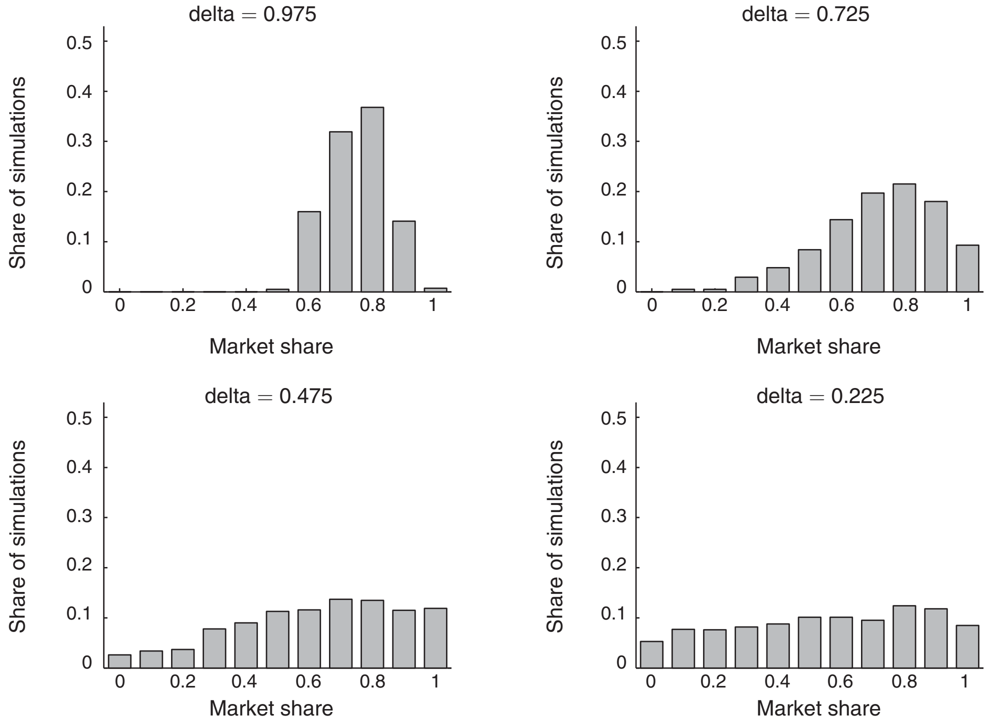

Figure 8. Persistence of Market Shares under Exogenous Shocks

Notes: Each panel shows the distribution of the long-run predicted market shares, y_t, after 100 years of market evolution. The initial stock of brand capital (and thus the initial market share) is initialized at 0.75. The four panels show outcomes of independent simulations using different values of the persistence parameter δ.

Figure 8 plots the distribution of the market shares in the final year of the simulation across 1,000 replications. The first panel shows that when we fix δ at its estimated value (0.975), long-run market shares remain closely concentrated around their initial value of 0.75, even after 100 years of shocks. The probabilities that market shares are within 10 or 20 share points of their initial value after 100 years are 69 percent and 99 percent, respectively. The mechanism generating the resistance to shocks is the persistent effect of past purchase shares y on the consumer's stock of brand capital stock. If δ is high, the shocks on past purchase shares y do not translate into shocks on the stock of brand capital because they cancel out over time. It is thus the persistence δ in the stock of brand capital that buffers against the reinforcement of demand and supply shocks. Accordingly, the resistance to market shocks weakens when we consider lower values for δ. The probability that market shares are within 10 share points of the initial values drops from 69 percent with $\delta = 0.975$, to 22 percent with $\delta = 0.225$, which is barely above the 20 percent one would expect if shares after 100 years attain a uniform distribution. As δ decreases towards 0, historical advantages are all but erased.

From this simple simulation, we conclude that our estimates of preference persistence, combined with complementarity between current investment and brand capital, can rationalize stable market shares over long periods of time even in the presence of large shocks.

VI. Mechanisms

A. *Brand Capital*

We estimate that 40 percent of current geographic variation in purchase shares is explained by variation in consumers' brand capital stocks. For tractability and ease of exposition, we have modeled brand capital formation in a habit framework, assuming the current capital stock is a function only of past consumption. As mentioned in the introduction, however, the brand capital stock may be partly a function of other variables, such as past exposure to advertising (Schmalensee 1983; Doraszelski and Markovich 2007), or past observations of consumption by peers (Ellison and Fudenberg 1995).

To provide a first look at the mechanism behind brand capital, we ask how our parameter estimates depend on whether a category has high or low levels of advertising. Recall that we define a category to have high advertising if total expenditure by the top two brands is greater than the 75th percentile among all categories in our dataset. We reestimate our main model allowing both the weight on brand capital $(1 - \alpha)$ and the rate of persistence in brand capital δ to differ by advertising intensity.

We also divide categories by the extent to which their consumption is socially visible. We code this measure subjectively. We judge products to be socially visible if (i) they are frequently consumed together with others in social situations, and (ii) they are frequently consumed or served directly from a package with the brand name visible. Products such as beer, soda, chips, ketchup, and cigarettes are therefore coded as socially visible. Products such as baby food, toothpaste, and cold remedies are not socially visible because they fail criterion (i). Products such as gravy mixes, frozen pasta, and shredded cheese are not socially visible because they fail criterion (ii). See Table A3 for the module-by-module coding.

As with advertising, we allow both $(1 - \alpha)$ and δ to differ by social visibility. Note that the correlation between the dummy for high advertising and the dummy for high visibility is low, so the sample splits by advertising and visibility should capture independent variation.

Table 8 presents the results. We find that advertising-intense categories have a significantly lower value of α and, thus, a significantly larger weight on the brand capital stock in utility. We cannot interpret this difference as causal, but it is consistent with the stock of past advertising exposure influencing current willingness to pay above and beyond the effect of past consumption. We find no significant differences in δ, consistent with the influence of past consumption and past advertising decaying at a similar rate.

We see a similar pattern with social visibility. We find that categories with a high degree of social visibility have a smaller estimated α, implying greater weight on brand capital. This finding is consistent with past observations of peer consumption exerting an independent influence on current willingness to pay. We again find no significant difference in δ.[14]

[14] In Appendix C, we show the existence of additional heterogeneity in our structural parameters across modules with different concentration ratios, purchase intensities, and degrees of spatial variation in preferences.

