ECONSTOR Make Your Publications Visible.

A Service of

ZBW

Leibniz-Informationszentrum Wirtschaft Leibniz Information Centre for Economics

Anderson, Soren T.; Kellogg, Ryan; Langer, Ashley; Sallee, James M.

Working Paper

The Intergenerational Transmission of Automobile Brand Preferences: Empirical Evidence and Implications for Firm Strategy

Working Paper, No. 247

Provided in Cooperation with:

George J. Stigler Center for the Study of the Economy and the State, The University of Chicago Booth School of Business

Suggested Citation: Anderson, Soren T.; Kellogg, Ryan; Langer, Ashley; Sallee, James M. (2013) : The Intergenerational Transmission of Automobile Brand Preferences: Empirical Evidence and Implications for Firm Strategy, Working Paper, No. 247, The University of Chicago, George J. Stigler Center for the Study of the Economy and the State, Chicago, IL

This Version is available at: https://hdl.handle.net/10419/262649

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



WWW.ECONSTOR.EU

Working Paper No. 247

"The Intergenerational Transmission of Automobile Brand Preferences: Empirical Evidence and Implications for Firm Strategy"

SOREN T. ANDERSON

RYAN KELLOGG

ASHLEY LANGER

JAMES M. SALLEE

George J. Stigler Center for the Study of the Economy and the State The University of Chicago

The Intergenerational Transmission of Automobile Brand Preferences: Empirical Evidence and Implications for Firm Strategy^{*}

Soren T. Anderson Michigan State University Ryan Kellogg University of Michigan

Ashley Langer University of Michigan James M. Sallee[†] University of Chicago

May 6, 2013

Abstract

Do children inherit product preferences from their parents? If so, how might such inheritance influence firm strategy and market outcomes? In this paper, we document a strong correlation in the brand of automobile chosen by parents and their adult children. Our preferred estimates imply that children are 46% more likely to choose an automobile brand if their parents also own that brand than if they do not. Correlations in intrafamily brand choice could represent a causal transmission of brand preference, or it could be due to correlated family characteristics that determine brand choice. We present a variety of empirical specifications which lend support to the causal interpretation. We then develop a model of Bertrand competition among producers in the presence of brand switching costs that are transmitted across generations to demonstrate how a causal link might influence firm strategy. We find that intergenerational transmission of brand preferences lowers equilibrium prices—particularly for vehicles targeted at parents—and firm profits because preference transmission limits the ability of firms to capitalize on existing brand loyalty by increasing the importance of investing in future loyalty.

^{*}We thank Evan Hernstadt and Katie Lim for excellent research assistance and Gary Solon, Clifford Winston and seminar participants at Arizona, Chicago, Maryland, Michigan, Stanford, UC Berkeley, UC Davis, UC San Diego, Yale, Washington University, Wayne State and Western Michigan for helpful comments.

[†]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

Economic models generally treat consumer preferences as exogenous and fixed. However, recent research on tastes for food (Birch 1999; Logan and Rhode 2010; Atkin 2011), female labor supply (Fernández, Fogli and Olivetti 2004), packaged goods (Bronnenberg, Dubé and Gentzkow Forthcoming), and preferences for redistribution (Luttmer and Singhal 2011) suggests that tastes and preferences may form endogenously through prior behavior, consumption, and experience. Such endogeneity has the potential to transform our view of market behavior and public policy by creating a dynamic link between consumption today and welfare tomorrow. For example, long-lived firms considering their pricing strategy to-day may think ahead to how market share, and consequent brand loyalty, will affect demand among future consumers.

The focus in economics on models of stable exogenous preferences is at least partially pragmatic. Models that allow for endogenous preference formation present technical and empirical challenges, and it is difficult to identify factors that influence preferences in most settings. One influence that is both plausibly important and potentially manageable is the family, which has proven to be fertile ground in prior literature (Fernández et al. 2004; Logan and Rhode 2010; Atkin 2011). Families may transmit preferences across generations, not only through immutable genetic endowments, but also through environment and experience. Thus, research on the intergenerational transmission of preferences may provide a better understanding of endogenous preference formation more generally.

In this paper, we empirically investigate the intergenerational transmission of preferences for automobile brands (e.g., Ford) and use these results to motivate a model of firm competition. Automobiles are particularly interesting because they represent a large, infrequent consumer purchase and because the industry remains crucial to the U.S. economy. Brands are particularly interesting in that they represent a "soft" attribute for which we might expect preferences to be especially malleable. That is, it is difficult to see why a consumer would inherently prefer a Ford over a GM for purely exogenous reasons, such as genetics. Instead, it seems natural to expect such preferences to arise from experience and context.

Our investigation makes use of the Panel Study of Income Dynamics (PSID), which is unique in that it follows multiple households within the same family over time. In particular, it surveys adult children who grew up in a PSID household but have since left and formed their own households. In several recent waves, the PSID has included questions about automobile ownership. Using these data, we find strong correlations in automobile choices across generations within a family. Specifically, a child whose parent has recently purchased a given brand is 46% more likely to choose that same brand (a 6.5 percentage-point increase on a base of 14.3%) than a demographically similar child whose parent did not choose that brand. To the best of our knowledge, we are the first to document this correlation.

In this paper, we define brand preference somewhat broadly, as representing a situation in which a consumer prefers one automobile brand over another holding constant major vehicle attributes (such as size and performance) and consumer characteristics (including demographics and geography).¹ This definition allows intergenerational brand choice correlation to be driven either by cross-household state dependence or direct brand preference transmission. In cross-household state dependence, the brand choices of parents influence child preferences (and therefore choices) because children have contact with their parents' vehicles and develop tastes for minor design details or nostalgic childhood associations with a brand, or because parental ownership generates information about performance and reliability that is conveyed to children. In direct brand preference transmission, the brand preferences of parents influence child preferences independent of parental brand choice. For instance, parents might have a long-term brand preference that is instilled in children early in life, or the parents might learn about a brand through friends or advertisements and then convey this information to their child prior to purchasing the brand themselves. Separate identification of these mechanisms in our data is difficult, though we will present evidence suggesting that at least cross-household state dependence is at work.

