The Formation and Persistence of Automobile Brand Preferences*

Soren T. Anderson
Michigan State University

Ryan Kellogg
University of Michigan

Ashley Langer
University of Michigan

James M. Sallee†
University of Chicago

March 30, 2012

Preliminary draft: Please do not cite without permission.

Abstract

Brand preferences in the automobile market demonstrate strong persistence, which implies that market conditions and policies that influence purchases today may reverberate far into the future. In this paper, we offer novel empirical evidence that (1) individuals tend to choose the same brand of automobile again and again and that (2) adult children tend to choose the same brands as their parents. Our results suggest that intergenerational correlations may arise both from information sharing within families and from endogenous preference formation in youth. These patterns of preference formation and persistence have implications for the long-term impacts of transportation policies and for automakers' pricing and product line strategies.

*We thank Evan Hernstadt and Katie Lim for excellent research assistance. Kellogg gratefully acknowledges financial support from the California Energy Commission.

†Anderson: Department of Economics and Department of Agricultural, Food, and Resource Economics, Michigan State University; email: sta@msu.edu; web: https://www.msu.edu/∼sta/. Kellogg: Department of Economics, University of Michigan; email: kelloggr@umich.edu; web: http://www-personal.umich.edu/∼kelloggr. Langer: Ford School of Public Policy and Department of Economics, University of Michigan; email: alanger@umich.edu; web: http://www.ashleylanger.com/. Sallee: The Harris School, University of Chicago; email: sallee@uchicago.edu; web: http://home.uchicago.edu/∼sallee.

1 Introduction

The automobile industry is seemingly obsessed with quantities. Press coverage focuses on market shares, sales volumes, and customer loyalty, often to the neglect of prices and profitability. Automakers are known to focus on "capture rates"—the fraction of new car buyers that previously purchased a different brand, and it is said that the Big Three will sometimes sell vehicles to rental companies at a loss to hit quarterly sales targets. The chairman of Toyota, Akio Toyoda, cited volume targets as the cause of Toyota's recent spate of safety recalls, telling the U.S. Congress in 2010 that, "Toyota has, for the past few years, been expanding its business rapidly. Quite frankly, I fear the pace at which we have grown may have been too quick . . . We were not able to stop, think, and make improvements as much as we were able to before."[1] At the time, Toyota had just recently overtaken GM as the global leader in sales, a title that GM had held for more than seven decades, although GM has since reclaimed that title.

While it is tempting to view a firm that maximizes sales instead of profits as misguided, strong brand loyalty may justify such a focus. If automobile buyers form strong attachments to the specific brands they own, then lowering prices to sell more vehicles today, perhaps even at the expense of today's profits, is an investment in brand loyalty that can pay dividends in the future. Our goal in this paper is to study the origins and persistence of consumers' automobile brand preferences and to explore the implications of these preferences for transportation policy and automakers' pricing and product offering decisions.

We study automobile brand preferences using information on vehicle ownership contained in the Panel Study of Income Dynamics (PSID), which follows individual households over time and continues to follow adult children even after they leave home to form their own households. This panel structure, which is unique among vehicle datasets in the United States, allows us to study both the persistence of an individual household's automobile brand choices as well as the correlation of brand choices across different households within the same family network—intergenerational correlations in particular.

We find that the persistence of brand choice within a household is quite strong: an individual that currently owns a particular brand is more than twice as likely as the average consumer to choose that same brand again when buying a new car. The intergenerational correlation in brand choice is similarly strong: an individual whose parent owns a particular brand is approximately twice as likely as the average consumer to purchase that same brand. The magnitude and robustness of this result across different specifications implies that intergenerational brand loyalty is an important feature of the car market. To the best

[1] "Toyota President Akio Toyoda apologizes for safety lapses," *Los Angeles Times*, February 24, 2010.

of our knowledge, we are the first to document this correlation.

We define brand preference, in its purest form, as a situation in which a consumer's experienced utility from owning a vehicle would change if the vehicle were produced under a different brand label, holding constant the vehicle's physical attributes. In practice, we also allow that brand preference could arise through a consumer growing accustomed to a brand's unique features over time (such as, for example, the contours of the seats or the positions of the radio dials), even though the consumer has no intrinsic preferences for such features ex-ante, making switching to another brand costly ex-post. In addition, brand preference could arise from uncertainty about the features or quality of different brands, leading to a preference for brands with which the consumer is more familiar. While it is challenging to distinguish between these different forms of brand preference, all three lead to correlation in brand choices over time, either directly or through dependence of the current brand choice on the choice made at the previous purchase.

Related to these different forms of brand preference, we posit that intergenerational transmission of brand preferences could be driven by three mechanisms: (1) the endogenous formation of preferences for particular brands themselves, perhaps as the result of nostalgia, (2) the endogenous formation of tastes for a brand's unique features through past experience, and (3) information that is shared over time within family networks. While it is difficult to assign particular empirical correlations to particular mechanisms, we find evidence suggesting that both endogenous taste formation and information sharing are at work. In support of the information mechanism, we find that parents' recent vehicle purchases influence the brand choices of their adult children. In support of endogenous taste formation, we show that there is a strong relationship between the brands owned by parents when their children were still living at home and the future brand choices of those children when they become adults and start their own households.

While intergenerational correlations in brand choice could arise from the transmission of these forms of brand preference from parent to child, they could also arise from correlations in household characteristics or geographic factors that influence the demand for vehicle attributes commonly offered by the brand. For example, both an adult child and her parents might have high incomes, leading both independently to choose European luxury brands. We address this concern by including an extensive set of geographic and demographic controls in our empirical specifications and ultimately by limiting our sample to pairs of automobile brands, such as Ford and GM, that offer a similar lineup of vehicle types. We continue to find a strong intergenerational correlation in brand choice in all of our empirical specifications. While we cannot rule out the possibility that these correlations are driven by unobserved household characteristics, the robustness of our results across a range of specifications argues

in favor a true brand preference.

After presenting our empirical evidence, we provide a heuristic model of the car market to show how the transmission of brand preferences from parent to child might influence the competitive behavior of automakers. In our model, consumers live for two periods, shopping for a car in a different market each period. For example, consumers might purchase an entry-level vehicle in the first period and then an upscale vehicle in the second. At the end of the second period, consumers have children, who may or may not inherit their parents' brand preferences. Firms compete in both markets. When individuals carry brand preference into the second period based on their first-period choice, but brand preferences are not transmitted across generations, the game collapses to a two-period model, as in Klemperer (1987), in which firms "invest" in brand loyalty in the first period and "harvest" the rents in the second period. In this case, markups are low on entry-level vehicles but high on upscale vehicles.

When children inherit the brand preference of their parents, however, then firms are unable to harvest loyalty by raising prices for upscale vehicles because harvesting today shrinks the market for the next generation. In this case, the game can be characterized as an infinite period model, similar to the one analyzed by Dubé, Hitsch and Rossi (2009). For the magnitudes of brand loyalty we estimate in the PSID data, we find that the presence of brand loyalty lowers prices and profits, since the incentive to lower prices to invest in future profits outweighs the incentive to raise prices to increase current profits. In essence, automakers are continuously engaged in a price battle to invest in future loyal customers but are never able harvest this loyalty.

Our model also shows that, despite the result that brand loyalty ultimately results in lower prices and profits in equilibrium, competitive pressures provide firms with a strong unilateral incentive to encourage brand loyalty amongst their customers. Intuitively, a firm can increase its equilibrium market share and profitability if it can take actions (such as, for example, undertaking an aggressive advertising campaign or installing a distinctive, firm-specific trim on all its models) that more tightly tie its current customers to its brand. These benefits, however, come at the expense of the firms' rivals, and when all firms behave in this way they collectively fare worse in equilibrium.

Of course, the persistence of brand loyalty within and across generations is certainly not the only explanation for why automakers develop vertical product offerings and focus extensively on sales quantities. The strength of the brand preferences that we document, however, suggests that this mechanism is qualitatively important and deserves further attention.

Brand loyalty may also be an important determinant of the impacts of transportation policies such as fuel economy standards or fuel taxes. To the extent that fuel economy is dif-

ferentiated across brands, the persistence of brand preferences within and across generations suggests that changes to fuel taxes (or, more generally, to fuel prices) may require many years to elapse before their effects are fully realized. In addition, enforcement of aggressive fuel economy standards may require substantial changes in vehicle prices or brands' fuel economy offerings in order to meet the targets. In future work, we plan to address these issues more directly by studying the persistence of consumers' preferences for vehicle attributes such as fuel economy.

Our analysis relates to several existing literatures. First, previous work in economics and marketing has studied brand loyalty in the automobile market (Train and Winston 2007; Mannering and Winston 1985 1991). While these papers document within-household brand loyalty, the automobile literature has, to the best of our knowledge, never analyzed the intergenerational dimension of brand preference that we document here.

Second, recent work on brand preferences in consumer packaged goods has demonstrated that brand loyalty is responsible for much of the observed persistence in brand choices for individual consumers and the observed persistence in market shares within specific geographic areas (Bronnenberg, Dhar and Dubé 2009; Bronnenberg, Dubé and Gentzkow Forthcoming). As compared to the products studied in that literature, automobiles are much larger expenses, they are purchased much less frequently, and the product offerings are much more heterogeneous. Brand loyalty in the automobile sector typically involves individuals purchasing quite different products that share a brand label, whereas the literature on packaged goods is better characterized as repeat purchases of the same item. For small purchases, brand loyalty may better be understood as a heuristic to aid in quick decision-making, which is likely quite different from the role brands play in purchasing an automobile.

Third, our work relates to an emerging literature on endogeneous preference formation that analyzes how household experiences during childhood influence adult behaviors. Much of this literature is focused on food, which is surveyed in Birch (1999), with recent economic contributions in Logan and Rhode (2010) and Atkin (2011). Luttmer and Singhal (Forthcoming) argue that preferences for redistribution are formed in the cultural context of childhood. Fernández, Fogli and Olivetti (2004) provide a related analysis of how the labor supply of mothers influences the labor supply of their sons' wives, which the authors explain as owing to preferences developed in childhood.

Fourth, there is of course a considerable literature on the intergenerational transmission of earnings and education, much of which utilizes the PSID, as well as datasets from other countries that link outcomes of parents and children across generations.[2] Closely re-

[2]See Solon (1992) for a seminal contribution and Solon (1999) and Black and Devereux (2011) for recent reviews.

lated literatures study the intergenerational transmission of IQ, occupations, welfare status, health, attitudes, social behavior, consumption, and wealth, nearly always finding a strong correlation between the outcomes of parents and their children (see Black and Devereux 2011).

Finally, our work has parallels in the extensive peer-effects and social interactions literatures (see Manski 1993 2000). Whereas much of this literature studies how individuals are influenced by the aggregate behavior and characteristics of a reference group, we focus on how parents and children are influenced by the behavior and choices of a small number of individual family members.