Table 8—Structural Parameters by Advertising Intensity and Social Visibility

	Advertising			Visibility		
	Low	High	Difference	Low	High	Difference
α	0.656	0.494	0.163	0.676	0.546	0.129
	(0.031)	(0.030)	(0.043)	(0.034)	(0.026)	(0.044)
δ	0.976	0.965	0.011	0.983	0.964	0.019
	(0.010)	(0.012)	(0.016)	(0.009)	(0.011)	(0.014)

Notes: The first through third columns report parameter estimates from a specification in which α and δ are allowed to differ for "low advertising" and "high advertising" categories. The fourth through sixth columns report parameter estimates from a specification in which α and δ are allowed to differ for "low social visibility" and "high social visibility" categories. See Section VI for details.

B. *Baseline Demand*

The remaining 60 percent of geographic variation in purchase shares is driven by differences in baseline demand $\mu(\mathbf{X}, \xi)$. Recall that the source of this result is the observation that when migrants move, their consumption shifts immediately toward the dominant brand in the destination market, closing 60 percent of the gap in purchase shares. Because **X** does not change with a move, our model would imply that this must be explained by differences in unobserved product characteristics, ξ. In other words, consumption jumps because migrants encounter some combination of lower prices, higher advertising, widespread distribution, or other advantages of the dominant brand when they arrive in their destination state.

We use the aggregate IRI data to get a feel for the role of specific supply-side variables. Details of this exercise are provided in Appendix E. First, for each category, we compute the share of cross-market variation in the log difference in purchase shares explained by the following independent variables: (i) log relative prices, (ii) relative display advertising intensity, (iii) relative feature advertising intensity, (iv) difference in the share of stores with nonzero sales of each brand ("availability"), and (v) difference in the average number of UPCs with nonzero sales for each brand ("UPCs"). We also compute the share of cross-market variation explained by these five supply-side variables together. Finally, we compute the mean of this share across categories.

We find that the cross-market correlation between log share differences and prices is -0.50 in the average category. The average share of cross-market variation explained by prices is 32 percent. We find that the average cross-market correlations between log share differences and feature and display advertising are 0.44 and 0.42 respectively, explaining 28 percent and 24 percent of cross-market variation. The average cross-market correlations of log share differences and availability and UPCs are 0.51 and 0.76 respectively, explaining 32 percent and 62 percent of cross-market variation on average. Together, prices, feature, display, availability, and UPCs explain 83 percent of cross-market variation on average.

These correlations suggest that a migrant is likely to find that the locally popular brand in her destination state has relatively lower prices, more advertising, and wider distribution. We expect these factors explain some portion of the jump in her consumption. We cannot tell how much, however, because all the supply-side variables are clearly endogenous, and the availability and UPCs measures are mechanically related to sales. The correlations and shares of variation explained have no causal interpretation.

To address this issue partially, we exploit the panel structure of our data to obtain an alternative estimate of the share of variance explained by prices and advertising. For each category, we run separate regressions of the log difference in purchase shares at the category-market-week level on market and week dummies, plus log relative prices, relative display intensity, and relative feature intensity respectively. We omit availability and UPCs because these variables' mechanical relationship to sales would make the panel estimates difficult to interpret. For each independent variable of interest, we compute the share of cross-market variance explained and then compute the average of this share across categories.

From these specifications, we estimate that variation in relative prices explains 8 percent of cross-market variation on average. Variation in relative feature intensity explains 0.9 percent. Variation in relative display intensity explains 3 percent.

A final possibility is that baseline demand depends on the observed consumption of others. This role for peer effects differs from the contribution to the brand capital stock discussed above. It would imply we might expect to see faster adjustment (higher α) for highly visible categories. As already discussed, Table 8 shows the opposite is true. This could mean that peer effects are not an important contributor to baseline demand, or that this effect is outweighed by their contribution to brand capital.

VII. Conclusions

Our results suggest that much of consumers' observed willingness to pay for brands may reflect the influence of past experiences. Heterogeneity in brand capital explains a substantial share of geographic variation in purchases. Brand capital evolves endogenously as a function of consumers' life histories and decays slowly once formed. Brand capital can explain large and long-lasting advantages to first movers. Brand preferences play an especially important role in categories with high levels of advertising and social visibility.

Appendix A: Derivation of Equation

We first write y_{A+1} recursively as a function of y_A. Define $\zeta_a = \hat{y}_a - y_a$. For any $A > a^*$, we can expand equation (4) as

$$y_A = \alpha\mu(\mathbf{X}, \xi') + (1 - \alpha)\frac{\sum_{a=1}^{A-1} \delta^{A-a}(y_a + \zeta_a)}{\sum_{a=1}^{A-1} \delta^a}. \tag{A1}$$

Combining equation (A1) with the analogous expression for y_{A+1} we can show that

$$y_{A+1} = \alpha\mu(\mathbf{X}, \xi')\frac{\delta}{\sum_{a=1}^{A} \delta^a} + \left(1 - \alpha\frac{\delta}{\sum_{a=1}^{A} \delta^a}\right)y_A + \frac{(1 - \alpha)}{\sum_{a=1}^{A} \delta^a}\delta\zeta_A. \tag{A2}$$

Next, we write $\beta(a^*, t^* + 1)$ as a function of $\beta(a^*, t^*)$. We know from equation (7) that for each a^* and t^* there exists $\beta(a^*, t^*)$ such that

$$y_A = \beta(a^*, t^*)\,\mu(\mathbf{X}, \xi') + (1 - \beta(a^*, t^*))\,\mu(\mathbf{X}, \xi) + \varepsilon_A.$$

Using this fact along with equation (A2), we can show that

$$\beta(a^*, t^* + 1) = \frac{\alpha}{\sum_{a=0}^{a^*+t^*} \delta^a} + \left(1 - \frac{\alpha}{\sum_{a=0}^{a^*+t^*} \delta^a}\right)\beta(a^*, t^*).$$

Starting from the fact that $\beta(a^*, 1) = \alpha$, it is then straightforward to show that

$$\beta(a^*, t^* + 1) = 1 - (1 - \alpha)\prod_{r=1}^{t^*}\left(1 - \frac{\alpha}{\sum_{\ell=0}^{a^*+r-1} \delta^\ell}\right).$$

Appendix B: Robustness Checks

Table A1 reports the results of several robustness checks. For each case, we first report our estimates for α and δ. We omit standard errors; in all cases they are comparable in magnitude to those reported in Table 6. We also report for each case a counterfactual: the number of years it takes for a second entrant with a baseline demand (μ) of 0.65 to catch up to a rival who had a ten-year head start. For our main specification, this is the value reported in the third column of the third row of Table 7 and is equal to ten years. Finally, we report the half-life of a year of brand capital.

