APPENDIX

A THEORY MODEL DETAILS

This appendix provides additional information and results for the theory model that is described briefly in section V. We focus on implications unique to intergenerational brand preference transmission by considering markets, such as those for automobiles, in which firms can discriminate between younger and older consumers. In particular, we explore how firms offering multiple products—each catering to a different age class—will want to price and market these products differently in the presence of intergenerational state dependence.[36]

The related literature (see especially Klemperer [1987], Dubé et al. [2009], and Somaini and Einav [2013]) has typically used 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. This is identical to how we implement what we here call state dependence.[37] Our model differs from the existing literature in that we allow the choices of older consumers to affect the preferences of their children (next period's younger consumers) and that we model firms as offering different products to consumers of different ages.[38] 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 state dependence can play in determining automobile prices in equilibrium.

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. We forego a richer model that would more closely match the current automobile industry—a model

[36]As discussed in section II above, intergenerational brand preference transmission will only have implications for pricing strategies if it operates via the state dependence mechanism. In contrast, the direct preference inheritance mechanism will affect automakers' advertising and product line strategies but will not affect their pricing strategies.

[37]Some papers model switching costs as an increase in utility from purchasing the same brand that was purchased in the previous 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.

[38]A further generalization of our model would be to allow for a broader set of peer effects so that, for example, older consumers' choices could influence other older consumers. While this generalization is beyond the scope of this paper, we believe that such peer effects would have similar effects to what we find here for intergenerational state dependence: a relative reduction in the markup for upscale vehicles (since within-household brand loyalty becomes less important in the presence of cross-household state dependence). Moreover, our focus on parent-to-child transmission is motivated by our empirical finding that the parent-to-child channel is stronger than that for other family links and by the extensive literature showing intergenerational correlations in many economic measures [Black and Devereux 2011].

with more than two major firms, many products per firm, and richly differentiated consumer preferences—for several reasons.[39] 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.[40] 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.[41]

A(i) A simple model of automobile pricing under brand loyalty

In our model, there are two symmetric firms, denoted j and k, that compete in a differentiated Bertrand pricing game with an infinite horizon and overlapping generations of consumers. In each period, there are unit masses of two types of households: young (type A) and old (type B). All consumers are born as type A, become type B in the second period of their lives, and then die, creating a new type A consumer (offspring) upon death. All consumers purchase exactly one vehicle in each period of their lives, and there is no outside good. We have in mind that children are present in the type B households, are exposed to their parents' brand choice, and then become type A consumers upon leaving their parents' home. A key feature of the model is that the type A and type B consumers purchase different kinds of cars. Both firms are aware of this fact, and both sell two vehicle models catering to the two types. Thus, there are four vehicles in the market: jA, jB, kA, and kB. Car types A and B 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 consider only vehicles jA and kA and type B households consider only vehicles jB and kB. Clearly this is an abstraction, as there will be some substitution by households across vehicle types.[42]

[39] The treatment of consumers as living for two periods is also an abstraction, though one that is in line with much of the switching cost literature (Klemperer [1987] and Somaini and Einav [2013], for example). We revisit the question of the duration of consumers' lifetimes and the time gap between periods when we discuss the model's discount factor further below.

[40] Dubé 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.

[41] 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.

[42] Allowing for some cross-age substitution has essentially no impact on models in which intergenerational

Still, in a survey of over 22,000 consumers by a market research firm described in Langer [2012], 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.

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. The parameter μ_A denotes the strength of intergenerational brand preferences. This formulation assumes that the parents' type A car—which we imagine to be owned by parents before the next generation is born—does not influence the child's utility function. (Thus, our two-period formulation does not distinguish between short- and long-run state dependence, but it does exclude direct brand preference inheritance that could cause parental choices that occurred before a child was born to influence child choice.) 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

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 continues to hold. This is true even in the extreme case in which there is no intergenerational brand preference and consumers have no systematic preference for their own type of vehicle. This last model is similar to that of Doganoglu [2010], in which consumers live for two periods and the (single product) firms cannot distinguish between young and old.

purchase vehicle jA or kA.[43] 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.

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

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

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 using value function iteration

[43] Per the intuition of Somaini and Einav [2013], we expect that allowing for forward-looking behavior by type A consumers would result in higher prices for type A vehicles because type A consumers will become less sensitive to current price changes. As a second-order effect, prices for type B vehicles should then fall in equilibrium because the continuation value of future type A consumers will have increased.

techniques.[44] In the simulations presented below, we fix $V = 1$, $\alpha = 8$, 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 et al. [1995].

The choice of discount factor merits discussion. We treat automakers as having an annual real discount factor of 0.9 and treat the time between periods in the model as five years, which roughly corresponds to the average vehicle holding time in our data. Thus, the discount factor δ used in our model is $0.9^5 \approx 0.59$. An obvious tension here is that consumers live longer than ten years. One alternative approach would be to use a discount factor that reflects the time gap between generations (with a value of $0.9^{25} \approx 0.07$, for example). In this case, automakers would care little about future generations when setting prices. However, this alternative approach neglects the fact that consumers purchase vehicles more frequently than every 25 years and that children begin purchasing vehicles of their own soon after they are exposed to purchases their parents made while they were teenagers still living at home.

Ideally, we would resolve this issue by studying a model in which consumers live for many periods, each five years apart, and transfer brand preferences to new consumers (their children) via their brand choices over one or several purchases relatively late in life. In such a model, even though generations would be far apart in time, the chain of purchases occurring every five years—and in particular across the short, potentially overlapping transition from one generation to the next—would give automakers a sufficient incentive to consider the next generation when setting prices for upscale vehicles. While the high dimensionality of the state and decision spaces of such a model precludes its implementation, a previous version of this paper did present a computationally feasible version of such a model that forces automakers to group consumers into two broad classes (young and old), selling one vehicle type to each class and forbidding within-class price discrimination. Despite this alternative model's large time gap between generations, we nonetheless found qualitatively similar results to those discussed below. In particular, intergenerational state dependence eliminates differences in markups between vehicles targeted at young versus old consumers.