The results presented in this paper do establish that the automobile preferences of children are influenced by their parents' choices, but we are still pursuing two extensions that will significantly add to the analysis. First, we are exploring instrumental variables strategies that use changes in gasoline prices and the dates of openings and closings of automobile dealers to more credibly identify the causal effect of parental choices on child choices. Second, we are building a structural model of consumer vehicle choice, which jointly models the parent's and child's discrete choice problem, in order to simulate how brand choice responds to quality shocks and policy interventions. This modeling exercise will require that we make additional assumptions about the structure of preferences, but we believe the results are valuable in complementing our reduced-form estimates and allowing for internally consistent simulations.

The balance of the paper is structured as follows. We describe our data in 2. In section 3, we document the strength of brand persistence within households, cross-sectionally and over time. In section 4 we provide a framework for interpreting across household correlations in brand choices, which we document empirically in section 5. Section 6 lays out a simple theoretical framework which provides intuition for the implications of intergenerational brand preference for automakers' prices and profits in equilibrium. We expand on the work currently in progress in section 7. Section 8 concludes.

2 Data

Our data on vehicle ownership come from the Panel Study of Income Dynamics (PSID). In 1968, the PSID selected a nationally representative sample of households to survey, and since then it has asked them a battery of economic and demographic questions every year until 1997 and then every two years thereafter. The PSID collects information on everyone who lives in a PSID household, but it also follows members of the original PSID sample households and their children whenever they join or create a new household. As a result, the

survey now collects information on many households that are members of the same extended family.

The PSID began collecting information on vehicles in 1999. Respondents report the total number of vehicles that they own or lease and additional detailed information on up to three vehicles, including vehicle make, model, and vintage, as well as the date of purchase, purchase price, and whether the vehicle was a gift. These data are available from surveys conducted in 1999, 2001, 2003, 2005, and 2007. These data allow us to examine how vehicle choices by individual households evolve over time as well as how the vehicle choices of one household relate to the vehicle choices of other households linked by family relationships. To the best of our knowledge, the PSID is unique in providing such information for families in the United States.

Below, we focus on the relationship between an individual's brand choice and the brand holdings of his or her parents. To do this, we limit the PSID sample to individuals who own a vehicle and whose mother is also in the PSID and is also a vehicle owner. We link individuals to their mothers' households because we were able to match more individuals consistently with their mothers than with their fathers. Consistent with PSID terminology, we refer to males and single females as household heads and female spouses (including live-in significant others) as wives.

Table 1 shows sample means for several variables for the full sample of individuals with mothers in the PSID, for the individuals who remain after limiting the sample to those with vehicles, and for the mothers themselves. There are 6,798 heads or wives whose mothers are present in the data. Of these, 4,044 are vehicle owners and have mothers who own a vehicle. Compared to all household heads and wives with mothers present in the data, these households are slightly older, are slightly more educated, and have higher income. Because many of these individuals are siblings, there are only 2,213 unique mothers in the data. Mothers are obviously older, with a mean age of 59 compared 36 for their children. Heads and wives have larger households, have more kids living at home, and own slightly more vehicles than their mothers.

3 The strength of brand preferences within households

Brand preference, or brand loyalty, is a familiar term in marketing and sales. But what precisely does it mean in a microeconomic choice model?

We define brand preference for automobiles broadly to include several potential forms of preference. First, a consumer might experience a different level of utility from owning a vehicle produced under a different brand label, even holding fixed the physical attributes of

Table 1: Variable means in PSID

	All heads and wives with mothers	Final sample of heads and wives	Mothers in final sample
Age	33.6	35.8	59.3
Education	13.2	13.6	12.5
Annual family income	48,514	73,281	62,527
Number of kids	0.9	1.0	0.3
Household family size	2.7	3.0	2.3
Number of vehicles	1.5	2.1	2.0
Number of unique individuals	6,798	4,044	2,213
Number of individual-years	26,659	10,032	6,280

Table reports sample means. The first column reports means for all household heads and wives in the PSID whose mother also appears in the PSID. The next column focuses on the sub-sample of these households that own a vehicle and whose mother owns at least one vehicle. The final column focuses on the mothers of this sub-sample.

the vehicle itself. For example, an individual might have grown up in a household that drove Fords and therefore experience nostalgic pleasure when driving a Ford that would not be experienced when driving a physically identical GM. This form of brand preference would be analogous to a preference for locally or domestically produced goods—indeed, automobile brand preferences themselves may reflect biases either for or against imported cars.

Second, the consumer might have developed, through past experience, a taste for the unique features of the brand, such as the interior styling, the feel of the front seats, or the layout of the instrument panel. In addition, the consumer might have developed a good relationship with brand's local dealers and mechanics. In short, the consumer may simply have grown accustomed to the brand, and would therefore face a cost in switching. Implicit in this perspective is the assumption that there is nothing intrinsic to the consumer or special about these unique features of the brand, other than the consumer's previous experience itself, that leads to a difference in expected utility.

Finally, the consumer might possess, through past experience, better information about a particular brand. Thus, a risk-averse consumer might prefer to stick with the brand with which she is more familiar rather than switch to a brand for which her future utility is more uncertain. For example, an individual who has always driven Fords might feel confident that he or she knows the true reliability of Fords and dislike the greater uncertainty from buying a GM. Actual differences in reliability and performance are not brand preference, but subjective *beliefs* about reliability and variability are.[3]

[3]There is a large marketing literature on brand loyalty, which relates to our definition of brand preference. A seminal contribution is Jacoby and Chestnut (1978), which offers many definitions of brand loyalty that share with us a presumption that consumer choice is influenced directly by brand, but they differ from our

In this section, we study brand preferences by demonstrating the persistence of brand choices *within* households, looking both at the distribution of brand choices within a household's current fleet and at the correlation of the household's brand choices over time. In subsequent sections, we study the correlation of brand choices *across* households within family networks. This analysis allows us to test for brand preference spillovers and to pursue a different set of identification strategies that establish the existence of individual brand preference.

Mathematically, we can write the expected utility of individual i from choosing vehicle j as a function of the individual's characteristics, the vehicle's attributes, and brand preferences. Let vehicles be indexed by $j = 1, 2, ..., J$. Each vehicle has a brand B_j and a vector of characteristics $\mathbf{X_j}$ (of length K) that does not include brand. Brands are indexed by $b = 1, 2, ..., \mathcal{B}$. Each consumer has a vector of characteristics $\mathbf{D_i}$, which includes demographics such as age, income, and family structure, as well as location-specific characteristics $\mathbf{L_i}$, such as weather, terrain, and the presence of nearby automobile dealerships. In keeping with tradition in the discrete choice literature, we write the utility of individual i choosing vehicle j as an interaction between the vehicle's attributes and consumer's characteristics, along with a random error term:

$$E[U_{ij}] = \beta_{\mathbf{i}}\mathbf{X_j} + \sum_{b=1}^{\mathcal{B}}(\theta_{ib} \cdot 1(B_j = b)) + \varepsilon_{ij}, \tag{1}$$

where

$$\beta_{ik} = \alpha_k + \alpha_{\mathbf{dk}}\mathbf{D_i} + \alpha_{\mathbf{lk}}\mathbf{L_i} + \pi_{ik} \tag{2}$$

$$\theta_{ib} = \xi_b + \eta_{ib}. \tag{3}$$

and where $1(\cdot)$ denotes the indicator function, ε_{ij} is an idiosyncratic random utility error term, π_{ik} is a preference for characteristic k that is independent of demographics D_i and location L_i, and $\xi_b + \eta_{ib}$ denotes a preference for brand b that may have a component ξ_b shared by all consumers and an individual-specific component given by η_{ib}. Uncertainty is from the point of view of the individual consumer, whom we suppose is uncertain of the actual performance of a vehicle and has subjective beliefs regarding performance (which may be a function of brand). In this framework, individual i has a brand preference that differs from the average brand preference if $\eta_{ib} \neq 0$ for some b.

In the case of a pure brand preference, as in our first definition above, the interpretation of η_{ib} is straightforward: it is a preference shifter holding fixed *all* observed and unobserved

definition here in not beginning with a model of utility and in focusing on repeated purchases.

Table 2: Distribution of brand choices for two-vehicle households in the PSID

				Vehicle 2			
Vehicle 1	GM	Ford	Chrysler	Toyota	Honda	Other Asian	European
GM	**49**%	21	11	5	4	7	3
Ford	29	**39**	11	7	5	8	2
Chrysler	30	21	**23**	6	5	10	4
Toyota	20	17	9	**25**	9	14	6
Honda	20	16	7	16	**22**	12	6
Other Asian	24	17	12	9	9	**25**	5
European	17	16	10	10	10	8	**28**
Vehicle 2 Market Share	33	23	13	9	7	11	5

Table shows the percentage of vehicles listed second in the PSID survey that are of a given brand name (given by the column headings) conditional on the brand name of the vehicle listed first (given by the row headings). Since each row is a conditional distribution, the rows sum to one hundred. Sample is restricted to households that own exactly 2 cars; N=24,758. Individual households may appear multiple times if they respond in multiple years of the PSID survey (and own exactly two cars in those years).

vehicle attributes. In the case of a brand preference related to experience, the interpretation of η_{ib} is somewhat subtler: a preference shifter that only activates when the consumer has previous experience with the brand. While η_{ib} may be tied to specific observed and unobserved vehicle attributes, $\eta_{ib} = 0$ before the consumer develops experience with the brand and $\eta_{ib} \neq 0$ afterward. Finally, in the case of information, we can interpret η_{ib} as a risk premium (or better, certainty discount) that applies to a particular brand, due to greater certainty relative to other brands.

Expressions (1), (2), and (3) make clear that correlations in brand choices, even within an individual, are not proof of brand preference according to any of these definitions. Demographic characteristics D_i (e.g., income or family size) and location characteristics L_i (e.g., near a Ford dealer or live in the mountains) in general will lead some consumers to consistently value certain attributes more or less than the average consumer. Thus, if attributes vary systematically by brand (e.g., Toyotas are more fuel efficient), then individual consumers may consistently prefer certain brands, even in the absence of brand preference ($\eta_{ib} = 0$). We address this concern in our empirical specifications by progressively adding controls and ultimately by restricting the sample to brands with similar product offerings.

To begin, table 2 shows the within-fleet distribution in brand choices for households that own exactly two vehicles at the time of the survey. Respondents are asked to list up to three vehicles, in an order that they themselves determine, and to provide additional details on these vehicles. We focus on households that own exactly two vehicles to keep the analysis simple and transparent.

The table shows the percentage of second vehicles that are of the column brand, con-

ditional on the first vehicle being the row brand. For example, the first entry in the table indicates that, of all households who list a GM as their first vehicle, 49% also list a GM as their second vehicle. For comparison, only 33% of respondents overall list a GM as their second vehicle. Every brand shows a similar pattern; conditional on listing a particular brand as the first vehicle, the probability of listing that same brand as the second vehicle is substantially higher, in every instance. For smaller brands, the proportional difference is striking. For example, while only 5% of households overall list a European brand as their second car, a full 28% of households whose first car is a European brand also list a European brand as their second car.