Intergenerational correlation in brand choices could also arise from familial correlations of demographic or geographic factors that also determine brand choice. We attempt to insulate our estimates from these more mundane correlations in several ways. First, we demonstrate that controlling for the rich set of demographic factors available in the PSID has little effect on our estimated brand choice correlations. Second, we show that correlations remain strong when fine geographic controls, at the county or census tract level, are introduced nonparametrically. Finally, we repeat our analyses for a pair of firms (GM and Ford) that produce very similar models with common attributes, which limits the degree to which unobserved characteristics could be determining choice. That is, while there are many characteristics that might lead an individual to choose a BMW over a Ford, it is more difficult to imagine characteristics that might cause an individual to prefer a GM over a Ford, other than those that are included under our definition of brand preference. We find strong intergenerational correlations for Ford and GM buyers as well. Thus, even though we cannot definitively rule out the possibility that the observed brand choice correlations are driven by unobserved household characteristics, the robustness of our results across a range of specifications argues in favor of true cross-household state dependence or brand preference transmission.

After presenting our empirical evidence, we provide a heuristic model of the car market to

¹The distinction between attributes and brand preference can be blurry, as we discuss below.

show how the dependence of children's brand choices on the previous choices of their parents 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 more limited in their ability 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 likely lowers prices and profits because 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 to harvest this loyalty sufficiently to compensate for the initial lower prices.

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 (for example, advertising aggressively 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.

Our analysis relates to several existing literatures. First, as discussed above, our research relates to the nascent literature on endogenous preference formation. Much of that literature centers on preferences for food, which in contrast to automobiles is characterized by small and frequent repeat purchases.

Second, previous work in economics and marketing has studied brand loyalty in the automobile market (Mannering and Winston 1985, 1991; Train and Winston 2007). While these papers document within-household brand loyalty, the automobile literature has, to the

best of our knowledge, never considered the intergenerational dimension of brand preference that we document here.

Third, 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 et al. Forthcoming). As compared to the products studied in that literature, automobiles are much larger expenses, they are purchased less frequently, and the product offerings are 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 that brands play in purchasing an automobile.

Fourth, there is 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 related literatures that study the intergenerational transmission of IQ, occupation, 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 balance of the paper is structured as follows. In section 2 we provide a framework for interpreting interhousehold correlations in brand choices. We then describe our data in section 3, and we report our empirical results regarding the correlation in brand choice across generations in section 4. Section 5 lays out a simple theoretical framework that shows the implications of intergenerational brand preference for automakers' prices and profits in equilibrium. Section 6 concludes.

 $^{^{2}}$ See Solon (1992) for a seminal contribution and Solon (1999) and Black and Devereux (2011) for recent reviews.

2 Conceptual model of intergenerational vehicle choice

We begin by presenting a simple model of household vehicle choice that clarifies possible mechanisms by which choices may be correlated across families and the empirical challenges of separately identifying them. Consider a household i in family f that purchases vehicle j at time t. Let the utility that household i derives from this purchase be denoted by:

$$U_{ifjt} = f(D_{ift}, X_j; \beta) + \theta_{ifjt}, \tag{1}$$

where D_{ift} denotes a vector of observed and unobserved demographic and location-specific characteristics of household *i*, such as income, education, climate, and terrain. These characteristics interact with X_j , which denotes the attributes (including brand) of vehicle *j*, through function $f(\cdot)$ and parameter vector β . This interaction allows observable and unobservable characteristics of households and their locations to influence vehicle choice in a variety of ways. For example, rural households may tend to choose pickup trucks, wealthy households may tend to purchase large SUVs, pro-union households may tend to purchase U.S. brands, and households living close to a Ford dealership may tend to purchase Fords. Finally, θ_{ifjt} denotes a preference for vehicle *j* that is unrelated to demographic or locationspecific factors. We focus on influences from other family members as determinants of θ_{ifjt} , but other factors may exist, such as exposure to advertisements, prior driving experiences, idiosyncratic tastes (e.g., for a particular color or trim), or vehicle and gasoline market conditions at the time of purchase.

If parents' preferences and choices for vehicles have a direct impact on their children's preferences, this is expressed in our model as a correlation in θ_{ifjt} across households within families, which leads to correlation in vehicle choices. Cross-household correlation in θ_{ifjt} could stem directly from experience with a particular brand, whereby parental *choices* influence child choices. For example, if a child's parents purchased a string of GM vehicles, then that child may have nostalgic feelings for GM, a developed taste for the unique features of GM's design (e.g., the layout of the instrument panel or the feel of the seats), superior information about GM's performance and reliability, or simply a "comfort level" with the brand. Alternatively, the correlation in θ_{ifjt} could arise from transmission of preferences from one household to another in a way that is not mediated by brand choice itself. For instance, parents might be exposed to a positive review of GM and tell their children about what they read. This type of mechanism may lead ultimately to a correlation of parent and child choice, but it operates not through state dependence and experience but rather through information.

Correlations in households' vehicle choices may also arise through cross-household corre-

lations in D_{ift} . It is natural to expect such correlations to exist; Solon (1992), for example, documents strong intergenerational transmission of income. If households with high incomes are more likely to purchase SUVs and European luxury brands, then this correlation in income across generations will lead to correlations in vehicle choices across generations. Thus, a fundamental empirical challenge of our work is to identify vehicle choice correlations that arise from state dependence and preference transmission rather than those that arise from similarities in demographic and geographic characteristics. This separate identification is important because it is only the former set of channels that is relevant for automakers' strategies.

We use several approaches to separately identify "true" transmission of vehicle preferences across households within the same family. First, we leverage the wealth of demographic and location information within the PSID dataset to control directly for potential confounding factors. Despite being able to use a rich set of covariates (including census tract fixed effects), one might nonetheless be concerned that influential unobserved factors remain. For instance, if a family is unobservably pro-union, all of that family's households might have a preference for U.S. brands. Thus, we also investigate subsets of the data for which unobserved factors are less likely to be important, focusing in particular on choices between Ford and GM.

A more challenging empirical problem is the separate identification of the extent to which children's brand preferences are determined by their parent's choices or by their parent's preferences. This distinction is important because, while both mechanisms are potentially relevant for automakers' advertising, design, and product line strategies, only the choice dependence mechanism is relevant for automakers' pricing strategies. This identification problem is similar to one common in the marketing literature, in which one observes a series of brand choices by a single household and then tries to determine whether that household's choices are state-dependent or whether they simply reflect a preference for a particular brand (see, for example, Dubé, Hitsch and Rossi (2010)). Ideally, we would solve this problem using an instrumental variable, such as vehicle prices, that affect households' brand choices but not their preferences. However, our experimentation with this strategy has found it to be severely under-powered. Therefore, we adopt the alternative strategy of searching for systematic patterns in brand choice correlations that speak differentially to choice versus preference dependence. In particular, sub-section 4.3 studies whether the observed choice correlations are stronger when households have a relatively high level of exposure to the choices of their parents or other households in their family.