The first row of the table repeats the estimates from our main specification reported in Tables 6 and 7. The following three rows report estimates in which we replace the number of purchases with alternative quantity measures in the definition of our dependent variable $\hat{y}_{ij}$, i.e., purchase volume normalized to account for differences in package size, dollars spent, and the number of discrete units bought. The fifth row reports estimates in which we weight observations to account for heteroskedasticity due to variation in the number of observed purchases across household-modules. The sixth row reports estimates in which we omit state-module pairs where one of the top two brands was not purchased. The final row reports estimates in which we drop households with a nonzero "gap" between leaving their birth state and arriving in their current state (see Sections VB and VD).

In all cases, our results are qualitatively similar.

Appendix C: Additional Evidence on Heterogeneity

Table A2 reports estimates from specifications in which we allow heterogeneous structural parameters by module. In each specification, we divide modules into two groups and allow each group to have separate α and δ parameters. The analysis thus follows the same format as the splits by advertising intensity and social visibility reported in Table 8. For each specification, we report the parameter estimates for each group, the number of years it takes for a second entrant with a baseline demand

Table A1—Robustness of Structural Parameters

Robustness check	α	δ	Years until convergence	Half-life of brand capital
Baseline	0.626	0.975	10	27
Alternative quantity variables:				
Equivalent units	0.627	0.975	10	27
Expenditure	0.626	0.975	10	27
Units	0.627	0.975	10	27
Error variance depends on number of purchases	0.625	0.971	10	23
Only module-states where both brands available	0.617	0.974	10	26
Only households with gap = 0	0.606	0.969	11	22

Notes: The table reports parameter estimates and counterfactuals for the robustness checks discussed in Appendix B. Equivalent units measures purchase volume normalized to account for differences in package size. Expenditure measures dollars spent. Units measures the number of discrete units bought. Error variance depends on number of purchases reports FGLS estimates where we weight observations to account for sampling variability in purchase shares. Denoting the share of the top brand in module j by $\bar{y}_j$ and the total number of purchases of the top two brands for household i by n_{ij}, we regress squared residuals from the unweighted model on a constant and a coefficient times $\bar{y}_j(1-\bar{y}_j)/n_{ij}$. We then weight observations by the inverse of the predicted values from this regression. Only module-states where both brands available excludes state-modules where one of the top two brands had zero sales. Only households with gap = 0 reports on estimates from a subsample of migrants with no gap between the year they report leaving their state of birth and the year they report moving to their current state.

Table A2—Additional Evidence on Heterogeneity

Sample	α	δ	Years until convergence	Half-life of brand capital
Joint share of top-two brands				
Below median	0.668	0.982	7	39
Above median	0.553	0.962	15	18
Purchase frequency				
Below median	0.580	0.975	14	27
Above median	0.641	0.974	9	27
Cross-state variation in shares				
Below median	0.480	0.971	25	24
Above median	0.667	0.972	6	25

Notes: The table reports results from models where we allow the structural parameters α and δ to differ for modules either above or below median along some dimension. See Appendix C for details. Joint share of top-two brands is the sum of purchases of the top two brands divided by the total number of purchases in the module. Purchase frequency is the number of purchases made in the module. Cross-state variation in shares is the standard deviation across markets within module of the market level mean of individual purchase shares.

(μ) of 0.65 to catch up to a rival who had a ten-year head start (the value reported for our main specification in the third column of the third row of Table 7), and the half-life of a year of brand capital.

In the first two rows, we split modules by the combined market share of the top two brands. For modules in which this share is low, we observe a higher value of α (i.e., less weight on brand capital) and a higher value of δ.

In the following two rows, we split modules by the number of purchases made per household. For modules that are less frequently bought, we observe a lower value of α (i.e., more weight on brand capital). The value of δ does not differ materially between the two groups.