These within-fleet distributions do not control for demographics or location. Thus, we run regressions of new brand choice on past ownership in a sample of vehicles that represents new additions to a household's fleet. We isolate new additions by selecting the survey respondents who report a net increase in their fleet size between survey waves and a number of vehicles acquired since the last wave equal to the net increase in fleet size. Restricting to this sample alleviates concerns that brand persistence might result from dealer preferences regarding trade-ins; if dealers offer higher trade-in prices for the brands they sell, this could induce correlation between the brand of the newly acquired vehicle and previously owned vehicle. Our sample isolates vehicles that represent an expansion of a households fleet and is therefore unlikely to have involved a trade-in.[4]

For this sample, we run linear probability models of the new brand choice on past brand ownership:

$$B_{ijb} = \alpha + \beta \cdot \text{EVER}_{ib} + \delta' \mathbf{X_i} + \varepsilon_{ib}, \tag{4}$$

where B_{ikb} is a dummy equal to 1 if newly acquired vehicle j purchased by individual i is brand b; EVER_{ib} is a dummy equal to 1 if i has ever reported owning a vehicle of brand b in prior years of the survey (including vehicles observed in the current wave that were reported to have been purchased in the past); $\mathbf{X_i}$ are individual controls; and ε_{ib} is an error term. If β is positive, it implies that having owned a brand in the past is predictive of brand choice for new acquisitions.

Table 3 reports results from regressions with varying controls. Each coefficient is from a separate regression in which the dependent variable is equal to one if the newly acquired vehicle is the brand listed in the row. The final column reports the unconditional sample

[4]We have performed a very similar analysis on all vehicles (not just newly acquired ones), on all newly acquired vehicles (not just net fleet size increases), and using a measure of current brand ownership as the independent variable. In all instances, our results are qualitatively similar; the correlation between new brand and past ownership is positive, statistically significant, and economically large.

Table 3: Correlations between new vehicle brand and past ownership of that brand

	No controls	State fixed effects	Full controls	Brand share
GM	0.165***	0.147***	0.164***	0.32
	(0.0145)	(0.0145)	(0.0157)	
Ford	0.122***	0.118***	0.137***	0.22
	(0.0147)	(0.0149)	(0.0159)	
Chrysler	0.0711***	0.0654***	0.0681***	0.13
	(0.0147)	(0.0146)	(0.0153)	
Toyota	0.150***	0.142***	0.139***	0.08
	(0.0186)	(0.0186)	(0.0196)	
Honda	0.154***	0.146***	0.142***	0.07
	(0.0198)	(0.0199)	(0.0213)	
Other Asian	0.0922***	0.0827***	0.0956***	0.13
	(0.0162)	(0.0164)	(0.0176)	
European	0.229***	0.220***	0.192***	0.04
	(0.0277)	(0.0272)	(0.0285)	
N	10,013	10,013	9,266	

Each entry is from a separate linear probability model regression in which the dependent variable is a dummy coded as 1 if the vehicle is of the brand listed in the column. The table entry is the coefficient on a dummy variable coded as 1 if the individual has ever previously owned a vehicle of the same brand. The sample is restricted to vehicles that were newly acquired by the individual since the prior survey wave and that represent a net increase in fleet size. Individuals may appear more than once if they purchased multiple new vehicles. Standard errors (in parentheses) are clustered on the individual level. The "no controls" column includes only a constant term. The "state fixed effect" column includes state fixed effects. The "full controls" column includes demographic controls (age, education, income, urban dummy, gender, number of children in household and family size), state fixed effects, survey year fixed effects, and dummies for the number of vehicles owned by the individual.

means. Individuals can appear more than once if they purchase multiple vehicles, so standard errors are clustered on individual. The first column includes no controls (beyond the constant term). For every brand, there is an economically large and statistically significant correlation between new brand choice and past ownership. For example, GM has a 32% share of all new vehicles, but someone who previously owned a GM is 16 percentage points more likely to buy a new GM than someone who has never owned a GM.

These correlations could be caused by brand preference, or they could be due to location or demographic characteristics that consistently lead individuals to prefer certain brands. We take several steps to explore the importance of such factors. First, in the second column of table 3, we add state fixed effects to control for geographic factors. These controls have very little impact on the estimated coefficients. While our current data do not allow us to measure market shares at a finer level of aggregation, and while relevant geographic factors likely vary within state, the fact that state fixed effects have almost no impact suggests

Table 4: Correlations between new vehicle brand and past ownership of that brand for similar brands

Sample with only GM & Ford vehicles

	No controls	State fixed effects	Full controls	Brand share
GM	0.169***	0.163***	0.186***	0.60
	(0.0196)	(0.0198)	(0.0225)	
Ford	0.160***	0.162***	0.194***	0.40
	(0.0214)	(0.0218)	(0.0239)	
N	5427	5427	5067	

Sample with only Toyota & Honda vehicles

	No controls	State fixed effects	Full controls	Brand share
Toyota	0.176***	0.197***	0.223***	0.52
	(0.0435)	(0.0471)	(0.0526)	
Honda	0.190***	0.204***	0.198***	0.48
	(0.0446)	(0.0468)	(0.0511)	
N	1530	1530	1387	

Each entry is from a separate linear probability model regression in which the dependent variable is a dummy coded as 1 if the vehicle is of the brand listed in the column. The table entry is the coefficient on a dummy variable coded as 1 if the individual has ever previously owned a vehicle of the same brand. The sample is restricted to vehicles that were newly acquired by the individual since the prior survey wave and that represent an increase in fleet size, and where is acquired vehicle is one of the two brands listed in each panel. Individuals may appear more than once if they purchased multiple new vehicles. Standard errors (in parentheses) are clustered on the individual level. The "no controls" column includes only a constant term. The "state fixed effect" column includes state fixed effects. The "full controls" column includes demographic controls (age, education, income, urban dummy, gender, number of children in household, and family size), state fixed effects, survey year fixed effects, and dummies for the number of vehicles owned by the individual.

that brand persistence is not largely a story of geography.[5] In the third column of table 3, we add demographic controls, including age, education, family income, a dummy for urban residence, gender of respondent, number of children in the household, and household family size, as well as year (of survey) fixed effects and dummies for the number of vehicles owned. These demographic controls have almost no effect on the coefficients, suggesting that brand persistence is not simply due to demographic observables.

Finally, we also restrict our regressions to individuals who made a new vehicle choice from a set of very similar brands, to test whether or not brand persistence could be driven by individuals consistently preferring a type of vehicle that is offered by one brand but not the others. The two sets of brands that are the most similar in their offerings are GM and

[5] We are in the process of obtaining information about local market shares and the proximity of dealerships, but these data are not yet available.

Ford, which both offer full lines emphasizing light trucks, and Toyota and Honda, which offer similar sedans and have limited light trucks. Table 4 repeats the specifications of table 3 but limits the sample to GM and Ford acquisitions in the top panel and to Toyota and Honda in the bottom panel. Coefficient estimates continue to be large and highly significant in these selected samples. Overall, the coefficients are slightly larger in magnitude than when looking across all brands.

These within household estimates indicate very strong brand persistence. While it is impossible to eliminate the possibility that these correlations are driven by unobservable demographic characteristics or local geographic factors, the fact that controlling for state fixed effects, controlling for demographic variables, and limiting the choice set to similar brands has minimal impact on the estimates makes a strong *prima facie* case for brand preference. We will add to the evidence in favor of significant brand preference below by showing that correlations across households within a family suggest brand preference spillovers, which are hard to rationalize in the absence of brand preference. Before doing so, however, we first provide a framework for understanding what factors might lead to brand correlations across households, including demographic and geographical correlations, in section 4.

4 Why might brands be correlated across households?

Above, we demonstrated that there are very strong correlations in brand choice within a household. Below, we show that there is also a very strong correlation in brand choice across households within a family. We are interested in determining whether or not this intrafamily correlation implies that brand preferences, in the sense defined above, are correlated across these households. If so, it implies that there are interesting network or experience effects that lead individuals to act upon brand preferences in ways determined by the choices made by other family members. Under what conditions does an intrafamily brand choice correlation imply correlated brand preferences? In this section we provide a formal framework of vehicle demand and brand preference in a simplified choice situation to clarify the interpretation of our empirical results and make clear the identification challenges we face.

We begin by modifying slightly the notation that we introduced in section 3 for the utility that an individual i receives from choosing vehicle j of brand b. We add a subscript to indicate that an individual is in family f, and we assume that each family has two members, so that

$i = 1$ or $i = 2$. We write the utility of the two members of family f as:

$$E[U_{1fj}] = \beta_{\mathbf{1f}}\mathbf{X_j} + \sum_{b=1}^{\mathcal{B}}(\theta_{1fb} \cdot 1(B_{1fj} = b)) + \varepsilon_{1fj} \tag{5}$$

$$E[U_{2fj}] = \beta_{\mathbf{2f}}\mathbf{X_j} + \sum_{b=1}^{\mathcal{B}}(\theta_{2fb} \cdot 1(B_{1fj} = b)) + \varepsilon_{2fj}, \tag{6}$$

where:

$$\beta_{1fk} = \alpha_k + \alpha_{\mathbf{dk}}\mathbf{D_{1f}} + \alpha_{\mathbf{lk}}\mathbf{L_{1f}} + \lambda_{fk} + \pi_{1fk} \tag{7}$$

$$\beta_{2fk} = \alpha_k + \alpha_{\mathbf{dk}}\mathbf{D_{2f}} + \alpha_{\mathbf{lk}}\mathbf{L_{2f}} + \lambda_{fk} + \pi_{2fk} \tag{8}$$

$$\theta_{1fb} = \xi_b + \phi_{bf} + \eta_{1bf} + \gamma \cdot 1(B_{2fj} = b) \tag{9}$$

$$\theta_{2fb} = \xi_b + \phi_{bf} + \eta_{2bf} + \gamma \cdot 1(B_{1fj} = b). \tag{10}$$

These equations differ from equation 1 in two ways. First, the β parameters are now functions of a λ_{fk} term that is common across households within each family. This term denotes unobserved demographic and geographic features that are shared across these households that can induce cross-household correlation in brand choices. Second, the θ terms are now functions of a ϕ_{bf} term and a $\gamma \cdot 1(B_{2fj} = b)$ term. The first term captures correlated brand preferences while the second allows for a direct impact of family ownership on the expected utility that an individual derives form a particular brand.

This formulation makes clear that two individuals in the same family might have correlated utility for a given vehicle j because the individuals have similar observable demographic characteristics ($\mathbf{D_{1f}}$ correlated with $\mathbf{D_{2f}}$), similar locational characteristics ($\mathbf{L_{1f}}$ correlated with $\mathbf{L_{2f}}$), similar unobservable characteristics ($\lambda_{fk} \neq 0$), correlated brand preference ($\phi_{bf} \neq 0$), or a direct ownership effect ($\gamma \neq 0$). Our interest is in isolating effects that operate through either correlated brand preferences or direct ownership effects, rather than correlated characteristics.