Finally, we note that the distinction between brands and attributes is blurry. It is tempting to define a brand as something that is independent of all vehicle attributes, as if, for instance, Ford and GM vehicles were identical apart from the logo stamped on the grill. In practice, vehicles of different brands will differ in "minor" features, including trim style, dashboard layouts, and perceived reliability, even for cars that share identical observable characteristics, such as size, power, and cargo space. We define a brand in a way that encompasses these "minor" characteristics so that a brand preference might be derived from, for example, a preference to have the dashboard controls laid out in a particular way. In contrast, when we speak of preferences for attributes, we refer specifically to major vehicle characteristics, such as class, horsepower, size, and fuel economy. We believe that making this distinction between attributes and brands, thusly defined, is useful because the transmission of preferences for brands has very different implications than the transmission of preferences for attributes. The former is primarily relevant for automakers' pricing, marketing, and product line strategies, while the latter is primarily relevant for public policies addressing the externalities of vehicle use.

3 Data

Our data on vehicle ownership come from the Panel Study of Income Dynamics (PSID). In 1968, the PSID surveyed a nationally representative sample of households, and since then it has asked them a battery of economic and demographic questions every year until 1997 and 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, 2007 and 2009. To the best of our knowledge, the PSID is unique in providing such information for families in the United States.

Our primary focus is on how parental vehicle brand choices influence the choices of their adult children. Accordingly, our baseline sample is limited to adult heads of household (or spouses) who purchase a car in the sample, for whom we can identify a parent who owned a vehicle prior to their child's vehicle purchase. We identify 4,006 unique adult children matched to 2,381 unique parents. The difference in the number of parents and children is due to the fact that there are many siblings in our sample.³

 $^{^{3}}$ In our regression analysis, we cluster standard errors on the original 1968 PSID family in all regressions to allow for correlated errors across relatives.

	Adult children	Parents	
Age (years)	35.8	59.5	
Years of education	13.4	12.4	
Annual family income (\$)	77,231	60,419	
Number of people in household	3.1	2.3	
Number of vehicles owned	2.3	2.1	
Number of unique individuals	4,006	2,381	
	Matched Pairs		
Number of child vehicle choices matched to parent choice	$13,\!535$		
Total number of parent to child vehicle matches	14,530		

Table 1: Variable means and sample sizes in PSID

Table 1 shows sample means for both children and parents in this sample. Adult children are on average 36 years old, whereas parents are 60. Adult children have higher household income, one year more education (13.4 as opposed to 12.4), and larger family sizes, which accords with the higher probability that they have children living with them, as compared to their parents. We observe 13,535 unique vehicle purchases by these 4,006 adult children. Excluded from this sample are vehicles that were received as gifts and vehicles that are likely to have been within-family cross-household sales.⁴ In cases where parents are separated, but both are present in the PSID and both have a prior vehicle purchase available, we match the child's vehicle choice with data from both parents.⁵ There are 995 such cases, which leaves us with 14,530 parent-child vehicle pairs in our main estimation sample.

4 Empirical evidence of intergenerational brand preference transmission

In this section, we develop and estimate a linear probability model (LPM) of the relationship between brand choices of children and the choices of their parents, as well as other covariates. We employ an LPM rather than a structural discrete choice model because it relies on fewer assumptions about the structure of the error terms and is more forgiving to the extensive geographic and time fixed effects that we will use in our estimation. This comes at the cost of not being able to interpret our coefficient estimates as parameters of a utility function.

To operationalize our brand choice data in the LPM framework, we first categorize all

⁴Specifically, we drop a child's vehicle purchase if the parent household owned the same make, model, and model-year in the previous survey wave and subsequently no longer owns the vehicle following the child's purchase.

⁵In these cases we weight each vehicle-parent pair by half so that choices that appear twice in our data are weighted equally to those that appear once.

vehicle choices as being one of seven "brands": 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 all choice probabilities lie in the 4%–33% range in the raw data.⁶

We build a linear probability model with multiple choice possibilities by stacking a set of binary linear probability models for each of our 7 brands. To motivate our approach, we first consider a linear probability model for a single brand. Our hypothesis is that parental ownership of a given brand will make a child more likely to choose that brand. For example, to test this hypothesis for Ford we could run a linear probability model where the dependent variable is coded as one if the child was observed to choose a Ford. In addition to controls for the child's demographics and parent's demographics, our regressor of interest would be a dummy variable for the parent's brand choice, which would be coded as 1 if the parent's most recent vehicle purchased prior to the child's purchase was a Ford. The one-brand estimation equation would be:

$$Ford_{ift} = \gamma \cdot Ford_{pft} + \mathbf{X_{ift}}'\beta_j + \mathbf{X_{pft}}'\delta_j + \alpha_t + \epsilon_{ift}, \tag{2}$$

where $Ford_{ift}$ is coded as 1 if the child's vehicle is a Ford, $Ford_{pft}$ is coded as 1 if the parent's vehicle is a Ford, α_t is a period-specific constant to capture Ford's average market share in each period, \mathbf{X}_{ift} are the child's observable demographic characteristics, and \mathbf{X}_{pft} are those of the parent. In this regression of "Ford against the field", all observations in our data are included.

Instead of running 7 separate one-brand linear probability models, we stack them and run pooled OLS. For each of our seven brands, we create an observation for each observed choice situation (each car purchase) where the dependent variable is coded as 1 if the brand is chosen and zero otherwise. We do this for all seven brands and stack the data, which generates a final data set that has seven times the number of observations as our original data set that included one observation per choice.⁷

⁶In 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 subsets of our regressions with 41 sub-brands instead of the 7 brands, interacting our control variables with all 41 sub-brands. In general, we find that both the overall brand and the sub-brand of the parent have a statistically significant correlation with the sub-brand chosen by the adult child.

⁷Our procedure can also be described as expanding each observed choice as follows. 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.