Table A3—Modules, Top Two Brands, and Selected Module Characteristics

Module	Brand 1	Brand 2	Aggregate purchase share	Cross-state SD	Ad intense	Socially visible
Abrasive clnsr-liq	Soft Scrub	Comet	0.90	0.07	0	0
Abrasive clnsr-pwdr	Comet	Ajax	0.78	0.08	0	0
Adult incont. prod	Poise	Tena Serenity	0.68	0.15	0	0
Analgesic/chest rubs	Icy Hot	Vicks Vaporub	0.55	0.12	0	0
Antacids	Prilosec	Rolaids	0.71	0.08	1	0
Anti-gas products	Beano	Gas-X	0.52	0.13	0	0
Auto. dishwshr cmpnd	Cascade	Electrasol Jet-Dry	0.73	0.08	0	0
Baby food-strained	Gerber	Beechnut Stages	0.70	0.17	0	0
Bakery bagels	Thomas'	Sara Lee	0.74	0.29	0	0
Bakery bfast rolls	Little Debbie	Entenmann's	0.64	0.24	0	0
Bakery bread	Nature's Own	Sara Lee Soft and Smth	0.50	0.32	0	0
Bakery buns	Sara Lee	Wonder	0.61	0.32	0	0
Bakery cakes	Little Debbie	Hostess	0.91	0.07	0	0
Bakery cheesecake	The Father's Table	Cheesecake Factory	0.59	0.24	0	0
Bakery doughnuts	Hostess	Entenmann's	0.52	0.27	0	0
Bakery misc.	Homestyle	Flatout	0.51	0.26	0	0
Bakery pies	Little Debbie	JJ's	0.52	0.29	0	0
Bakery rolls	King's Hawaiian	Martin's	0.51	0.36	0	0
Baking cups and liners	Reynolds	Wilton	0.78	0.07	0	0
Bath additive-liq	Lander	Mr. Bubble	0.73	0.20	0	0
Beer	Budweiser	Miller High Life	0.64	0.19	1	1
Bouillon	Wyler's	Knorr	0.61	0.25	0	0
Breath sweetener	Tic Tac	Breath Savers	0.72	0.07	0	1
Butter	Land O Lakes	Challenge	0.86	0.27	0	0
Candy-choc minis	M&M Mars Snickers	Reese's Pnt Bttr Cup	0.51	0.07	0	1
Candy-chocolate	M&M Mars M&M Plain	Reese's Pnt Bttr Cup	0.52	0.06	1	1
Candy-diet. non choc	Life Savers	Baskin-Robbins	0.68	0.14	0	1
Candy-dietetic choc	Russell Stover	Whitman's Wgt Wtchrs	0.81	0.14	0	1
Candy-hard rolled	Pez	Smarties	0.52	0.11	0	1
Candy-lollipops	Tootsie Roll Pops	Spangler Dum Dum Pop	0.67	0.11	0	1
Candy-non choc minis	Tootsie Roll	M&M Mars Skittles	0.76	0.08	0	1
Candy-non chocolate	Y&S Twizzlers	Just Born	0.51	0.13	0	1
Candy-special choc	Hershey's Kisses	Russell Stover	0.54	0.07	0	1
Caramel corn	Crunch 'n Munch	Cracker Jack	0.71	0.09	0	1
Cat food-dry	Meow Mix	Purina Cat Chow	0.50	0.07	0	0
Catsup	Heinz	Hunt's	0.66	0.13	0	1
Cereal-dry	G M Cheerios	Post Hny Bnchs Oats	0.54	0.07	1	0
Cereal-granola	Sunbelt	Nature Valley	0.55	0.16	0	0
Cheese-amrcn cheddar	Kraft	Cracker Barrel	0.66	0.33	0	0
Cheese-amrcn colby	Kraft	Crystal Farms	0.81	0.23	0	0
Cheese-grated	Kraft	4C	0.92	0.06	0	0
Cheese-misc.	Kraft	Sargento	0.66	0.12	0	0
Cheese-mozzarella	Frigo Cheese Heads	Kraft Snkbls Polly-O	0.68	0.18	0	0
Cheese-muenster	Sargento	Finlandia	0.79	0.22	0	0
Cheese-shredded	Kraft	Sargento	0.72	0.15	0	0
Cheese-specialty	Athenos	Sargento	0.52	0.16	0	0
Cheese-swiss	Sargento	Kraft Deli Deluxe	0.61	0.15	0	0
Cigarettes	Marlboro	Doral	0.83	0.11	0	1
Cleaner-bathroom	Scrubbing Bubbles	Arm and Hammer Cn Shwr	0.81	0.05	1	0
Cleaner-disinfectant	Clorox	Lysol	0.52	0.11	1	0
Cleaner-metal	Jet-Dry	Dishwasher Magic	0.58	0.21	0	0
Cleaner-non disnfct	Pine-Sol	Mr. Clean	0.51	0.16	0	0
Cleaner-window	Windex	Sprayway	0.92	0.07	0	0
Coffee and tea filters	Melitta	Brew Rite	0.60	0.13	0	0
Coffee-grnd/bean	Maxwell House	Folgers	0.50	0.17	1	0
Coffee-soluble	Folgers	Nescafe Taster's Chc	0.54	0.11	0	0
Coffee-soluble flv	General Foods Int'l	Hills Bros	0.74	0.15	0	0
Cola-diet	Diet Coca-Cola	Diet Pepsi	0.55	0.11	1	1
Cola-regular	Coca-Cola Classic	Pepsi	0.52	0.12	1	1
Cold remedies-adult	Benadryl	Vicks Nyquil	0.55	0.10	1	0
Cold remedies-child	Tylenol Plus	Benadryl	0.50	0.21	0	0
Conditioner	Pantene Pro-V	Suave Naturals	0.51	0.09	1	0
Contact lens soln	Alcon Opti-Free Rpl	B&L Renu Multiplus	0.62	0.17	0	0
Cookies	Little Debbie	Nabisco Oreo	0.55	0.12	1	1
Corn chips	Fritos	Wise Dipsy Doodles	0.99	0.02	0	1
Corn dogs	State Fair	Foster Farms	0.69	0.21	0	0
Cough drops	Halls	Ricola	0.91	0.08	0	0

(*Continued*)

TABLE A3—MODULES, TOP TWO BRANDS, AND SELECTED MODULE CHARACTERISTICS (*Continued*)

Module	Brand 1	Brand 2	Aggregate purchase share	Cross-state SD	Ad intense	Socially visible
Cough syrups/tablets	Mucinex DM	Delsym	0.66	0.13	1	0
Crackers-butter	Nabisco Ritz	Keebler Townhouse	0.76	0.09	1	0
Crackers-cheese	Sunshine Cheez-It	Pepprdge Fm Goldfish	0.65	0.06	1	0
Crackers-flake	Keebler Club	Lance	0.97	0.04	0	0
Crackers-oyster	Nabisco	Dandy Vista	0.76	0.15	0	0
Crackers-sandwich	Austin	Lance	0.55	0.22	0	1
Crackers-soda	Nabisco Premium	Keebler Zesta	0.76	0.19	0	0
Dental floss	J&J Reach	Crest Glide	0.59	0.08	0	0
Denture cleanser	Polident	Efferdent	0.56	0.15	0	0
Deodorant-misc.	Secret	Mitchum	0.58	0.15	0	0
Deodorant-solid	Degree	Secret	0.54	0.07	1	0
Depilatories-women's	Nair	Sally Hansen	0.54	0.16	0	0
Detergent-heavy duty	Tide	Purex	0.56	0.05	1	0
Detergent-light duty	Dawn	Palmolive	0.56	0.09	1	0
Detergent-packaged	Tide	Gain	0.60	0.21	0	0
Dip mix	Hidden Valley Ranch	Concord Foods	0.78	0.11	0	0
Dip-canned	Frito-Lay	Tostitos	0.65	0.14	0	1
Dishwshr rinsing aid	Jet-Dry	Cascade Crystal Clr	0.80	0.07	0	0
Disinfectants	Lysol	Clorox	0.80	0.06	0	0
Disposable cups	Dixie	Dart	0.77	0.16	0	0
Disposable diapers	Pampers	Huggies	0.53	0.13	1	0
Disposable dishes	Dixie	Hefty	0.68	0.16	0	0
Dog and cat treats	Whiskas Temptations	Milk-Bone	0.60	0.12	0	0
Dog food-dry	Purina Beneful	Iams	0.52	0.14	1	0
Dog food-wet	Purina	Alpo Pedigree	0.53	0.13	0	0
Eye drops and lotions	Visine	Alcon Systane	0.53	0.16	0	0
Facial tissue	Kleenex	Puffs	0.63	0.07	1	0
Floor care cleaner	Swiffer Wet Jet	Clorox Ready Mop	0.87	0.11	1	0
Foot cmfrt products	Gold Bond	Dr Scholl's	0.63	0.18	0	0
Foot prepn-athlts ft	Lamisil AT	Tinactin	0.54	0.21	0	0
Foot prepn-misc.	Dr Scholl's	Pro Foot	0.85	0.07	0	0
Frozen dinners	Banquet	Healthy Chc Cmpt Slc	0.67	0.09	1	0
Frozen pot pies	Banquet	Marie Callender's	0.52	0.11	0	0
Frozen snacks	Totino's	Superpretzel	0.76	0.12	0	0
Fruit drinks-misc.	Minute Maid	Tropicana	0.65	0.18	0	1
Fruit juice-misc.	Dole	Tropicana	0.78	0.14	0	1
Fruit juice-orange	Tropicana	Minute Maid	0.67	0.16	0	1
Fruit spread	Smucker's Simply Frt	Polaner	0.52	0.27	0	1
Frzn Asian entrées-1	Weight Watchers	Tai Pei	0.59	0.17	0	0
Frzn Asian entrées-2	Lean Csn Cafe Clsscs	Banquet	0.56	0.13	0	0
Frzn Italn entrées-1	Weight Watchers	Bertolli	0.63	0.08	1	0
Frzn Italn entrées-2	Weight Watchers	Healthy Chc Simp Slc	0.51	0.17	1	0
Frzn meat entrées-1	Banquet	On-Cor	0.56	0.22	0	0
Frzn meat entrees-2	Lean Csn Cafe Clsscs	Boston Market	0.51	0.13	0	0
Frzn Mexcn entrées-1	El Monterey	Jose Ole	0.67	0.16	0	0
Frzn Mexcn entrées-2	Weight Watchers	Banquet	0.60	0.18	0	0
Frzn misc. entrées-1	Stouffer's	Mrs. T's	0.57	0.18	1	0
Frzn pltry entrées-1	Tyson	Banquet	0.68	0.10	0	0
Frzn pltry entrées-2	Weight Watchers	Boston Market	0.62	0.15	0	0
Frzn seafd entrées-1	Gorton's	Weight Watchers	0.64	0.16	0	0
Gelatin salad-refrig	Jell-O Ref	Winky Ref	0.89	0.09	0	0
Gravy mix	McCormick	Pioneer	0.74	0.15	0	0
Gravy-canned	Heinz Homestyle	Campbell's	0.56	0.12	0	0
Gum-bubble	Dubble Bubble	Adams Bubblicious	0.73	0.11	0	1
Hair color-women's	Clairol Nice 'n Easy	Revlon Colorsilk	0.55	0.08	1	0
Hair prepn-women's	Sunsilk	Pantene Pro-V	0.54	0.18	0	0
Hair spray-women's	Suave	White Rain	0.55	0.10	0	0
Hand sanitizer	Germ-X	Purell	0.52	0.13	0	0
Health bars/sticks	Zone Perfect	Clif	0.52	0.20	0	1
Hominy grits	Quaker	Jim Dandy	0.88	0.11	0	0
Honey	Sue Bee	Golden Nectar	0.68	0.24	0	1
Horseradish	Silver Spring	Gold's	0.59	0.38	0	0
Ice cream cones	Joy	Keebler	0.53	0.13	0	1
Ice cream-bulk	Breyers	Dreyer/Edy's Slw Chn	0.64	0.12	0	1
Ice milk and sherbet	Dreyer's/Edy's	Blue Bell	0.66	0.36	0	1
Insoles	Dr Scholl's	Pro Foot	0.77	0.10	1	0
Jam	Smucker's	Welch's	0.76	0.10	0	1