Clearly, demographic characteristics matter for vehicle choice. Members of the same family are known to have correlated characteristics, such as income, education, and family size. These demographic characteristics will determine demand for particular vehicle attributes, such as price, size, and fuel economy. To the degree that vehicle attributes are distributed differently across different brands, brand choices will be correlated across individuals in the same family. Location-specific factors, such as rural versus urban residence, weather, or local supply for particular vehicles, will also influence vehicle choices for similar reasons. To the degree that members of the same family live in the same or similar places, this will also

generate correlation in attribute demand that may result in correlated brand choices.

Family members may also have correlated brand preferences in the sense that we have defined it here ($\lambda_{fk} \neq 0$). A correlation in brand preference could be the result of past experiences. Suppose, for example, that a household had a strong preference for GM vehicles. A child who grew up in that household would have accumulated experiences with GMs. Driving a GM might conjure up nostalgic feelings, so that they experience driving a GM differently than driving a Ford with identical attributes. Or, in that same situation, the child may feel that they have significant experience with a GM, so that they are confident that they know what the GM's quality and reliability will be. If individuals within a family have correlated brand preferences for any of these reasons, it will lead to correlation in the brand choices we observe.

These long-run correlations have close parallels in the intergenerational mobility literature, which estimates long-run correlations in earnings, education, and other outcomes of parents and their children. As this literature emphasizes (see especially Solon (1992)), many studies give downward-biased estimates for the correlation in lifetime earnings because lifetime earnings are usually not observed and earnings over shorter periods are a noisy measure of lifetime earnings. In our case, observed correlations in brand *choice* will likely underestimate the intergenerational persistence in brand *preference* if short-run choices are a noisy measure of long-run preferences.

There could be a correlation of brand choices without brand preference, however, if there are direct ownership effects ($\gamma \neq 0$). Imagine an intervention that allocated a particular brand to one member of a family. Would this influence the brand choice of other members of the same family? Recent brand acquisitions by one household in a family could influence other members of the family by making them aware of the new brand or providing experience with it, directly or by word of mouth. This type of short-term phenomenon is potentially distinct from the long-term preference formation we mean to embody in the correlated preferences. Such an effect is, however, conceptually indistinguishable from a more general model in which brand preference is a function of family brand ownership.

These direct, short-run ownership effects have close parallels in the peer effects literature, which often estimate the contemporaneous or near-contemporaneous effects that the choices or characteristics of one's reference group has on one's own choices. A key challenge to identification in this context, as first formalized by Manski (1993), is the reflection problem. In a large-group setting, with aggregate choices and characteristics of one's reference group influencing one's own choice, the reflection problem typically manifests as an inability to determine whether it is the average *choices* of one's reference group or rather the group's average *characteristics* that determine one's own choice. This indeterminacy arises because

the individual choices of one's peers are also a direct function of their individual characteristics. In our small-group setting, since a mother's preferences are potentially a function of her child's choice, and vice-versa, the reflection problem manifests as a correlation between the child's unobserved preference shifters and the mother's choice, which enters as a right-hand-side variable for the child. In effect, our estimates are subject to standard simultaneous equations bias. As described below, we attempt to mitigate this bias by exploiting the distinct timing of when mothers and their children make their choices.

Below, we walk through a number of empirical exercises that attempt to tease apart brand correlations driven by correlated attributes and those driven by brand preference correlation, be it from long-run endogenous preference formation or from information transmission, which may occur in the long or short run. We begin by showing the strength of the raw correlation, which may be driven by brand preferences or correlated characteristics. We then control for a rich set of demographic variables and state fixed effects, demonstrating that these controls account for only a modest fraction of the raw correlation. Then, paralleling our within-household analysis, we focus our attention to two brands that have very similar attributes, Ford and GM. Finally, we attempt to distinguish short and long-run mechanisms by studying how parental vehicle ownership when a child lived in his parents' household impacts the brand choices that he makes when he forms his own independent household.

5 How and why are brand choices related across households?

In this section, we develop a linear probability model for the brand choice of children and demonstrate how children's brand choices are related to their parents' brand choices, using a variety of specifications that try to distinguish correlated brand preference from more mundane causes such as correlated demographics. Future versions of this paper will use a discrete choice model, which confers the advantages that its estimated parameters can be interpreted as parameters of a utility function and that predicted choice probabilities are constrained to the unit interval. For now, however, the linear probability model is valuable in that it can establish the qualitative strength of intra-family correlation in brand choices and explore the robustness of that correlation to a variety of control structures. The linear probability model is also computationally simpler to estimate, particularly given the large number of fixed effects in many of our specifications.[6]

[6]With most of the fixed effects considered here, such as the brand-by-year or brand-by-state dummy variables, the only concern with a discrete choice model is computational time. With the brand-by-individual fixed effect specification, however, the discrete choice model also suffers from inconsistent estimates due to

In section 3, we demonstrated the persistence of brand choices within households using linear probability model regressions for each brand, one at a time. This approach is useful in retaining an intuitive setup, but it is limiting in other respects. Here, we generalize this approach to allow us to study all brands in a single regression. Specifically, we categorize all vehicle choices as being one of the seven "brands" that we used above: GM, Ford, Chrysler, Toyota, Honda, Other Asian, and European. Grouping smaller Asian automakers and European manufacturers together ensures that each brand is chosen frequently enough to yield meaningful estimates in a linear probability framework. In particular, these brand definitions imply that their choice probabilities lie in the 3%–20% range in the raw data. For each vehicle purchase by every individual in our data, we expand the original data sample to include seven lines of data. The first is for the brand that was chosen by the individual, and this line has the dependent variable coded as one. The other six are observations with a zero dependent variable, one for each of the six brands not chosen. To avoid this expansion of the data set unduly shrinking our standard error estimates, we cluster all standard error calculations on the individual (child) level. This clustering also accounts for correlation in the individual's choices across choice situations (vehicles purchased).

We do not allow for an outside good, which would be interpreted as the option not to purchase a vehicle at all. Inclusion of an outside good is a standard step in discrete choice modeling, but here we are interested in knowing whether or not a child, conditional on purchasing a vehicle, decides to buy a brand that is the same as the one owned by members of his or her family. Inclusion of an outside good would mix together correlations in choice that determine whether or not individuals purchase vehicles with correlations in the brand chosen when purchasing a vehicle, which are distinctly different economic phenomena.

5.1 Brand correlations within family networks

We begin by showing simple correlations in order to demonstrate the strength of the intrafamily relationship, and then demonstrate the impact that controlling for observable characteristics has on the results. Table 5 reports results from several specifications of our linear probability model, where the dependent variable for each observation (each brand choice possibility for a child) is coded as 1 if the child selected that brand. The independent variable of interest is coded as 1 if the mother chose that brand for her most recent vehicle purchase that was prior to the child's purchase.[7] Column 1 reports the raw correlation,

the incidental parameters problem.

[7]Throughout, we focus on the mother-child relationship because of our superior ability to match such households in the data. A regression on just the father's choice yields an estimate that is similar to the mother's choice alone; the coefficients are not statistically different. For a regression in which both the mother's and father's most recent purchases are included, each coefficient is around half of the coefficient

Table 5: Correlations between child brand choice and parental brand choice

	Dependent Variable: Child's Brand				
VARIABLES	(1)	(2)	(3)	(4)	(5)
Mother's Brand	0.104***	0.098***	0.097***	0.085***	0.063***
	(0.006)	(0.006)	(0.007)	(0.007)	(0.006)
Mother's Previous Brand					0.053***
					(0.006)
Child's Previous Brand					0.164***
					(0.007)
Brand x year fixed effects	Yes	Yes	Yes	Yes	Yes
Child's demographics	No	Yes	Yes	Yes	Yes
Mother's demographics	No	No	Yes	Yes	Yes
State x brand fixed effects	No	No	No	Yes	Yes
Observations	95,914	86,170	77,777	77,777	77,777
R-squared	0.084	0.098	0.102	0.124	0.146

Standard errors clustered by individual in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Each column is a linear probability model where each individual-year-vehicle choice enters the data 7 times, once for each brand. The brands include: GM, Ford, Chrysler, Toyota, Honda, Other Asian, and European. Child's and mother's demographics include age, education, income, urban dummy, gender, number of children in household, and family size.

with only year-brand fixed effects included, which control for the overall share of each brand during each period. All standard errors are clustered on the individual level to allow for arbitrary correlation across the choices of each individual, which are mechanically related. The coefficient estimate is 0.104, precisely estimated, and economically large. There are 7 brands in our choice set, so the average brand is selected 14.3% of the time. Our estimate therefore indicates that a consumer's mother purchasing a brand increases the likelihood that the consumer purchases the brand by over 70%, on average.

In columns 2 through 5 of table 5, we introduce progressively richer controls. Recall from our discussion in section 4 that we wish to isolate correlations across households that are not due to demographic and location-specific factors that cause related households to demand similar attributes. To assess the importance of such factors, we first introduce a set of demographic controls for the child and his or her household in column 2: family income, urban versus rural dummy, age, sex, education, number of kids in the household, household size, and the gasoline price when the child purchased the vehicle. We interact each of these

in column 1 of table 5. Identification in this case comes from individuals whose parents live in separate households. We have also experimented with a variety of ways of characterizing the mother's choice given that mothers may own multiple vehicles. We have used the share of mother's vehicles of a specific brand, and we have done one-to-one vehicle matching across mothers and children using the order in which vehicles are listed in the survey. In all cases, our qualitative results are quite similar. We prefer the most recent vehicle the mother purchased prior to the child's purchase because the interpretation of a dummy variable is more transparent and because this selection best accords with the causal model we have in mind.

with our seven brand dummies, and we include year-brand fixed effects to capture average market shares. We drop observations that are missing information in the relevant survey year for any of these demographic controls, which reduces our sample size. The coefficient changes very little; even after controlling for a rich set of demographics, a child whose mother's most recent vehicle acquisition was of a particular brand is on average 69% more likely to choose that same brand than a child whose mother did not.

In column 3, we introduce a matching set of demographic controls for the mother and her household. These controls reflect the fact that the mother's attributes might influence the child's optimal choice of vehicle; for example, if the mother is elderly, the child might prefer a vehicle that is easy to enter and exit. These controls reduce the intergenerational correlation by a very small amount. In column 4 we add state by brand fixed effects that allow the average market share of each brand to differ for each state. The purpose of these controls is to capture location-specific factors in a nonparametric fashion. These fixed effects cause the coefficient to fall, but again by a small amount. Comparing the raw correlation of 0.104 with the estimate of 0.085 in column 4, we see that our rich set of controls moves the raw correlation less than 25%.[8] While these regressions cannot rule out unobservable demographic and location-specific factors that might be correlated with preferences for vehicle attributes, we believe that the limited sensitivity of intergenerational brand correlations to these rich controls is important evidence that they are driven by brand preferences rather than confounders. Additional evidence is provided in section 5.3, in which we restrict our attention to brands offering vehicle fleets with similar attributes.