We interact all of our regressors with brand dummies, thereby allowing observed demographic and geographic factors to affect the choice probability of each brand differently. Thus, all covariates are denoted with a j subscript or interacted with a coefficient vector of j length. The one restriction that we impose is that the effect of the parent's past brand choice is common across brands.⁸ Thus, instead of seven dummy variables for parental brands, there is only one dummy variable coded as 1 when the parental choice matches the brand represented in that row of data for the child. For each choice situation, exactly 1 of 7 observations will have the parental dummy variable coded as 1.

This leads us to the following estimation equation:

$$b_{ifjt} = \gamma \cdot 1(b_{pfjt} = b_{ifjt}) + \mathbf{X_{ift}}'\beta_j + \mathbf{X_{pft}}'\delta_j + \alpha_{jt} + \epsilon_{ifjt}, \tag{3}$$

where the dependent variable, b_{ifjt} , is a dummy coded as 1 if child *i* of family *f* chose brand *j* in choice *t*. The independent variable of interest is a dummy variable that indicates whether the parent's most recent prior purchase is of that same brand $1(b_{pfjt} = b_{ifjt})$.⁹

Our hypothesis is that the γ coefficients will be positive—that is, children are more likely to purchase a given brand if their parents have purchased that brand in the recent past. We control for both child characteristics X_{ift} and parent's characteristics X_{pft} , which enter with brand-specific coefficient vectors, β_j and δ_j , that we estimate by interacting child and parent characteristics with brand dummies. Finally, we allow for brand-by-month of sample fixed effects, α_{jt} to capture overall market shares, leaving ϵ_{ifjt} as the error term.

This setup expands each observed brand choice into seven observations. To avoid this expansion of the data set unduly shrinking our standard error estimates, we cluster all standard error calculations at the level of the 1968 PSID family. This clustering accounts for the mechanical correlation in the residuals between the seven observations that represent a single choice, the correlation in the individual's choices across choice situations, and the correlation across siblings or cousins. We also weight each observation using PSID-provided sampling weights so that our regression results can be interpreted as being representative of

⁸By using the linear probability model, we do not impose a restriction that predicted values must be between zero and one, nor do we require that the sum of the predicted values across the seven brands must equal one for each choice situation. We have checked our predicted values for our baseline specifications, and we find that the vast majority of predicted values are between zero and one, and those that deviate are very small negative numbers. Similarly, the sum of the predicted values across the seven brands for each choice situation are closely distributed around 1.

⁹We have experimented with a variety of ways of characterizing the parent's choice given that parents may own multiple vehicles. We have used the share of parent's vehicles of a specific brand; we have done one-to-one vehicle matching across parents and children using the order in which vehicles are listed in the survey; and we have included up to three lags of the parent's choices simultaneously. In all cases, our qualitative results are quite similar. We prefer the dummy variable specification for ease of interpretation and because it accords with the causal model we have in mind.

the U.S. population.

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 standard 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 distinct phenomena.

4.1 Baseline results

We begin by showing simple correlations in order to demonstrate the strength of the intrafamily relationship and then demonstrate how the correlation is affected by various controls. We focus here and in section 4.2 on separating mechanisms of interest (choice and preference dependence) from choice correlation driven by demographic or geographic factors. In section 4.3 we then attempt to distinguish among choice and preference-dependence mechanisms.

Table 2 includes coefficient estimates of γ from equation 3, which regresses the brand chosen by the child on a dummy for whether or not the parent's most recent purchase is the same brand. The estimate in column 1, which includes only month of purchase-by-brand fixed effects that control for the overall share of each brand during each period, indicates that a child is 10 percentage points more likely to choose the brand that their parent's choose. There are seven brands in our choice set, so the probability that the average brand is selected is 0.143. Thus, our estimate implies that a child whose parents chose a particular brand is 70% more likely to choose that brand than another child whose parents chose differently.

This is a remarkably strong relationship, but it may reflect not only the intergenerational transmission of brand preference that we are interested in, but also familial correlations in demographic and location-specific factors that cause related households to demand similar attributes in vehicles, in turn causing a correlation in brand choice.

To address this issue, we introduce progressively richer controls in columns 2 through 7 of table 2 and examine how the coefficient estimates change. In column 2 we add demographic controls (including family income, an urban versus rural dummy, age, sex, education, number of kids in the household, and household size) for the child's household as well as the parent's household.¹⁰ The addition of these controls lowers the estimated coefficient from 0.101 to 0.093. While the modest impact of these controls on the estimated coefficient is encouraging,

¹⁰Each of these characteristics is interacted with a dummy for each brand, which is a flexible analog to the traditional approach in the automobile demand literature of interacting vehicle attributes with buyer characteristics.

	Dependent Variable: Child's Brand						
VARIABLES	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Parent's brand $=$ child's brand	0.101	0.093	0.083	0.081	0.065	0.070	0.056
	(0.008)	(0.008)	(0.008)	(0.008)	(0.008)	(0.014)	(0.014)
Month of purchase fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Child's demographics	No	Yes	Yes	Yes	Yes	Yes	Yes
Parent's demographics	No	Yes	Yes	Yes	Yes	Yes	Yes
Child's state fixed effects	No	No	Yes	Yes	Yes	Yes	Yes
Parent's state fixed effects	No	No	No	Yes	Yes	Yes	Yes
Child's county fixed effects	No	No	No	No	Yes	Yes	Yes
Child's census tract fixed effects	No	No	No	No	No	No	Yes
Number of choices	$14,\!530$	$14,\!530$	$14,\!530$	$14,\!530$	$14,\!530$	6,017	6,017
R^2	0.093	0.107	0.122	0.128	0.229	0.293	0.412

Table 2: Correlations between child brand choice and parental brand choice

Standard errors clustered by 1968 PSID family are in parentheses. Each column is a linear probability model where each individual-year-vehicle choice enters the data 7 times, once for each brand (GM, Ford, Chrysler, Toyota, Honda, Other Asian, and European). Child's and parent's demographics include age, education, income, urban dummy, gender, number of children in household, and family size. All control variables are interacted with 7 dummies, one for each brand. Columns 6 and 7 limit the sample to households living in census tracts that contain more than one PSID family.

the regression's R^2 also changes only slightly, from 0.093 to 0.107. This combination of results raises the possibility that, following the logic of Altonji, Elder and Taber (2005), if unobserved demographic factors that affect children's brand choices are as correlated with parents' choices as the observed demographic factors, then these unobservables may explain much of the observed brand choice correlation. We further address this issue in section 4.2 in which we focus on two brands—Ford and GM—that have very similar attributes and are therefore unlikely to be affected by demographic factors.