(*Continued*)

Table A3—Modules, Top Two Brands, and Selected Module Characteristics (*Continued*)

Module	Brand 1	Brand 2	Aggregate purchase share	Cross-state SD	Ad intense	Socially visible
Jelly	Welch's	Smucker's	0.59	0.12	0	1
Laxatives	Metamucil	Benefiber	0.56	0.18	1	0
Lemon/lime-diet	Sprite Zero	Diet Seven Up	0.51	0.16	0	1
Lemon/lime-regular	Sprite	Seven Up	0.66	0.14	1	1
Light beer	Bud Light	Miller Lite	0.56	0.17	1	1
Lighters	Bic	Scripto	0.78	0.07	0	1
Lip remedies-misc.	Carmex	Blistex	0.70	0.16	0	0
Lip remedies-solid	Chap Stick	Blistex	0.76	0.05	0	0
Lunches-refrig	Osc Mayer Lunchables	Armour Lunch Makers	0.85	0.09	1	0
Margarine and spreads	Shedd's	Blue Bonnet	0.51	0.12	0	0
Marshmallows	Kraft Jet Puffed	Campfire	0.94	0.05	0	1
Mayonnaise	Hellmann's	Kraft	0.55	0.25	0	1
Meat snacks	Jack Link's	Slim Jim	0.55	0.15	0	1
Medical accsry-misc.	Ezy-Dose	Apex	0.53	0.17	0	0
Medical wrap/brace	Mueller Sport Care	Ace	0.66	0.14	0	0
Minerals	Nature Made	Caltrate 600 + D	0.60	0.13	0	0
Misc. carb. bev-diet	Diet Dr Pepper	Diet Mountain Dew	0.51	0.13	1	1
Misc. carb. bev-reg	Mountain Dew	Dr Pepper	0.53	0.14	1	1
Mustard	French's	Gulden's	0.86	0.10	0	1
Nasal product	Afrin	Zicam	0.62	0.13	0	0
Nutritional supplmt	Nature Made	Rexall	0.56	0.16	0	0
Oral rnse/antiseptic	Listerine	Crest Pro-Health	0.74	0.05	1	0
Pain remedy-chld	Liq Children's Motrin	Chldrn's Tylenol Liq	0.63	0.12	0	0
Pain remedy-headache	Tylenol	Aleve	0.52	0.07	1	0
Paper napkins	Mardi Gras	Vanity Fair	0.57	0.11	0	0
Paper towels	Bounty	Kleenex Viva	0.70	0.06	1	0
Pasta-frzn	Rosetto	Celentano	0.60	0.36	0	0
Pasta-refrig	Buitoni	Monterey Pasta Co.	0.85	0.10	0	0
Peanut butter	Jif	Skippy	0.64	0.19	0	1
Pet care-bird food	Pennington	Morning Song	0.58	0.31	0	0
Pet care-pet food	Wardley	Kaytee	0.58	0.12	0	0
Petroleum jelly	Vaseline	Personal Care	0.72	0.14	0	0
Pizza-frozen	DiGiorno	Red Baron	0.51	0.07	1	0
Pizza-refrig	Mama Rosa's	Uno	0.85	0.20	0	0
Popcorn-popped	Smartfood	O-Ke-Doke	0.65	0.35	0	1
Popcorn-unpopped	Orville Rdnbacher's	Act II	0.63	0.09	0	0
Pork rinds	Baken-Ets	Mac's	0.74	0.23	0	1
Potato chips	Lay's	Pringles	0.54	0.07	0	1
Pre-moistened towels	Kleenex Cttnlle Frsh	Huggies	0.61	0.09	0	0
Precut salad mix	Fresh Express	Dole Fresh Favorites	0.70	0.21	0	0
Preserves	Smucker's	Polaner	0.89	0.09	0	1
Pretzels	Snyder's of Hanover	Rold Gold	0.60	0.17	0	1
Proc. cheese slices	Kraft Singles	Borden	0.67	0.22	0	0
Proc. cheese snacks	The Laughing Cow	Kraft Easy Cheese	0.57	0.15	0	0
Proc. cheese-amrcn	Kraft Singles	Borden	0.76	0.11	0	0
Razor-disposable	Schick Xtrme 3 Cmft+	Bic Comfort 3	0.54	0.16	0	0
Razor-nondisposable	Bic Soleil	Gillette Venus Embrc	0.52	0.21	1	0
Rectal medication	Preparation H	Tucks	0.77	0.13	0	0
Refrig entrees	Tyson	Perdue	0.57	0.30	0	0
Rug cleaner	Resolve	Bissell	0.53	0.14	0	0
Salad dressing mix	Hidden Valley Ranch	Good Seasons	0.50	0.21	0	0
Salad dressing-light	Kraft Free	Wish-Bone Sld Sprtzr	0.67	0.10	0	1
Salad dressing-liq	Kraft	Ken's Steak House	0.64	0.17	0	1
Salad dressing-refrg	Marie's	Marzetti	0.56	0.31	0	1
Salad toppings-dry	Hormel	Oscar Mayer	0.67	0.13	0	1
Salads-misc.	Reser's Ready Pac	Bistro Sld	0.63	0.24	0	0
Sandwiches-frzn/ref	Lean Pockets	Hot Pockets	0.52	0.07	0	0
Sauce mix-taco	Old El Paso	McCormick	0.54	0.21	0	0
Sauce-Asian	Kikkoman	La Choy	0.70	0.11	0	1
Sauce-barbecue	Kraft	Sweet Baby Ray's	0.61	0.17	0	1
Sauce-chili	Heinz	Tuong Ot Sriracha	0.81	0.18	0	0
Sauce-cocktail	Kraft	McCormick	0.64	0.26	0	1
Sauce-cooking	Hunt's Manwich	Del Monte	0.92	0.06	0	0
Sauce-dipping	Marzetti	Litehouse	0.81	0.28	0	1
Sauce-hot	Louisiana	Texas Pete	0.59	0.34	0	1
Sauce-marinara	Prego	Hunt's	0.52	0.08	0	0
Sauce-meat	A.1.	Heinz 57	0.80	0.15	0	0