The mother's brand coefficient in column 4 incorporates the possibility that mothers and children may have similar past experiences with vehicle brands. Thus, this coefficient may be capturing something about the child's own brand loyalty. We investigate this possibility by including lagged purchases of both the mother and child in column 5 of table 5. The lagged variables are equal to one if the previous purchase observed in the data was of the brand and zero otherwise. If there is no previous vehicle purchase by the mother or child in our dataset, then the lagged variables are equal to zero for all observations. We find that both the mother's and child's lagged purchases are statistically and economically significant determinants of the child's brand choice, with estimated coefficients of 0.053 and 0.164,

[8] In order to test whether the correlation across generations is coming from a correlated preference for brand (e.g., Ford) or sub-brand (e.g., Ford, Lincoln, or Mercury), we have run the regression in column 4 of Table 5 with 41 sub-brands instead of the 7 brands (and the consumer demographics interacted with the 41 sub-brands). We find that both the overall brand and the sub-brand of the mother have a statistically significant correlation with the sub-brand chosen by the adult child. The brand coefficient is 0.00633 (0.00097) and the sub-brand coefficient is 0.0489 (0.0043). The number of observations jumps to 466,662 and the R^2 jumps to 0.146.

respectively.[9]

5.2 Suggestive Heterogeneity

One interesting aspect of the PSID is that while we can follow an extended family over time, new consumers are constantly marrying into the family. We do not observe the family histories of these new family members; thus, within each couple we observe the vehicle purchases of the mother of one member but not the other. If couples are influenced by the vehicle choices of both sets of parents, then we should observe that mothers' brand choices are more important for single children than for children living with spouses or significant others.[10] Table 6 shows exactly this result. Column 1 replicates column 5 from table 5, and then in column 2 we run this regression only for single adult children. Column 3 shows the result only for coupled adult children. While the child's own lagged brand choice is not statistically different between single and coupled children, the effect of the mother's previous brand purchase coefficient drops from 0.109 to 0.052 for single versus coupled children, while the mother's second most recent brand purchase correlation drops from 0.073 to 0.049.

5.3 Similar Brands

Of principle concern is that demographic or locational characteristics of children and mothers will be correlated and that these characteristics drive demand for vehicle attributes that are correlated with brand. For example, individuals who live in rural areas, work in construction jobs, or live in areas that receive heavy snowfall may be more likely to prefer light trucks to passenger cars. Because GM's fleet is more heavily tilted towards light trucks than Toyota's, such people will be more likely to buy a GM, even in the absence of any brand preference. Above, we showed that controlling for many observable characteristics had limited impact on the coefficients, but concerns about unobserved factors remain.

[9]When we add a second lag of both the mother's and child's brand choices, we find that all coefficients are statistically different from zero and that, for both the mother and child, the size of the correlation with the child's current brand choice is decreasing in the number of lags. Nonetheless, including these lags in the model causes only a mild decrease—from 0.085 to 0.063—in the coefficient on the mother's most recent purchase. We take this result as evidence that the passing of brand preferences from parents to children operates, at least in part, through a short-run mechanism of information transmission. The significance of mothers' lagged purchases is suggestive that a longer-run preference formation mechanism may also be at work, but it is difficult to cleanly distinguish this mechanism from that of information transmission. We further explore the issue of long-run preference formation and inheritance in section 5.4, in which we study individuals whom we observe as teenagers living with their parents early in the sample and then as independent household heads or wives later in the sample.

[10]Coupled adults include both those who are married as well as those who are living with a significant other but not married.

Table 6: Correlations between child brand choice and parental brand choice for subsamples

	Dependent Variable: Child's Brand		
VARIABLES	All	Single	Coupled
Mother's Brand	0.063***	0.109***	0.052***
	(0.006)	(0.013)	(0.007)
Mother's Previous Brand	0.053***	0.073***	0.049***
	(0.006)	(0.014)	(0.007)
Child's Previous Brand	0.164***	0.153***	0.160***
	(0.007)	(0.015)	(0.008)
Brand x year fixed effects	Yes	Yes	Yes
Child's demographics	Yes	Yes	Yes
Mother's demographics	Yes	Yes	Yes
State x brand fixed effects	Yes	Yes	Yes
Observations	77,777	16,590	61,187
R-squared	0.146	0.182	0.149

Standard errors clustered by individual in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Each column is a linear probability model where each individual-year-vehicle enters the data 7 times, once for each brand. The brands include: GM, Ford, Chrysler, Toyota, Honda, Other Asian, and European. Child's and mother's demographics include age, education, income, urban dummy, gender, number of children in household and family size.

Here, we use a different strategy, isolating the choice set to brands that are very similar. Ford and GM are both full-line automakers that compete directly in every vehicle segment. Because their vehicle lineups are very similar, we would expect intrafamily brand choice correlations to be quite weak in the absence of intrafamily brand preference transmission when we limit our sample to children who choose either a GM or a Ford. Similarly, Toyota and Honda both produce a full range of high-quality cars as well as fuel efficient SUVs and trucks. Thus, if correlated preferences for vehicle attributes are driving our results in Table 5 then we would expect the correlation to disappear when we limit our sample to children who chose either a Toyota or a Honda.

Table 7 repeats the specification from table 5, column 5 for these two subsamples. Specifically, we keep all instances in which a child chose either a Ford or GM (or a Toyota or Honda),[11] and we code the mother's choice as before. The sample size is different than the results in table 5 both because we limit the sample to these choice situations and because we restrict the choice set to Ford or GM (or Toyota or Honda), moving us from 7 brand choices (and hence 7 observations per choice) to 2 brand choices (and hence 2 observations per choice).

[11] Around 56% of our original sample chose a Ford or GM and 14% of our original sample chose a Toyota

Table 7: Correlations between new vehicle brand and parental brand choice for similar brands

	Dependent Variable: Child's Brand		
VARIABLES	All Brands	Just Ford/GM	Just Toyota/Honda
Mother's Brand	0.063***	0.064***	0.124***
	(0.006)	(0.013)	(0.035)
Mother's Previous Brand	0.053***	0.056***	0.097**
	(0.006)	(0.013)	(0.040)
Child's Previous Brand	0.164***	0.202***	0.181***
	(0.007)	(0.013)	(0.032)
Brand x year fixed effects	Yes	Yes	Yes
Child's demographics	Yes	Yes	Yes
Mother's demographics	Yes	Yes	Yes
State x brand fixed effects	Yes	Yes	Yes
Observations	77,777	12,672	2,974
R-squared	0.146	0.128	0.150

Standard errors clustered at the child level in parentheses. *** p<0.01, ** p<0.05, * p<0.1 Columns 2 and 3 are linear probability models in which each individual-year-vehicle enters the data 2 times, once for GM and once for Ford or once for Toyota and once for Honda. The data are restricted to children who own a GM or a Ford and whose mother owns a GM or a Ford (or a Toyota or Honda). Child's and mother's demographics include age, education, income, urban dummy, gender, number of children in household, and family size.

Column 1 of table 7 reproduces column 5 of table 5, for comparison. Columns 2 and 3 show that, for both the Ford/GM and Toyota/Honda samples, mothers' brand choices have a positive, economically large, and statistically significant effect on their children's brand choices. Column 2 shows that a child whose mother's most recently purchased vehicle was a GM is 6.4 percentage points more likely to buy a GM, conditional on the child's previous brand choice and the mother's second most recent brand choice. In this subsample, the choice share of GM is 61% and that of Ford is 39%. This intrafamily brand correlation would therefore raise the probability that a child purchased his mother's brand by about 10% for GM and 15% for Ford. These impacts are smaller than the corollaries from table 5, which are approximately 20% for GM and 30% for Ford. This result is intuitive given that Ford and GM are generally close substitutes, and it does suggest that correlations in observable household characteristics that influence demand for vehicle attributes may have contributed to the brand correlations in table 5. However, a strong intrafamily brand correlation remains, even after limiting the sample to similar automakers.

or Honda.

5.4 Intrafamily brand influence through past experience

A link between the brand preference of parents and the brand preference of children could operate through an information transmission mechanism and/or a long-run preference inheritance mechanism. The results above, which show that a mother's recent purchases are important determinants of a child's purchases, even after conditioning on both the mother's and child's lagged purchases, suggest that the information transmission mechanism is at work. Here, we explore the long-run preference inheritance mechanism by examining how the brand choices of parents when their children still live at home influence the future brand choices of those children after they become adults and move out.

To study long-run preference formation and inheritance, we would ideally be able to observe, for each respondent, both his parents' purchases while he was still living with them and his own purchases many years after he moved out. However, because vehicle ownership data are only available in the PSID from 1999-2007, we are constrained to studying young adults who only recently moved out of their parents' home. Thus, our ability to distinguish a preference inheritance mechanism from an informational mechanism is not as clean as we would like.

We carry out our analysis by selecting a sample of children whom we observe living in their parents' household in one year but observe living in their own household in a later year. We restrict this sample further to include individuals who did not form their own household until they were at least 18 years old and were not living with their parents beyond age 23. For these individuals, we then identify the brand ownership of their parents at the last time we observe them living in their parents' household.

Table 8 reports linear probability model regressions that predict the child's brand choice (once they have established their own household) using a dummy variable for whether or not any of the parents' vehicles at the time the child left the household were of a given brand. To avoid counting vehicles that children took with them from their parents' home, we drop all children who reported that any of the vehicles in their fleet had been received as a gift. The difference in the independent variable in these regressions relative to those discussed above means that we cannot directly compare the estimated coefficients to our previous estimates. However, the results from this analysis are qualitatively similar. Child brand choice is positively correlated with mother's ownership of a brand when the child was most recently in the mother's household, and the relationship is economically large and statistically significant. Table 8, column 1, which controls for brand-year fixed effects and the child's demographics, indicates that an individual is 6 percentage points more likely to own a given brand if his mother's household owned that brand somewhere in its fleet when he last lived in the household. Adding the mother's characteristics to the regression, in column

Table 8: Correlations between new vehicle brand and parental ownership at time when child lived with parents

VARIABLES	(1) Child's Brand	(2) Child's Brand	(3) Child's Brand
Mother Owned Brand	0.0617***	0.0606***	0.0375***
	(0.0147)	(0.0150)	(0.0172)
Brand x year fixed effects	Yes	Yes	Yes
Child's demographics	Yes	Yes	Yes
Mother's demographics	No	Yes	Yes
State x brand fixed effects	No	No	Yes
Observations	2,912	2,898	2,898
R-squared	0.119	0.132	0.250

Standard errors clustered at the child level in parentheses. *** p<0.01, ** p<0.05, * p<0.1 Each column is a linear probability model where each individual-year-vehicle enters the data 7 times, once for each brand. The data are restricted to children who were observed as a dependent in a household in the sample when he or she was under 23 and then as the head (or wife) of a household later in the sample. Child's and mother's demographics include age, education, income, urban dummy, gender, number of children in household and family size.

2, does not substantially affect this result, and in column 3 the coefficient falls by more than one third when the state fixed effects are added. This decline is slightly larger than in prior tables. Still, these results suggest a link between brand experience during adolescence and brand choice in adulthood, suggestive of long-run endogenous brand preference formation. In future work, we aim to further probe the importance of local characteristics through use of data on dealer networks and local market shares.