The remaining columns of table 2 address common geographic factors that might lead to choice correlation. Column 3 adds state-by-brand fixed effects, which control for differences in market shares and location-specific factors that vary by state. Column 4 adds analogous fixed effects for parents. These fixed effects cause the coefficient to fall modestly, from 0.093 to 0.081.

Location specific factors like dealer location, local prices, weather and terrain may vary significantly within some states. In column 5 we add county-by-brand fixed effects for the child's county of residence, which is intuitively a small enough geographic area to control for most omitted factors that we have in mind. County fixed effects lower the point estimate to .065, which is approximately two-thirds the size of the raw correlation and still highly significant, both statistically and economically. The coefficient estimate of .065 in column 5 implies that parental ownership boosts the conditional probability that a child buys a given

brand by 46%.

We interpret the difference across columns in the estimated effects as evidence that some location-specific factors are both important in determining brand choice and are correlated across family members (who tend to live in similar places). Weather, terrain, urbanization, and culture are important determinants of the demand for attributes, which are different on average across brands. Even conditional on demand, there may be a different availability of brands across geographic areas due to the location of dealerships. In our view, most of these differences should be captured by county-level fixed effects, so column 5 is our preferred specification.

The original PSID sample design drew stratified samples from particular geographic areas. The legacy of that original sample design is that PSID households are still more geographically clustered than would be the case for a random sample of households. As a result, despite having a modest sample size, we have enough distinct families that live in the same census tract (a unit of approximately 2,500 to 8,000 people) to include census tract fixed effects on the subsample of households that live in tracts with multiple families in our sample. Column 6 re-estimates the model with county level fixed effects on this subsample, and column 7 estimates a model with census tract fixed effects.

Census tract fixed effects are fine enough to control for factors such as local repair shops, mechanics, and other network effects. Moreover, they are fine enough to control for correlations in tastes or characteristics of people who choose to locate in the same neighborhood with a city. Even when including these effects, our coefficient of interest is economically and statistically significant. The census tract fixed effects are powerful, increasing the R^2 from 0.293 to 0.412, but they only lower the point estimate modestly between columns 6 and 7, so that the coefficients are comfortably within each other's confidence intervals.

Our baseline regressions focus on parental choices determining child choices, which we believe to be the strongest intrafamily channel of brand preference transmission. We can, however, configure our data to examine the relationship that prior purchases by any family member have on subsequent choices by their relatives. To do so, we take every vehicle choice observed in the data and match it to the most recent purchase made by all family members (including parents, children, siblings, cousins, etc.). We then include all of these bilateral relationships in one regression, down-weighting vehicles that are matched to multiple family members' vehicles so that they have equal influence on the estimate as those that have only one match. This alternative construction expands our sample size considerably and delivers more precise, but modestly smaller effects. For example, the all family matched analog of column 1 from table 2 produces a coefficient (standard error) estimate of 0.071 (0.005), and the county fixed effects analog to column 5 produces estimates of 0.056 (0.005). These results

are consistent with our intuition that parent to child influences are particularly strong, but it also suggests that broader family network effects have influence.

4.2 Estimates limited to similar brands

The principal concern with our baseline regressions is that demographic or location-specific characteristics of children and parents will be correlated and that these characteristics drive demand for vehicle attributes that are correlated with brand. While we believe that controlling for county—and ultimately census tract—fixed effects adequately addresses locationspecific confounders, the possibility remains that our estimates are contaminated by demographic confounders, even after controlling for observables. For example, individuals who work in construction jobs 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.

Here, we address this issue by isolating the choice set to two brands that are very similar: Ford and GM. Ford and GM are both full-line, U.S.-based automakers that compete directly in every vehicle segment.¹¹ 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.

Table 3 repeats the specifications in table 2 for a subset of choices limited to Ford and GM. Specifically, we keep all instances in which a child chose either a Ford or GM, which accounts for around 55% of our original sample.¹² As in the full sample, the results are all positive, statistically significant, and economically large. This corroborates our baseline results and casts doubt on the possibility that the correlation in brand choice across households is due entirely to demographic confounders. In particular, the estimated coefficient actually increases slightly going from column 1 to 2 as the demographic controls are added. If our "headline" demographic variables are not correlated with the choice between Ford and GM, it seems unlikely that unobserved factors would be correlated either.

The magnitudes of the coefficients are somewhat larger in the subsample in table 3 than

¹¹Popular perception holds that Ford and GM are similar brands. We have confirmed this empirically using a measure of "distance" between brands that borrows from Langer and Miller (Forthcoming), who calculate the distance between pairs of vehicles in attribute space based on vehicle segment, price, number of passengers, wheelbase, fuel economy and horsepower for GM, Ford, Toyota and Chrysler. Using their metric, we have confirmed that Ford and GM vehicles are on average substantially closer to each other than they are to Toyota and Chrysler.

¹²We do not restrict the sample based on whether the parents chose Ford or GM. We do, however, add an additional control variable for whether the parents' choice was one of these two brands. This control helps ensure that the sum of the child's choice probabilities for the two brands is close to one.

	Dependent Variable: Child's Brand						
VARIABLES	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Parent's brand $=$ child's brand	0.143	0.147	0.124	0.118	0.098	0.082	0.089
	(0.021)	(0.021)	(0.019)	(0.020)	(0.020)	(0.033)	(0.037)
Month of purchase fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Child's demographics	No	Yes	Yes	Yes	Yes	Yes	Yes
Parent's demographics	No	Yes	Yes	Yes	Yes	Yes	Yes
Child's state fixed effects	No	No	Yes	Yes	Yes	Yes	Yes
Parent's state fixed effects	No	No	No	Yes	Yes	Yes	Yes
Child's county fixed effects	No	No	No	No	Yes	Yes	Yes
Child's census tract fixed effects	No	No	No	No	No	No	Yes
Number of choices	7,983	$7,\!983$	7,983	$7,\!983$	7,983	$3,\!117$	3,117
R^2	0.079	0.083	0.106	0.121	0.279	0.325	0.454

 Table 3: Correlations between new vehicle brand and parental brand choice among those owning a Ford or GM