(*Continued*)

TABLE A3—MODULES, TOP TWO BRANDS, AND SELECTED MODULE CHARACTERISTICS (*Continued*)

Module	Brand 1	Brand 2	Aggregate purchase share	Cross-state SD	Ad intense	Socially visible
Sauce-Mexican	Pace	Tostitos	0.53	0.19	1	1
Sauce-misc.	Prego	Kraft	0.59	0.20	0	1
Sauce-pepper	Tabasco	Frank's Redhot	0.57	0.19	0	1
Sauce-pizza	Ragu	Contadina	0.70	0.18	0	0
Sauce-worcestershire	Lea and Perrins	French's	0.69	0.15	0	1
Sauces and gravies	Buitoni	Garden Fresh Gourmet	0.61	0.26	0	0
Seasoning mix-chili	McCormick	Carroll Shelby's	0.84	0.12	0	0
Seasoning mix-misc.	McCormick	Sun Bird	0.54	0.13	0	0
Shampoo	Suave Naturals	Pantene Pro-V	0.53	0.07	1	0
Shave cream-men's	Edge Advanced	Barbasol	0.51	0.10	0	0
Shave cream-women's	Skintimate	Gillette Satin Care	0.65	0.07	0	0
Sinus remedies	Tylenol Sinus	Sudafed PE	0.66	0.14	0	0
Snacks-misc.	SunChips	GM Chex Mix	0.52	0.05	0	1
Snacks-variety pk	Frito-Lay	Wise	0.98	0.04	0	1
Soap-bar	Dove	Dial	0.53	0.09	0	0
Soap-liq	Softsoap	Dial	0.77	0.06	0	0
Soap-specialty	Suave Naturals	Dove	0.52	0.11	1	0
Soda straws	Forster	Glad	0.75	0.19	0	0
Soup mix-dry/bases	Maruchan	Lipton	0.61	0.11	0	0
Soup-canned	Campbell's	Progresso	0.80	0.06	1	0
Soup-frzn/refrig	Tabatchnick	Skyline	0.57	0.32	0	0
Throat lozenges	Ricola	Halls Breezers	0.64	0.12	0	0
Toast/breadsticks	Old London	Wasa	0.51	0.16	0	0
Toilet bowl cleaner	Lysol	Clorox	0.52	0.06	0	0
Toilet tissue	Charmin	Angel Soft	0.54	0.07	1	0
Toothbrushes	Colgate 360	Oral-B Indicator	0.55	0.11	1	0
Tortilla chips	Doritos	Tostitos	0.64	0.06	0	1
Trail mix	Planters	GM Chex Mix	0.79	0.13	0	1
Vinegar	Heinz	Pompeian	0.73	0.15	0	0
Vitamins-children	Flintstones	L'il Crttrs Gummy Vt	0.71	0.13	0	0
Vitamins-misc.	Nature Made	Nature's Bounty	0.71	0.13	0	0
Vitamins-multi	One A Day	Centrum Silver	0.60	0.08	1	0
Water-sparkling	Vintage	Perrier	0.62	0.30	0	1
Water-still	Glaceau Vitmn Water	Nestle Pure Life	0.52	0.13	1	1
Wave setting product	Garnier Fructis Styl	Pantene Pro-V Style	0.66	0.11	0	0
Yogurt-frozen	Turkey Hill	Wells Blue Bunny	0.57	0.37	0	1
Yogurt-refrig	Yoplait	Dannon	0.62	0.10	1	0

Notes: Brand 1 and brand 2 in each module defined by total purchases. Aggregate purchase share for a given module is total purchases of brand 1/(total purchases of brand 1 + total purchases of brand 2), and is calculated using all households in the Nielsen Homescan data. Cross-state standard deviation of the average purchase share for non-migrants is computed by averaging purchase share within each state-module pair, and then taking the mean of the standard deviation across states for each module. Cross-state standard deviation is calculated using the final sample as described in Section VD.

In the final two rows, we split modules by the cross-state standard deviation of average purchase shares. For modules in which the variation is low, we observe a lower value of α (i.e., more weight on brand capital). Again, the value of δ does not differ materially between the two groups.