6 Implications of brand preferences for the vehicle market

What might the persistence of brand preferences imply for market outcomes in the automobile industry? We begin to address this question by considering the implications of brand preferences for automakers' pricing strategies. We study a simple, symmetric model in which two firms compete in two different product markets and consumers live two periods, purchasing once in each market. Our modeling approach follows the tradition of the switching costs and brand loyalty literature—particularly Klemperer (1987) and Dubé et al. (2009)—focusing on whether brand loyalty increases or decreases prices in equilibrium. We follow this analysis by studying a model in which consumers may be loyal to one firm but not the other, motivating a discussion of automakers' incentives to encourage brand loyalty amongst

their customers.

We forego a richer model that would more closely match the current automobile industry—a model with more than two major firms, many products per firm, and richly differentiated consumer preferences—for several reasons. First, the simultaneous estimation of the parameters needed to simulate the model (those governing households' preference heterogeneity, households' brand preference transmission, and firms' marginal costs) would be a substantial undertaking that is beyond the scope of this paper and likely beyond the power of our data.[12] Second, the computational challenges of simulating such a model would be immense. Finally, the simple model we present is close in spirit to most of the brand loyalty literature and provides clear, intuitive results that we believe would generalize qualitatively to a more complex model.[13]

6.1 A simple model of automobile pricing under brand loyalty

A substantial literature, dating back to Klemperer (1987), studies the effect of brand preferences on firms' prices in equilibrium. This literature most often refers to what we call brand preferences as "switching costs," though the concept is often made operational in the same way as is done here.[14] Our approach is closest in spirit to that of Dubé et al. (2009), hereafter DHR, in that we model an infinite-horizon game in which brand preferences are of a sufficiently modest magnitude that some households do switch brands in equilibrium. Our primary difference from DHR and the preceding literature is that we model firms as having two products that cater to two types of consumers: the young and the old. We model multi-product firms to relate the model more closely to the automobile market, in which nearly all manufacturers produce a range of models tailored to consumers in different stages of their lifecycle, and to highlight the role that intergenerational preference transmission can play in determining automobile prices in equilibrium.

In our model, there are two symmetric firms, denoted j and k, that compete in a differ-

[12] Dubé et al. (2009) and Dubé, Hitsch and Rossi (2010) are able to simultaneously estimate preference heterogeneity and within-household brand loyalty because they observe both a large number of repeat purchases per customer and rich price variation in their dataset on orange juice and margarine purchases. While our PSID dataset is well-suited for estimating intergenerational brand preference transmission, the limited number of purchases observed for each household and weak price data make it poorly suited for characterizing heterogeneous preferences for attributes or prices.

[13] We are encouraged here by the fact that Dubé et al. (2009) find qualitatively similar predictions from the simple and complex versions of their model.

[14] Some papers model switching costs as an increase in utility from purchasing the same brand as that purchased last period (our approach), while others model switching costs as a decrease in utility from purchasing a different brand. Dubé et al. (2009) examines both models and finds that they produce identical predictions in the absence of an outside good. In the presence of an outside good, the second formulation yields lower prices in equilibrium, as switching costs make the outside good relatively more appealing.

entiated Bertrand pricing game. We adopt an infinite horizon overlapping generation (OLG) framework in which households live for two periods. In each period, there are unit masses of two types of households: young (type A) and old (type B). All consumers are born as type A, become type B in the second period of their lives, and then die, creating a new type A consumer (offspring) upon death. A key feature of the model is that the type A and type B consumers purchase different kinds of cars. Both firms are aware of this fact, and both sell two vehicle models catering to the two types. Thus, there are four vehicles in the market: jA, jB, kA, and kB. A and B cars can be thought of as cars preferred by younger versus older consumers, or entry level versus upscale, or single person versus family vehicles.

For both brevity and clarity, we will focus on the case in which type A households only consider vehicles jA and kA and type B households only consider vehicles jB and kB. Clearly this is an abstraction, as there will be some substitution across any of the markets we described. Thus, this paper's appendix (to be added) will discuss how relaxing this assumption affects our findings.[15] Still, in a survey of over 22,000 consumers by a market research firm described in Langer (2011), the Cadillac Deville and Lincoln Town Car had more than 100 purchasers over the age of 60 and none under the age of 40, while the Scion tC had more than 100 purchasers under 40 and only 6 over 60.[16] Clearly there are vehicles that appeal strongly to specific age groups.

Let the utility of a particular consumer i of type B that purchases vehicle jB be given by:

$$U_{ijB} = V - \alpha P_{jB} + \mu_B 1\{b_{iA} = j\} + \varepsilon_{ijB},$$

where V is a baseline utility that is common across the two brands, P_{jB} is the price of vehicle jB, and $1\{b_{iA} = j\}$ is an indicator for whether consumer i purchased brand j when he/she was a type A last period. μ_B denotes the strength of within-consumer persistence of brand preferences. The utility from purchasing vehicle kB is given similarly. All type B households purchase exactly one vehicle, and there is no outside good.

[15] In brief, allowing for some cross-age substitution has essentially no impact on models in which intergenerational brand preference transmission is as strong as within-household transmission. In models in which intergenerational transmission is relatively weak, cross-age substitution reduces the gap between the type A and type B vehicle prices (and does so in a qualitatively symmetric way). The result that brand preferences (of a magnitude corresponding to our estimates above) reduce equilibrium prices continue to hold. This is true even in the extreme case in which there is no intergenerational brand preference and consumers have no systematic preference for their own type of vehicle. This last model is similar to that of Doganoglu (2010) in which consumers live for two periods and the (single product) firms cannot distinguish between young and old.

[16] In that same survey, only 5% of consumers who say they purchased a Buick are under the age of 40.

The utility of a consumer i of type A that purchases vehicle jA is similarly given by:

$$U_{ijA} = V - \alpha P_{jA} + \mu_A 1\{b_{iB} = j\} + \varepsilon_{ijA}.$$

Here, $1\{b_{iB} = j\}$ is an indicator for whether the parents of consumer i purchased brand j when they were type B last period. μ_A denotes the strength of intergenerational brand preferences. In accordance with our estimates above, we will model μ_A as less than or equal to μ_B. The random utility components ε_{ijB} and ε_{ijA} are assumed to be i.i.d. type I extreme value over individuals i, brands j and k, and types A and B.

For now, we assume that type A consumers are not forward-looking when deciding whether to purchase vehicle jA or kA; we will relax this assumption in future work.[17] We also assume that type B consumers are not forward looking in the sense that they do not consider the implications of the brand preferences they transmit to their children.

Let ϕ_A and ϕ_B denote the fraction of consumers loyal to brand j in the A and B markets, respectively. Given the price of each vehicle and ϕ_A and ϕ_B, the demand for each vehicle will be given by a weighted sum of standard logit choice probabilities. For example, the demand for vehicle jA is given by:

$$D_{jA} = \phi_A \frac{\exp(V - \alpha P_{jA} + \mu_A)}{\exp(V - \alpha P_{jA} + \mu_A) + \exp(V - \alpha P_{kA})} + (1 - \phi_A) \frac{\exp(V - \alpha P_{jA})}{\exp(V - \alpha P_{jA}) + \exp(V - \alpha P_{kA} + \mu_A)}.$$

We model the marginal cost of all four vehicles in the market as a constant, denoted by c. Firm j's per-period profits are then given by:

$$\pi_j(P_{jA}, P_{kA}, P_{jB}, P_{kB}, \phi_A, \phi_B) = (P_{jA} - c) \cdot D_{jA}(P_{jA}, P_{kA}, \phi_A) + (P_{jB} - c) \cdot D_{jB}(P_{jB}, P_{kB}, \phi_B)$$

In the infinitely repeated game, the firms' state variables are the brand loyalty shares ϕ_A and ϕ_B of the consumers of each type. The states evolve so that next period's loyalty of the type A consumers is given by the current period's demand of the type B consumers for vehicle jB: $\phi_A^{'} = D_{jB}(P_{jB}, P_{kB}, \phi_B)$. Similarly, $\phi_B^{'} = D_{jA}(P_{jA}, P_{kA}, \phi_A)$. We restrict the firms to Markov strategies so that, with a discount factor δ that is shared by the two firms, firm j's Bellman equation is given by:

$$V_j(\phi_A, \phi_B) = \max_{P_{jA}, P_{jB}} \{\pi_j(P_{jA}, P_{kA}, P_{jB}, P_{kB}, \phi_A, \phi_B) + \delta V_j(\phi_A^{'}, \phi_B^{'})\}$$

[17] Per the intuition of Somaini and Einav (2011), we expect that allowing for forward-looking behavior by the type A consumers will result in higher prices for the type A vehicles and lower prices for the type B vehicles in equilibrium.

Firm k's Bellman equation is defined similarly. These equations capture the tradeoff the firms face as the parameters μ_A and μ_B governing the strength of brand loyalty increase. The incentive to increase current-period profits by increasing prices is weighed against the incentive to increase future profits by lowering prices to boost the share of future loyal consumers.

For a given set of model parameters, the Markov Perfect Equilibrium (MPE) of the firms' dynamic Bertrand pricing game can be solved computationally (details to be provided in a future appendix).[18] In the simulations presented below, we fix $\delta = 0.9$, $V = 1$, $\alpha = 8$, and $c = 1$. The choice of V is immaterial in the absence of an outside good. The price preference α and marginal cost c parameters together yield, in the absence of any brand preferences, an equilibrium price for all vehicles of 1.25 and equilibrium own-price elasticities of -5. This markup and elasticity roughly correspond to typical markups and elasticities found by Berry, Levinsohn and Pakes (1995).

The range of brand loyalty parameters μ_A and μ_B that we consider spans zero to one. Values of zero collapse the model to a standard static Bertrand problem, for which the equilibrium price is 1.25. Values of within-household brand preference, μ_B, between 0.4 and 0.6 most closely correspond to the linear probability model estimates discussed above. These estimates also suggest that μ_A is roughly 1/2 to 3/4 the value of μ_B.

6.2 Optimal prices in a model with symmetric firms

We explore the impact of brand preferences on firms' equilibrium pricing strategies by increasing the brand preference parameters μ_A and μ_B from zero and examining the change in firms' equilibrium steady state prices. These prices are sufficient statistics for steady state profits, since in steady state the two firms always evenly share both the A and B markets (due to the symmetry of the firms' demand and cost parameters).

Figure 1 presents steady state equilibrium prices, over a range of brand loyalty strengths, for three cases. For all cases, the prices of firms j and k are equal within each of the markets A and B due to symmetry. In the first case, given by the solid lines, intergenerational brand transmission is turned off by holding $\mu_A = 0$ while the strength of within-household brand preference is varied by letting μ_B range from 0 to 1. In this case, we find that increasing μ_B raises the prices of the type B cars while lowering the prices of the type A cars. That is, when households develop brand loyalty but do not pass this loyalty to their children, then in equilibrium the prices for vehicles intended for older consumers will be high relative

[18] Without intergenerational brand loyalty ($\mu_A = 0$), the model reverts to a standard two-period game (akin to that of Klemperer (1987)) that can be solved for analytically, though the results presented below were nonetheless generated on the computer.