Standard errors clustered by 1968 PSID family are in parentheses. Sample is limited to the cases where the child chose Ford or GM. Each column is a linear probability model where each individualyear-vehicle choice enters the data 7 times, once for each brand (GM, Ford, Chrysler, Toyota, Honda, Other Asian, and European). Child's and parent's demographics include age, education, income, urban dummy, gender, number of children in household, and family size. All control variables are interacted with 7 dummies, one for each brand. Columns 6 and 7 limit the sample to households living in census tracts that contain more than one PSID family.

in the full sample in table 2. Nonetheless, brand preference correlation has a slightly *smaller* percentage effect on choice probabilities in the restricted sample because the baseline choice probabilities are higher. In the full sample, the market share for Ford is 22%, and the share is 33% for GM. In the subsample, the corresponding figures are 40% for Ford and 60% for GM. Thus, the coefficient of 0.098 in column 5 of table 3 (the county fixed effect specification in the restricted sample) implies that a child whose parent's most recent prior purchase was the same brand boosts the probability of purchase by 25% for Ford and 16% for GM, whereas the analogous coefficient in table 2 (the full sample) represents a 30% effect for Ford and 20% for GM. This difference is intuitive given that Ford and GM are generally close substitutes.

Toyota and Honda are also similar brands. They both produce a full range of sedans and fuel efficient SUVs, though Honda produces only a limited set of pickup trucks. Table 4 shows results from the same set of specifications for the sample of observations limited to children who purchased either a Honda or a Toyota, excluding all pickup trucks. The estimated effects in this subsample are even larger; the county fixed effects specification in column 5 (our preferred specification) indicates that having a parent who owns a Honda or Toyota increases the probability that a child chooses that brand by 53%.¹³ Honda and Toyota have a smaller market share than Ford and GM, which leaves us with a smaller sample

 $^{^{13}\}mathrm{A}$ 26.5 percentage point increase on an average market share of 50% is a 53% increase.

	Dependent Variable: Child's Brand						
VARIABLES	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Parent's brand $=$ child's brand	0.239	0.234	0.264	0.285	0.265	0.229	0.432
	(0.067)	(0.067)	(0.065)	(0.065)	(0.077)	(0.503)	(0.452)
Month of purchase fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Child's demographics	No	Yes	Yes	Yes	Yes	Yes	Yes
Parent's demographics	No	Yes	Yes	Yes	Yes	Yes	Yes
Child's state fixed effects	No	No	Yes	Yes	Yes	Yes	Yes
Parent's state fixed effects	No	No	No	Yes	Yes	Yes	Yes
Child's county fixed effects	No	No	No	No	Yes	Yes	Yes
Child's census tract fixed effects	No	No	No	No	No	No	Yes
Number of choices	1,899	$1,\!899$	$1,\!899$	$1,\!899$	$1,\!899$	229	229
R^2	0.132	0.146	0.180	0.224	0.505	0.822	0.933

 Table 4: Correlations between new vehicle brand and parental brand choice among those owning a Honda or Toyota

Standard errors clustered by 1968 PSID family are in parentheses. Sample is limited to the cases where the child chose Honda or Toyota. Each column is a linear probability model where each individual-year-vehicle choice enters the data 7 times, once for each brand (GM, Ford, Chrysler, Toyota, Honda, Other Asian, and European). Child's and parent's demographics include age, education, income, urban dummy, gender, number of children in household, and family size. All control variables are interacted with 7 dummies, one for each brand. Columns 6 and 7 limit the sample to households living in census tracts that contain more than one PSID family.

size and larger standard errors. Nevertheless, our estimates are statistically significant at any conventional level except when we limit the sample to those who live in a census tract common to another PSID family in our sample (columns 6 and 7), at which point we lose power.

4.3 Tests for cross-household state dependence

In section 2, we discussed three potential ways in which children's and parents' choices could be correlated. Sections 4.1 and 4.2 focused on distinguishing the effects of demographics and geography from our two mechanisms of interest: cross-household state dependence and cross-household direct preference transmission. Here, we explore the extent to which we can separate these two mechanisms from each other using two tests for the presence of state dependence. First, we attempt to identify cases in which a household had a positive versus negative experience with a particular vehicle and then assess whether other family members are less likely to buy that brand if the experience was negative. Second, we study whether or not children are more strongly influenced by vehicles that their parents owned while they still lived with their parents, under the presumption that children were more likely to have direct exposure to such vehicles. First, in terms of customer satisfaction, while we do not have direct measures of how satisfied each household is with its vehicle, we can proxy for satisfaction by measuring the duration of ownership. When a household quickly sells a vehicle that it recently purchased, we suspect that this is a sign of dissatisfaction, and when a household owns a vehicle for a very long time, we suspect that this implies a stronger positive attachment. Thus, we hypothesize that, in the presence of cross-household state dependence, when a parent owns a vehicle for a longer time period, the effect on the child's choice should be stronger.¹⁴ We test this hypothesis in our data by interacting our parental brand dummy variable with either a continuous measure of how long the parent owned the vehicle (measured in years) or a discrete variable indicating that the parent owned the vehicle for fewer than three years.¹⁵

Table 5 reports results from our preferred specification with county fixed effects with these added variables, but it does not use our baseline sample. Our baseline sample lacks statistical power to estimate both the main effect and the interactions, so we present here estimates from the expanded sample that includes all bilateral relationships among relatives that are available in the data, which we previewed above.¹⁶ This construction considerably increases our sample size by matching all family members (parents, siblings, cousins, etc.) that are part of the same 1968 PSID family. Vehicles that have multiple matches are included multiple times, but weighted so that they have the same influence on the regression as a vehicle that matches only once. In table 5, columns 1 and 2 include all brands, and columns 3 and 4 restrict the sample to observations for which the child purchased a Ford or GM.

Results are consistent with cars that are held longer having a greater positive impact on the future choices of family members. For regressions with all brands or just the Ford and GM sample, the coefficient on the interaction variable is positive and statistically significant, indicating that vehicles held longer have larger effects on subsequent family member choice. Similarly, the dummy variable specification in columns 2 and 4 indicates that vehicles owned for fewer than three years has a smaller effect on subsequent brand choice of children than other vehicles, though the estimated coefficient is imprecise for the Ford and GM only sample. In the all brands regression, the interaction suggests that a vehicle owned for fewer than three

¹⁴This effect could also arise from the fact that the child may have had more time to experience the parents' vehicle. This mechanism would also be consistent with state dependence.