APPENDIX D: ESTIMATION OF THE PRICE EFFECT ON BASELINE DEMAND

We use aggregate store-level data on 2001–2005 purchases and prices from the IRI Marketing Data Set (Bronnenberg, Kruger, and Mela 2008) to estimate the average effect of the relative prices of the top two brands in a typical consumer packaged goods category. These data cover sales in 30 consumer packaged goods categories for 260 weeks across 47 markets. We use total volume by brand-market-week as our measure of purchases. We compute prices by dividing expenditure for each brand-market-week by volume. We focus on the top two brands in each category by total volume across all markets and weeks. We index IRI markets by m and weeks by w. For the top two brands in category j, p_{1jmw} and p_{2jmw} are prices, $feat_{1jmw}$ and $feat_{2jmw}$ are the feature advertising intensity levels, $disp_{1jmw}$ and $disp_{2jmw}$ are the display advertising intensity levels, and y_{jmw} is the top brand's purchase share (as a fraction of all purchases of the top two brands). For each of the two promotional variables—display and feature—we define the intensity as the fraction of total volume sold in market m during week w under the given promotion. For our analysis, we use the relative promotional intensities: $\Delta feat_{jmw} = feat_{1jmw} - feat_{2jmw}$ and $\Delta disp_{jmw} = disp_{1jmw} - disp_{2jmw}$.

We extend equation (9) to model baseline demand as an explicit function of prices and advertising:

$$\text{(A3)} \qquad \mu(\mathbf{X}_i, \xi_{jmw}) = \gamma_0^{jm} + \gamma_1 \log\left(\frac{p_{1jmw}}{p_{2jmw}}\right) + \gamma_2^j \Delta feat_{jmw} + \gamma_3^j \Delta disp_{jmw} + \zeta_{jmw} + \mathbf{X}_i \lambda_j.$$

Assuming that the average capital stock k in the market is approximately constant over the period of this data, the expected purchase share is then

$$y_{jst} = \tilde{\gamma}_0^{jm} + \tilde{\gamma}_1 \log\left(\frac{p_{1jmw}}{p_{2jmw}}\right) + \tilde{\gamma}_2^j \Delta feat_{jmw} + \tilde{\gamma}_3^j \Delta disp_{jmw} + \tilde{\zeta}_{jmw}.$$

The constant $\tilde{\gamma}_0^{jm}$ is a category-market fixed effect that absorbs effect of the capital stock and the market average of $\mathbf{X}_i\lambda_j$. We assume the error term $\tilde{\zeta}_{jmt}$ is conditionally mean zero and cluster standard errors by category. Note that the price coefficient, $\tilde{\gamma}_1$, is just $\gamma_1\alpha$. Our estimate of the parameter $\tilde{\gamma}_1$ across the IRI categories is -0.307 with a standard error of 0.009. Using our previous estimate, $\alpha = 0.626$, this yields $\gamma_1 = -0.490$.

To apply these estimates to the counterfactual in Section C, consider a hypothetical category where brands 1 and 2 are identical in every respect except their prices, so that γ_0^{jm}, $\Delta feat_{jmw}$, $\Delta disp_{jmw}$, and the average of $\mathbf{X}_i\lambda_j$ are all zero. When $p_{1jmw} = p_{2jmw}$, market-level baseline demand μ is 0.5. Our estimates imply that baseline demand levels $\mu' \in \{0.55, 0.60, 0.65, 0.70, 0.75\}$ correspond to price discounts $1 - (p'_{1jmw}/p_{2jmw}) \in \{0.10, 0.18, 0.26, 0.34, 0.40\}$.

Appendix E: Estimation of Correlations between Shares and Marketing Variables Using IRI Data

We use the same IRI data as in Appendix D to assess the extent to which various supply-side variables drive the geographic variation in market shares.

Following the definitions in Appendix D we define the dependent variable of interest to be the log ratio of purchase shares $\log y_{jmt}/(1 - y_{jmt})$, and we define log relative prices, relative feature intensity, and relative display intensity to be $\log p_{1jmt}/p_{2jmt}$, $\Delta feat_{jmt}$ and $\Delta disp_{jmt}$ respectively. We define relative availability to be $\Delta avail_{jmt} = avail_{1jmt} - avail_{2jmt}$, where $avail_{kjmt}$ is the share of stores with nonzero sales of brand k in category j, market m, and week t. We define relative UPCs to be $\Delta UPC_{jmt} = UPC_{1jmt} - UPC_{2jmt}$, where UPC_{kjmt} is the average across stores of the number of distinct UPCs of brands k sold in category j, market m, and week t. In computing $avail_{kjmt}$ and UPC_{kjmt}, we weight stores by total volume.

We first collapse the data to the category-market level by taking means across weeks of each variable. We then estimate the raw cross-market correlation in each category between the log ratio of purchase shares and each supply-side variable. We also run a regression in each category of the log ratio of purchase shares on all five supply-side variables jointly and compute the R^2. We report the mean and standard deviation of the correlation and R^2 across categories.

We also run separate regressions by category of the log ratio of purchase shares on $\log p_{1jmt}/p_{2jmt}$, $\Delta feat_{jmt}$ and $\Delta disp_{jmt}$ along with market and week fixed effects. We collapse the data to the market level and compute predicted values by multiplying the market average of each independent variable of interest by its estimated coefficient. We estimate the share of variance explained by dividing the variance of the predicted value by the variance of the dependent variable. Finally, we compute the mean of the estimated shares across categories.

REFERENCES

Aaker, David A. 1991. *Managing Brand Equity*. New York: Free Press.

Allison, Ralph I., and Kenneth P. Uhl. 1964. "Influence of Beer Brand Identification on Taste Perception." *Journal of Marketing Research* 1 (3): 36–9.

Assmus, Gert, John U. Farley, and Donald R. Lehmann. 1984. "How Advertising Affects Sales: Meta-Analysis of Econometric Results." *Journal of Marketing Research* 21 (1): 65–74.

Atkin, David. 2010. "Trade, Tastes and Nutrition in India." Yale University, Department of Economics Working Paper No. 80.

Bagwell, Kyle. 2007. "The Economic Analysis of Advertising." In *Handbook of Industrial Organization Volume 3*, edited by M. Armstrong and R. Porter, 1701–1844. Amsterdam: Elsevier Science.

Bain, Joe S. 1950. "Workable Competition in Oligopoly: Theoretical Considerations and Some Empirical Evidence." *American Economic Review* 40 (2): 35–47.

Bain, Joe S. 1956. *Barriers to New Competition: Their Character and Consequences in Manufacturing Industries*. Cambridge, MA: Harvard University Press.

Becker, Gary S., and Kevin M. Murphy. 1988. "A Theory of Rational Addiction." *Journal of Political Economy* 96 (4): 675–700.

Berkman, Harold W., Jay D. Lindquist, and M. Joseph Sirgy. 1997. *Consumer Behavior: Concepts and Marketing Strategy*. Lincolnwood, IL: NTC Business Books.