Figure 1: Steady state prices with two symmetric firms

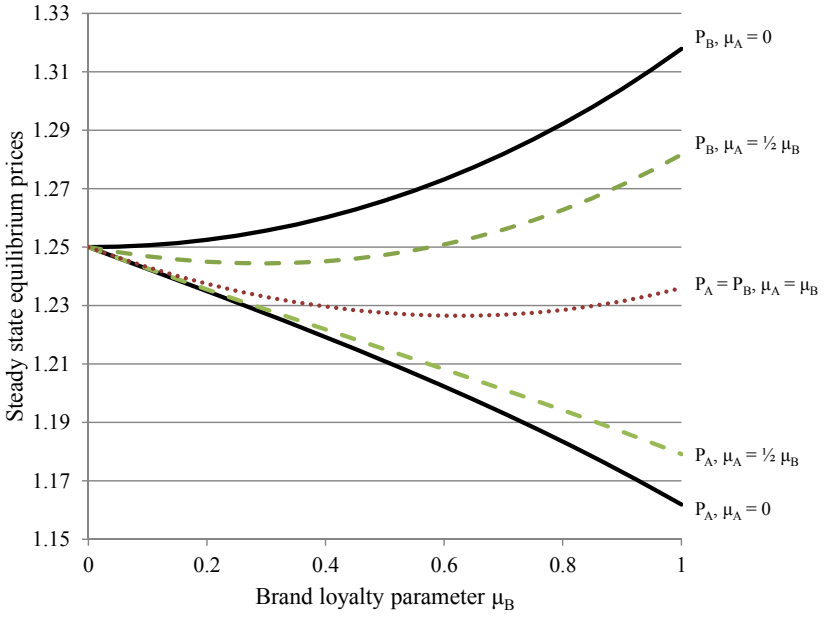

Note: Steady state equilibrium prices shown are from the model described in section 6.1 in which $\delta = 0.9$, $V = 1$, $\alpha = 8$, and $c = 1$. At steady state, the demand for each of the four cars jA, jB, kA, and kB is equal to 0.5. The solid line denotes the case in which there is no intergenerational brand loyalty, the dashed line denotes the case in which intergenerational brand loyalty is half the strength of within-household brand loyalty, and the dotted line denotes the case in which intergenerational and within-household brand loyalty are equal.

to prices for vehicles intended for younger consumers. The intuition for this result follows directly from Klemperer (1987): if first period choices determine brand loyalty in the second period, then firms will "invest" in customers in the first period by charging lower prices and "harvest" the consumer loyalty in the second period. Over the range of μ_B parameters plotted, the "investment" effect in the A market outweighs the "harvesting" effect in the B market in that the average vehicle price is less than the no-loyalty baseline price of 1.25.[19]

When intergenerational brand loyalty is equal to within-household brand loyalty—the case denoted by the dotted line in figure 1—the A and B markets behave identically to one another so that the prices for all four vehicles are equal in steady state, and the model collapses to that of DHR. Relative to the case with no intergenerational brand loyalty, type B consumers benefit and type A consumers lose as the firms no longer price their type B

[19]The lowest average price in this case is 1.238, occurring near $\mu_B = 0.7$.

cars higher than their type A cars. In particular, firms can no longer "harvest" brand loyalty through high mark ups in the B market because doing so reduces future demand and profits.

Across the full range of brand loyalty parameters plotted, steady state equilibrium prices are lower than in the case of no brand loyalty. The dashed line plots an intermediate case in which intergenerational brand preference transmission, μ_A, is one-half that of within-household transmission, μ_B. This case lies between the two other cases.

Given that the values of μ_A and μ_B that apply to the automobile industry almost certainly lie within the plotted range, these results suggest that the transmission of automobile brand preferences within and across households reduces automakers' prices and profits in equilibrium. Similar to DHR, we find that prices do increase with brand loyalty at the high end of the range. However, the point at which prices are actually higher than in the no-loyalty case—values of μ_A and μ_B of approximately 1.2—seems implausible given our estimates above.[20] The intuition for why moderate levels of brand loyalty reduce prices in equilibrium is given by Cabral (2009), which discusses why the "investment" incentive to lower prices is first-order while the "harvesting" incentive to raise prices is second-order.

Figure 2 displays equilibrium prices and demand out of steady state for the "intermediate" model given by the dashed line in figure 1. Figure 2 fixes the brand preference transmission parameters at $\mu_B = 0.5$ and $\mu_A = 0.25$, roughly corresponding to our empirical results above. The plot shows that, as the share of consumers loyal to firm j in both the A and B markets increases from zero to one, firm j's prices and demands in both markets increase. The market B price is more sensitive to the initial level of brand loyalty because the relative weakness of intergenerational brand preferences means that the "harvesting" incentive is relatively strong in this market. That is, if firm j finds itself with a large number of loyal type B consumers, it has a strong incentive to raise the price of vehicle jB in equilibrium to profit from these consumers. Figure 2 also demonstrates that, starting from any initial state, the price and demand dynamics will drive the market to the 50/50 steady state.

We believe that the results of our analysis speak to the automobile industry's apparent focus on sales volumes to the potential neglect of current profits. The industry media is filled with stories about market share, sales volumes, and conquest rates. Anecdotally, automakers are said to focus on hitting quarterly sales targets, which frequently leads them to discount vehicles and even dump some at a loss in fleet sales. It is natural for an economist to view such prioritization of sales volumes over profitability as a mistake. When brand preferences are present, however, firms must make trade-offs between current profits and future profits,

[20] A value of $\mu_B = 1.2$ implies that a household loyal to brand j facing equal prices for vehicles jB and kB would choose jB with a probability of 0.769. Such a choice probability would imply a coefficient of 0.537 in a linear probability model with two firms; this value far exceeds the magnitudes of the brand persistence we observe in the data.

Figure 2: Equilibrium non-steady state behavior; model with two symmetric firms

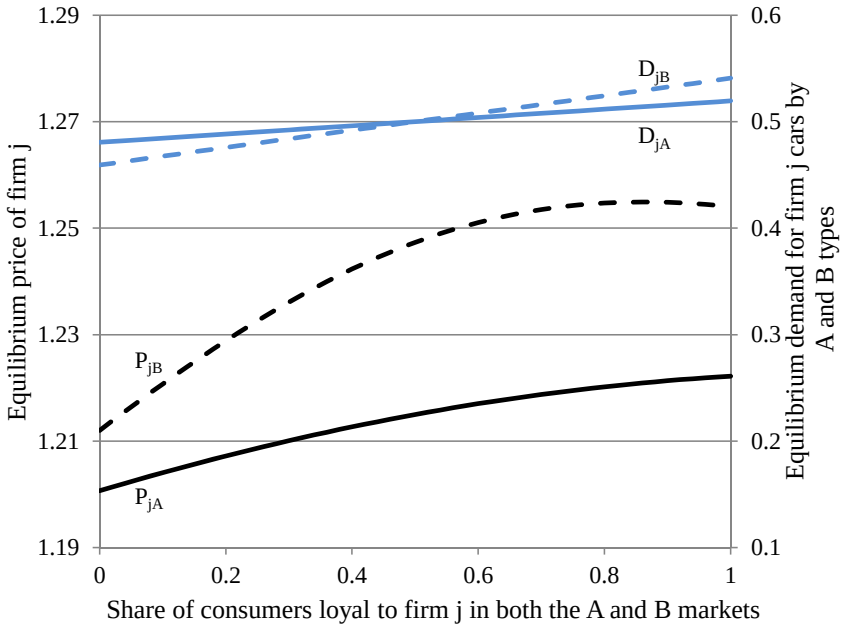

Note: Equilibrium prices shown are from the model described in section 6.1 in which $\delta = 0.9$, $V = 1$, $\alpha = 8$, and $c = 1$. μ_B, which governs the strength of within-household brand preference transmission, is 0.5, and μ_A, which governs intergenerational transmission, is 0.25. As one moves from zero to one on the horizontal axis, the state variables ϕ_A and ϕ_B denoting the shares of consumers loyal to firm j in the A and B markets both move from zero to one.

justifying a focus on volume. When brand preferences are transmitted across generations, the importance of brand preference is enhanced. In particular, intergenerational transmission sharply limits automakers' ability to harvest brand preference later in consumers' lifecycle, as harvesting jeopardizes the loyalty of future generations. Overall, automakers face strong pressure to cut prices today to compete on market share; however, it is difficult for them to ever reap the rewards of this customer loyalty.

6.3 Firms' incentives for encouraging brand loyalty

One of the main results from the analysis above is that the presence of brand loyalty likely reduces automakers' equilibrium prices and profits. This result begs the question of why, then, automakers appear to frequently encourage brand loyalty amongst their customers. For example, firms in this industry typically offer products that are vertically differentiated but nonetheless have similar brand-specific attributes (such as the location of the radio dials

or the front grill styling), often focus their marketing and advertising efforts on their overall brand rather than on individual products, usually have dealers that sell their full range of vehicles, and occasionally even place advertisements that explicitly appeal to nostalgia.[21] Clearly, however, the industry would be better off if its consumers did not develop brand preferences. To address this issue, we now explore firms' unilateral incentives to develop loyal consumers by studying a game in which the strength of brand preferences can vary across firms.

We begin with the model from the section above and, for simplicity, focus on the case in which intergenerational and within-household brand preferences are of the same magnitude. Thus, the A and B market prices will be equal in steady state. We create an asymmetry across the two firms by eliminating any brand preference for firm j. Mechanically, we do so by breaking the original brand loyalty parameters μ_A and μ_B into four parameters: μ_{jA}, μ_{jB}, μ_{kA}, and μ_{kB}. μ_{jA} and μ_{jB} apply if the consumer (or the consumer's parents) purchased brand j last period, while μ_{kA} and μ_{kB} if brand k was purchased. We set μ_{jA} and μ_{jB} equal to zero and vary μ_{kA} and μ_{kB} to examine how firm k is affected when consumers can be loyal to it but not to firm j.

The results of our analysis are presented in figure 3. This plot demonstrates that as μ_{kA} and μ_{kB} increase, firm k modestly decreases its prices to invest in brand loyalty, but firm j must substantially decrease its price in order to be competitive in the absence of any brand preference for its vehicles. Despite this price decrease, firm j's steady state market share still declines with μ_{kA} and μ_{kB}.

Because of these equilibrium pricing strategies, firm k's profits increase with μ_{kA} and μ_{kB} while firm j's profits decrease, as shown in figure 3. Intuitively, firm k's ability to build brand loyalty as it shuttles consumers through different vehicles over their lifecycle gives it a strong competitive advantage over firm j.

The fact that firm k's profits increase with μ_{kA} and μ_{kB} additionally implies that it has a unilateral incentive to encourage brand loyalty amongst its customers. Firm j also has this incentive as its profits are lower in this case than in the case from section 6.2 when it has loyal consumers as well. Thus, in an equilibrium in which firms have some control over the extent to which consumers develop brand preferences, firms will encourage brand preferences even though the resulting equilibrium leaves them worse off. That is, absent collusion, competitive forces push firms to compete not only in prices but also in the degree of loyalty of their consumers, further reducing their profits. This outcome has parallels in the literature on advertising in oligopoly markets, in which the equilibrium level of advertising

[21] For example, see two recent television advertisements for the Toyota Camry and Chevy Silverado at http://www.youtube.com/watch?v=46pmd-qq_6o and http://www.youtube.com/watch?v=Mrl-mm-7WM8.