¹⁵We do not observe the length of ownership for some vehicles because they are still held at the end of our sample in 2009. Still other vehicles are leased, and ownership length for these vehicles is therefore not a choice variable. For all vehicles for which ownership length is unknown or the vehicle is leased, we set ownership length in the regression to be equal to zero and code a "missing ownership length" dummy variable in the regression as a one. The average length of ownership for those vehicles in which ownership length is observed is 4.5 years.

¹⁶Results from our baseline sample are qualitatively the same, in that they have the anticipated sign and the magnitudes are similar to those in table 5, but the estimates are generally statistically insignificantly different from zero. Full results are available upon request.

	Dependent Variable: Household's Brand				
	All b	rands	Ford and GM only		
VARIABLES	(1)	(2)	(3)	(4)	
Relative's brand = household's brand	0.002	0.033	-0.057	0.050	
	(0.009)	(0.005)	(0.033)	(0.013)	
(Relative's brand = household's brand)	0.017		0.060		
$x \log(years owned)$	(0.005)		(0.020)		
(Relative's brand = household's brand)		-0.017		-0.045	
x years owned < 3		(0.007)		(0.027)	
Month of purchase fixed effects	Yes	Yes	Yes	Yes	
Child's demographics	Yes	Yes	Yes	Yes	
Parent's demographics	Yes	Yes	Yes	Yes	
Child's state fixed effects	Yes	Yes	Yes	Yes	
Parent's state fixed effects	Yes	Yes	Yes	Yes	
Child's county fixed effects	No	No	No	No	
Child's census tract fixed effects	No	No	No	No	
Number of choices	105,173	105,173	58,744	60,254	
R^2	0.210	0.210	0.233	0.230	

 Table 5: Correlations between child brand choice and relative brand choice interacted with relative's length of ownership

Standard errors clustered by 1968 PSID family are in parentheses. Each column is a linear probability model where each individual-year-vehicle choice enters once for each brand. Child's and relative's demographics include age, education, income, urban dummy, gender, number of children in household, and family size. All control variables are interacted with dummies for each brand. All regressions also include a dummy variable for cases where the length of ownership is not known because of censored or missing data.

years has half of the effect on subsequent family choices as does a vehicle owned longer.

Out second test for state dependence examines whether children are particularly likely to choose the same brand as their parents if they were still living with their parents when their parents purchased that brand. We observe a subset of our adult children living both in their parents' home and then subsequently in their own household during our sample period. For these individuals, we can determine whether or not vehicles that their parents purchased while they still lived in their household (and thus with which children were more likely to have direct experience) have a particularly strong influence and their future choices. We do so by including interactions indicating whether or not children lived at home when a vehicle was purchased by their parents.

All of our regressions above correlated an adult child's vehicle choice with the most recent brand choice of their parents. For this analysis, we are interested in the effect of vehicles purchased longer ago, so we reconstruct our data by matching each adult child vehicle purchase to all past purchases of parents that we observe in our data. We further

	Den en deut Verichle, Obildie Deen d					
	Dependent Variable: Child's Brand					
	A 11 1	1				
	All b	rands	Ford and GM only			
VARIABLES	(1) (2)		(3)	(4)		
Relative's brand $=$ child's brand	0.054	0.019	0.073	-0.026		
	(0.007)	(0.021)	(0.020)	(0.053)		
(Relative's brand = child's brand)	0.007	0.017	0.037	0.072		
x child at home	(0.010)	(0.012)	(0.030)	(0.034)		
(Relative's brand = child's brand)	-1.44e-05	3.69e-05	-5.75e-05	7.74e-05		
x months since purchase	(6.07e-05)	(7.06e-05)	(0.000170)	(0.000187)		
(Relative's brand = child's brand)		0.0010		0.0029		
x child's age at purchase		(0.0006)		(0.0015)		
Month of purchase fixed effects	Yes	Yes	Yes	Yes		
Child's demographics	Yes	Yes	Yes	Yes		
Parent's demographics	Yes	Yes	Yes	Yes		
Child's state fixed effects	Yes	Yes	Yes	Yes		
Parent's state fixed effects	Yes	Yes	Yes	Yes		
Child's county fixed effects	No	No	No	No		
Child's census tract fixed effects	No	No	No	No		
Number of choice pairs	44,919	44,919	24,296	24,296		
R^2	0.230	0.230	0.277	0.277		

 Table 6: Correlations between child brand choice and parental brand choice interacted with whether child was living with parents

Standard errors clustered by 1968 PSID family are in parentheses. Each column is a linear probability model where each individual-year-vehicle choice enters once for each brand. Child's and parent's demographics include age, education, income, urban dummy, gender, number of children in household, and family size. All control variables are interacted with dummies for each brand.

hypothesize that parental vehicles owned longer ago will have weaker effects, irrespective of whether the child was still in the house, so we control for the length of time between the parents' purchase and the child's. Finally, we also control for the child's age at the time of the parent's purchase, as we hypothesize that older children will be more attuned to vehicle holdings than younger children.

Table 6 reports linear probability model regressions that add these additional regressors in the expanded set of child-parent matched vehicles. We use our preferred specification with county fixed effects in all regressions. Our focus is on the interaction between the brand choices of parents and whether or not the child lived at home when the vehicle was purchased (the second variable in table 6). The estimated coefficient on the interaction term is positive in all specifications and is statistically significant in column 4 (the the Ford and GM subsample that includes a control for the child's age). These specifications also find that the brand choice correlation tends to be stronger for older children.

Neither the length of ownership specification nor the childhood experience specification

definitively identify the specific causal channel of brand preference transmission. However, these results do suggest (albeit with imprecision) that experience with a family member's vehicle is an important determinant of one's own vehicle brand choice. In section 5 below, we explore the implications of this cross-household choice dependence for automakers' pricing and marketing strategies.

5 Theoretical 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) by 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.¹⁷ 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.¹⁸

¹⁷Dubé et al. (2009) and Dubé et al. (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 price and other attributes.

 $^{^{18}}$ 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.

5.1 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 commonly uses the term "switching costs" to refer to a reduction in utility experienced by a consumer who switches from one brand to another in different periods of a model, which is identical to how we operationalize what we here call brand preference.¹⁹ Our approach is closest in spirit to that of Dubé, Hitsch and Rossi (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 differentiated 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 Aconsumer (offspring) upon death. A key feature of the model is that the type A and type Bconsumers 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 by households across vehicle types.²⁰ Still, in a survey of over 22,000 consumers by a market research firm described in

¹⁹Some 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.