Bowles, Samuel. 1998. "Endogenous Preferences: The Cultural Consequences of Markets and Other Economic Institutions." *Journal of Economic Literature* 36 (1): 75–111.

Bronnenberg, Bart J., Sanjay K. Dhar, and Jean-Pierre H. Dubé. 2007. "Consumer Packaged Goods in the United States: National Brands, Local Branding." *Journal of Marketing Research* 44 (1): 4–13.

Bronnenberg, Bart J., Sanjay K. Dhar, and Jean-Pierre H. Dubé. 2009. "Brand History, Geography, and the Persistence of Brand Shares." *Journal of Political Economy* 117 (1): 87–115.

Bronnenberg, Bart J., Jean-Pierre H. Dubé, and Matthew Gentzkow. 2012. "The Evolution of Brand Preferences: Evidence from Consumer Migration: Dataset." *American Economic Review*. http://dx.doi.org/10.1257/aer.102.6.2472.

Bronnenberg, Bart J., Michael W. Kruger, and Carl F. Mela. 2008. "The IRI Marketing Data Set." *Marketing Science* 27 (4): 745–48.

Bultez, Alain, and Philippe Naert. 1988. "SH.A.R.P.: Shelf Allocation for Retailers' Profit." *Marketing Science* 7 (3): 211–31.

Dekimpe, Marnik G., Jan-Benedict E.M. Steenkamp, Martin Mellens, and Piet Vanden Abeele. 1997. "Decline and Variability in Brand Loyalty." *International Journal of Research in Marketing* 14 (5): 405–20.

Doraszelski, Ulrich, and Sarit Markovich. 2007. "Advertising Dynamics and Competitive Advantage." *RAND Journal of Economics* 38 (3): 557–92.

Drolet, Aimee, Patrick Suppes, and Anand V. Bodapati. 2008. "Habits and Free Associations: Free Your Mind but Mind Your Habits." Unpublished.

Dubé, Jean-Pierre, Günter J. Hitsch, and Peter E. Rossi. 2010. "State Dependence and Alternative Explanations for Consumer Inertia." *RAND Journal of Economics* 41 (3): 417–45.

Einav, Liran, Ephraim Leibtag, and Aviv Nevo. 2010. "Recording Discrepancies in Nielsen Homescan Data: Are They Present and Do They Matter?" *Quantitative Marketing and Economics* 8 (2): 207–39.

Ellison, Glenn, and Drew Fudenberg. 1995. "Word-of-Mouth Communication and Social Learning." *Quarterly Journal of Economics* 110 (1): 93–125.

Erdem, Tülin. 1996. "A Dynamic Analysis of Market Structure Based on Panel Data." *Marketing Science* 15 (4): 359–78.

Foster, Lucia, John Haltiwanger, and Chad Syverson. 2010. "The Slow Growth of New Plants: Learning About Demand." Unpublished.

Heckman, James J., and James M. Snyder, Jr. 1997. "Linear Probability Models of the Demand for Attributes with an Empirical Application to Estimating the Preferences of Legislators." *RAND Journal of Economics* 28: S142–89.

Information Resources Incorporated. 2001–2005. "The IRI Academic Data Set." Chicago (accessed May 24, 2010).

Jones, John Philip. 1990. "Ad Spending: Maintaining Market Share." *Harvard Business Review* 68 (January–February): 38–42.

Keane, Michael P. 1997. "Modeling Heterogeneity and State Dependence in Consumer Choice Behavior." *Journal of Business and Economic Statistics* 15 (3): 310–27.

Keller, Kevin Lane. 1993. "Conceptualizing, Measuring, and Managing Customer-Based Brand Equity." *Journal of Marketing* 57 (1): 1–22.

Klemperer, Paul. 1987. "Markets with Consumer Switching Costs." *Quarterly Journal of Economics* 102 (2): 375–94.

Lambert-Pandraud, Raphaëlle, and Gilles Laurent. 2010. "Impact of Age on Brand Choice." In *The Aging Consumer: Perspectives from Psychology and Economics*, edited by Aimee Drolet, Norbert Schwarz, and Carolyn Yoon, 191–208. New York: Taylor and Francis, Routledge.

Ling, Davina C., Ernst R. Berndt, and Margaret K. Kyle. 2002. "Deregulating Direct-to-Consumer Marketing of Prescription Drugs: Effects on Prescription and Over-the-Counter Product Sales." *Journal of Law and Economics* 45 (2): 691–723.

Lodish, Leonard M., Magid Abraham, Stuart Kalmenson, Jeanne Livelsberger, Beth Lubetkin, Bruce Richardson, and Mary Ellen Stevens. 1995. "How T.V. Advertising Works: A Meta-Analysis of 389 Real World Split Cable T.V. Advertising Experiments." *Journal of Marketing Research* 32 (2): 125–39.

Logan, Trevon D., and Paul W. Rhode. 2010. "Moveable Feasts: A New Approach to Endogenizing Tastes." Unpublished.

Luttmer, Erzo F. P., and Monica Singhal. 2011. "Culture, Context, and the Taste for Redistribution." *American Economic Journal: Economic Policy* 3 (1): 157–79.

Moore, Elizabeth S., William L. Wilkie, and Richard J. Lutz. 2002. "Passing the Torch: Intergenerational Influences as a Source of Brand Equity." *Journal of Marketing* 66 (2): 17–37.

Nielsen Company. 2006–2008. "Homescan Disaggregated Data." Schaumberg (accessed November 11, 2008).

Nielsen Company. 2008. "Nielsen Panelviews Survey." Schaumberg (accessed October 1, 2008).

Pollak, Robert A. 1970. "Habit Formation and Dynamic Demand Functions." *Journal of Political Economy* 78 (4): 745–63.

Schmalensee, Richard. 1982. "Product Differentiation Advantages of Pioneering Brands." *American Economic Review* 72 (3): 349–65.

Schmalensee, Richard. 1983. "Advertising and Entry Deterrence: An Exploratory Model." *Journal of Political Economy* 91 (4): 636–53.

Smith, Michael D., and Erik Brynjolfsson. 2001. "Consumer Decision-Making at an Internet Shopbot: Brand Still Matters." *Journal of Industrial Economics* 49 (4): 541–58.

Thumin, Frederick J. 1962. "Identification of Cola Beverages." *Journal of Applied Psychology* 46 (5): 358–60.

TNS Media Intelligence. 2008. "TNS Ad$pender." New York (accessed June 9, 2010).

Williamson, O. E. 1963. "Selling Expense as a Barrier to Entry." *Quarterly Journal of Economics* 77 (1): 112–28.