Figure 3: Prices and profits when only firm k has loyal consumers

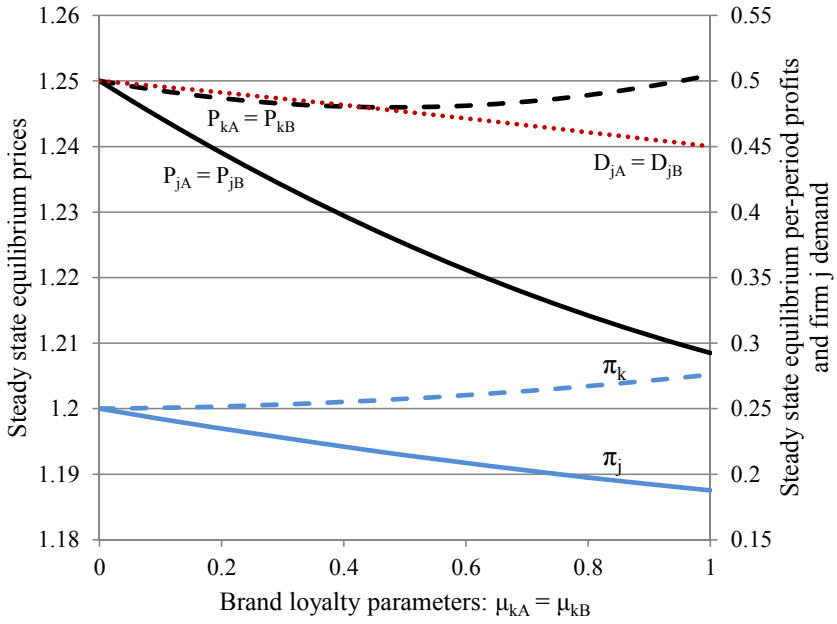

Note: Steady state equilibrium prices shown are from the model described in sections 6.1 and 6.3 in which $\delta = 0.9$, $V = 1$, $\alpha = 8$, $c = 1$, and brand preferences for firm j (μ_{jA} and μ_{jB}) are set to zero. For firm k, intergenerational brand preference parameter μ_{kA} is set to equal the within-household preference parameter μ_{jB}.

can exceed that which would maximize industry profits (see, for instance, Dixit and Norman (1978) and Grossman and Shapiro (1984)).

7 Ongoing Research

We have found evidence that brand preferences exist for vehicles and that these preferences are correlated across generations. We have also suggested that this intergenerational transmission of brand preference could have profound impacts on the structure of the automotive industry. We are currently pursuing supplementary results that reinforce and extend our findings in several directions. First, we have located a data set that details the openings and closings of auto dealers, by brand and year. We plan to use these openings and closings as instruments in order to add an instrumental variables analysis and improve our ability to distinguish among causal mechanisms.

In addition, we have thus far focused exclusively on intergenerational preference trans-

mission from parents to children, but the PSID allows us to also look at correlations across siblings and at preference formation that flows from children to parents. We are in the process of analyzing these relationships, and we anticipate that this additional analysis will aid us in testing auxiliary predictions of our model and in testing between transmission mechanisms.

We are also building a joint discrete choice model of *family* vehicle demand. This problem is challenging because joint modeling rapidly expands the choice set and creates computational barriers. A discrete choice model will require us to make stronger assumptions about the form of intrafamily brand preferences, but we believe that this exercise will complement our reduced-form analysis and enable us to perform simulations that show how vehicle choice would respond to brand quality shocks or various policy interventions.

8 Conclusion

Our analysis of the PSID shows that brand persistence in the automobile sector is strong. Within household persistence makes an individual more than twice as likely to purchase a vehicle of a particular brand if they have already purchased that brand. The strength of the correlation between children and their parents is nearly as strong and remains economically large and precisely estimated through various specifications. Together, our results paint a picture of strong brand preference that filters throughout family networks.

The strength of brand preferences for automobiles may be surprising, as compared to brand persistence in other goods. After housing, automobiles are the largest purchases that most households make, which means that a reliance on brand as a convenient heuristic for simplifying choice seem an unlikely explanation for brand persistence. Our results suggest that both endogenous taste formation (perhaps based on nostalgia) and information sharing are at work in explaining the brand loyalty we observe in the data.

The strength of brand preferences has important implications for industry dynamics. First, it suggests that market shares will evolve slowly, that firms that have experienced relative declines in quality may remain in the market for a long time, and that industry shocks will be mitigated through brand loyalty. For example, if fuel economy regulations are tightened significantly, or if retail gasoline prices rise sharply, the market may shift towards firms that have a comparative advantage in fuel economy. However, this shift will be mitigated by loyalty towards the incumbent firms. Second, using a model that captures firms' dynamic pricing incentives, we show that brand preferences are likely to depress automakers' prices and profits in equilibrium. In particular, intergenerational brand preference transmission forces firms to reduce prices of upscale vehicles intended for older

consumers. This outcome occurs despite firms' strong unilateral incentive to increase profits by encouraging their customers to stay loyal to their brand.

We conclude by noting that brand preferences are likely relevant to the broader question of why most major automakers carry a wide range of vertically differentiated products. There are, of course, many reasons for this strategy apart from its relationship to brand preferences: for instance, production economies of scope and the value of covering the wide range of consumers' attribute preferences are surely important. That said, it seems likely that giving consumers an opportunity to stay within a brand over the lifecycle, thereby allowing them to develop a brand preference, can substantially enhance the value of carrying a broadly differentiated vehicle fleet. It is intuitive to expect that brand preferences give producers an incentive to develop entry-level offerings that will "lead" consumers to their profitable upscale goods. The intergenerational link further implies that discount automakers will have an incentive to develop upscale product lines in order to "lead" future generations to their entry product. This link may, for instance, help explain the slow growth of firms such as Toyota and Honda in the 1980s and 1990s, and Hyundai today, which offer entry-level vehicles of high quality at low prices.

The product line question also harkens back to the debate over product strategies that shaped the initial competition between Ford and GM. Early in the twentieth century, Ford's strategy was to create a single vehicle that was affordable to all, driving down costs through economies of scale. There was no interest in vertical product lines, and Henry Ford famously quipped that "people can have a Model T in any color—so long as it is black." On the other hand, GM's strategic plan was to build a variety of cars that fit people at different life stages and income levels, embodied by the famous quote from Alfred Sloan that GM would sell "a car for every purse and purpose." Brand preference transmission within and across generations may be helpful in explaining why GM's strategy succeeded and why Ford ultimately deviated from its initial strategy and became a full-line automaker. Future research into brand preferences and firms' product line choices would be valuable in shedding light onto both these historic developments and more recent industry dynamics.

References

Atkin, David, "Trade, Tastes and Nutrition in India," August 2011. Manuscript: Yale University.

Berry, Steven, James Levinsohn, and Ariel Pakes, "Automobile Prices in Market Equilibrium," *Econometrica*, July 1995, *64* (4), 841–890.

Birch, Leann L., "Development of Food Preferences," *Annual Review of Nutrition*, 1999, *19*, 41–62.

Black, Sandra E. and Paul J. Devereux, "Recent Developments in Intergenerational Mobility," in "Handbook of Labor Economics," Vol. 4b, Elsevier, 2011, pp. 1487–1541.

Bronnenberg, Bart J., Jean-Pierre Dubé, and Matthew Gentzkow, "The Evolution of Brand Preferences: Evidence from Consumer Migration," *American Economic Review*, Forthcoming.

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

Cabral, Luís, "Small Switching Costs Lead to Lower Prices," *Journal of Marketing Research*, 2009, *46*, 449–51.

Dixit, Avinash and Victor Norman, "Advertising and Welfare," *Bell Journal of Economics*, Spring 1978, *9* (1), 1–17.

Doganoglu, Toker, "Switching Costs, Experience Goods, and Dynamic Price Competition," *Quantitative Marketing and Economics*, 2010, *8*, 167–205.

Dubé, Jean-Pierre, Gunter J. Hitsch, and Peter E. Rossi, "Do Switching Costs Make Markets Less Competitive?," *Journal of Marketing Research*, 2009, *46*, 435–45.

_ , _ , and _ , "State Dependence and Alternative Explanations for Consumer Inertia," *Rand Journal of Economics*, 2010, *41* (3), 417–45.

Fernández, Raquel, Alessandra Fogli, and Claudia Olivetti, "Mothers and Sons: Preference Formation and Female Labor Force Dynamics," *Quarterly Journal of Economics*, 2004, *119* (4), 1249–1299.

Grossman, Gene M. and Carl Shapiro, "Informative Advertising with Differentiated Products," *Review of Economic Studies*, Jan 1984, *51* (1), 63–81.

Jacoby, Jacob and Robert W. Chestnut, *Brand Loyalty: Measurement and Management*, New York: John Wiley and Sons, 1978.

Klemperer, Paul, "Markets with Consumer Switching Costs," *Quarterly Journal of Economics*, May 1987, pp. 375–394.

Langer, Ashley, "Demographic Preferences and Price Discrimination in New Vehicle Sales," Fall 2011. Manuscript: University of Michigan.

Logan, Trevon D. and Paul W. Rhode, "Moveable Feasts: A New Approach to Endogenizing Tastes," August 2010. Manuscript: The Ohio State University.

Luttmer, Erzo F.P. and Monica Singhal, "Culture, Context, and the Taste for Redistribution," *American Economic Journal: Economic Policy*, Forthcoming.

Mannering, Fred and Clifford Winston, "A Dynamic Empirical Analysis of Household Vehicle Ownership and Utilization," *Rand Journal of Economics*, Summer 1985, *16* (2), 215–236.

_ **and** _ , "Brand Loyalty and the Decline of American Automobile Firms," *Brookings Papers on Economic Activity. Microeconomics*, 1991, pp. 67–114.

Manski, Charles F., "Identification of Endogenous Social Effects: The Reflection Problem," *The Review of Economic Studies*, 1993, *60* (3), 531–542.

_ , "Economic Analysis of Social Interactions," *Journal of Economic Perspectives*, Summer 2000, *14* (3), 115–136.

Solon, Gary, "Intergenerational Income Mobility in the United States," *American Economic Review*, June 1992, *82* (3), 393–408.

_ , "Intergenerational mobility in the labor market," in O. Ashenfelter and D. Card, eds., *Handbook of Labor Economics*, Vol. 3 of *Handbook of Labor Economics*, Elsevier, October 1999, chapter 29, pp. 1761–1800.

Somaini, Paulo and Liran Einav, "A Model of Market Power in Customer Markets," September 2011. Manuscript: Stanford University.

Train, Kenneth E. and Clifford Winston, "Vehicle Choice Behavior and the Declining Market Share of U.S. Automakers," *International Economic Review*, 2007, *48* (4), 1469–1496.