²⁰Allowing 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

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. Similarly, only 5% of consumers who say they purchased a Buick are under the age of 40. 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 or she was a type A last period. The parameter μ_B denotes the strength of within-consumer persistence of brand preferences. The utility from purchasing the other brand's vehicle kB is given similarly. All type B households purchase exactly one vehicle, and there is no outside good.

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 the parents were type *B* last period. μ_A denotes the strength of intergenerational brand preferences. We 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*.

We assume that type A consumers are not forward-looking when deciding whether to purchase vehicle jA or kA; we will study the implications of relaxing this assumption in future work.²¹ 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

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.

²¹Per the intuition of Somaini and Einav (2011), we expect that allowing for forward-looking behavior by the type A consumers would result in higher prices for the type A vehicles and lower prices for the type B vehicles in equilibrium.

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:

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

Firm k's Bellman equation is defined similarly. These equations capture the tradeoff the firms face as the parameters μ_A and μ_B , which govern the strength of brand loyalty, vary. 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).²² In the simulations presented below, we fix $\delta = 0.5$, 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 discount factor of 0.5 between periods is consistent with

²²Without intergenerational brand loyalty ($\mu_A = 0$), the model reverts to a standard two-period game (akin to that of Klemperer (1987)) that can be characterized analytically, though the results presented below were nonetheless generated numerically.

an annual discount factor of 0.95 and reasonable assumptions for the timing of generations and the number of automobile purchases that individuals make while young and old.^{23}

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. A value of $\mu_A = 0.2$ leads to a 55% market share for brand jamong young adults whose parents also purchased brand j, which is consistent with our econometric estimates above showing that children whose parents bought a Ford have about a 10 percentage point higher probability of choosing a Ford over a GM, and vice versa. A larger value of about 0.5 would be consistent with our Honda versus Toyota results. Econometric estimates for within-household brand loyalty over time (not shown) imply that μ_B is about 1.5 to 2 times as large as μ_A .

5.2 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 because 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 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. The "investment" effect in the Amarket outweighs the "harvesting" effect in the B market (that is, average vehicle price is less than the no-loyalty baseline price of 1.25) for values of μ_B up to about 0.69. If brand

 $^{^{23}}$ If we assume that "young" adults purchase vehicles every five years from age 25 to 40 and that "old" adults purchase vehicles every five years from 45 to 75, that children are born when their parents are aged 30, and that automakers use an annual discount factor of 0.95, then the present-value number of purchases for old adults is 47% that of young adults, while the present-value number of purchases for young adults is 46% that of their parents.



Figure 1: Steady state prices with two symmetric firms

Note: Steady state equilibrium prices shown are from the model described in section 5.1 in which $\delta = 0.5$, 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 brand loyalty are equal.

loyalty is stonger than that, however, then the "harvesting" effect dominates in our model.

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 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. Steady state equilibrium prices are lower than in the case of no brand loyalty for values of μ_B up to about 0.68.

Finally, the dashed line plots an intermediate case in which intergenerational brand preference parameter, μ_A , is half as large as the within-household parameter, μ_B . Not surprisingly, this case lies between the two other cases. Here, average vehicle price is less than the no-loyalty baseline price of 1.25 for values of μ_B up to about 0.82.

Given that the values of μ_A and μ_B that apply to the automobile industry likely fall below 0.7, 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 seems unlikely given our estimates above.²⁴ The intuition for why moderate levels of brand loyalty reduce prices in equilibrium is given in 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 $\mu_B = 0.4$ and $\mu_A = 0.2$, which roughly corresponds to our empirical results above regarding intergenerational transmission of preferences and auxiliary results (not presented) on the strength of brand preference within the same household over time. 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, 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,

²⁴A value of $\mu_B = 0.7$ implies that a household loyal to brand *j* facing equal prices for vehicles *jB* and *kB* would choose *jB* with a probability of 0.67. Such a choice probability would imply a coefficient of 0.34 in a linear probability model with two firms; this value exceeds the magnitudes of brand loyalty 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 5.1 in which $\delta = 0.5$, V = 1, $\alpha = 8$, and c = 1. μ_B , which governs the strength of within-household brand preference transmission, is 0.4, and μ_A , which governs intergenerational transmission, is 0.2. 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.

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.

5.3 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: Why then do automakers appear to encourage brand loyalty amongst their customers? That is, why do automakers 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?²⁵ 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} . Parameters 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 is able to increase its price, while 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} .

These equilibrium pricing strategies cause firm k's profits to increase in 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 5.2 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 can exceed that which would maximize industry profits (see, for instance, Dixit and Norman (1978) and Grossman and Shapiro (1984)).

²⁵For example, see two recent television advertisements for the Toyota Camry and Chevy Silverado at http://www.youtube.com/watch?v=46pmd-qq_60 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 5.1 and 5.3 in which $\delta = 0.5$, 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} .

6 Conclusion

Our analysis of the PSID suggests that automobile brand preferences may be passed through the generations in ways that are important to the strategy of automobile producers. Our work documents a strong correlation across generations in brand choice. This relationship remains strong even when limiting the analysis to similar brands and controlling for a rich set of demographic factors and fine geographic fixed effects, which leads us to conclude that intrafamily correlations are not driven entirely by correlated characteristics but rather reflect an important role for what we call cross-household state dependence or direct brand preference transmission. Our results regarding duration of ownership and the differential impact of vehicles owned while a child lived with his or her parents are suggestive of an important role for cross-household state dependence, in which actual ownership of and experience with a vehicle is critical for influencing brand choice across generations.

Using a model that captures firms' dynamic pricing incentives, we show that, in theory,

the transmission of brand preferences are likely to depress automakers' prices and profits in equilibrium. Intuitively, 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 that we establish 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 relatively slow growth of firms, such as Toyota and Honda in the 1980s, and Hyundai today, that offer only entry-level vehicles, but do so at a low price given their high quality.

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

Altonji, Joseph G., Todd E. Elder, and Christopher R. Taber, "Selection on Observed and Unobserved Variables: Assessing the Effectiveness of Catholic Schools," *Journal* of Political Economy, February 2005, 113 (1), 151–184.

- 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.

- 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.
- and Nathan Miller, "Automakers' Short-Run Responses to Changing Gasoline Prices," *Review of Economics and Statistics*, Forthcoming.
- 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, February 2011, 3 (1), 157–79.
- 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.