Finally, 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. Our estimates from section IV correspond to values of μ_A within the zero to one range. For our preferred 'all brands' specification (column 5 of table II),

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

the corresponding value of μ_A is about 0.35.[45] For the Ford/GM regression, μ_A is about 0.17, while μ_A is about 0.65 for the Toyota/Honda regressions. To calibrate values of μ_B, we have estimated within-household brand choice correlations that parallel the intergenerational correlations presented in tables II and III. These regressions, which are discussed in more detail in appendix B, suggest that μ_B is about 0.75 in the regression with all brands and 0.50 in the Ford/GM regressions—so roughly 2-3 times as large as μ_A.

A(ii) 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 split the A and B markets equally (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, the equilibrium 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 A market 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 0.83. If brand loyalty is stronger 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 Dubé et al. [2009]. Relative to the case with no intergenerational state dependence, type B prices fall because high type B prices now reduce future demand and

[45] Given a value for μ_A and assuming the brands in the choice set yield equivalent utility in the absence of a brand preference, the effect of parents' ownership on the probability of brand choice in our model is given by $(e^{\mu_A} - 1)/(e^{\mu_A} + (N-1))$, where N is the number of brands in the choice set.

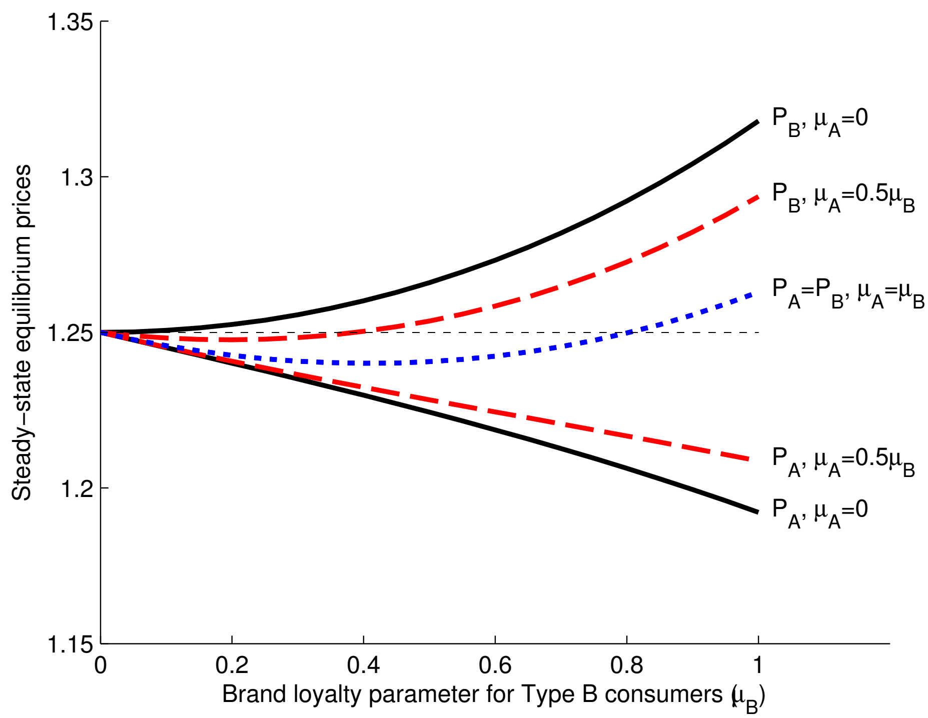

Figure 1: Steady state prices with two symmetric firms

Note: Steady state equilibrium prices shown are from the model described in section A(i) in which $\delta = 0.9^5$, $V = 1$, $\alpha = 8$, 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 and within-household brand loyalty are equal.

profits, and type A prices rise because investing in future type B consumers is no longer as profitable. The prices of the type A and B vehicles—now equal in steady state—are roughly equal to the average of the type A and B prices from the no intergenerational state dependence case. That is, intergenerational state dependence appears to primarily affect the distribution of prices across types rather than the average price in the market. Thus, similar to the no intergenerational state dependence case, steady state equilibrium prices are lower than in the case of no brand loyalty for values of μ_B up to about 0.80.

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. This case is consistent with our empirical estimates of the relative strength of intergenerational state dependence and within-household state dependence. 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 0.97, which exceeds our preferred estimates for μ_B of about 0.5 to 0.75. This implies that the existence of brand loyalty causes a net reduction in firm profits, which accords with the theoretical intuition of Cabral [2009].[46]

B WITHIN-HOUSEHOLD BRAND LOYALTY

In order to understand the size of intragenerational brand loyalty relative to within-household brand loyalty, we compute the size of within-household brand loyalty in our data. To that end, we use the sample of households whose purchases can be matched to their parents' prior vehicle purchase and estimate regressions analogous to those in table 2 of the text, with the brand of the household's most recent purchase in place of the brand of the parents' most recent purchase. Columns 1 through 4 of Table VIII present these results. We do not include specifications that include census tract fixed effects because these present an incidental parameters problem in what is essentially a lagged dependent variable regression.

Table VIII: Correlations between household brand choice and previous household brand choice

	Dependent Variable: Household's Brand			
VARIABLES	(1)	(2)	(3)	(4)
Household's brand = Lagged Household brand	0.225	0.218	0.203	0.137
	(0.007)	(0.007)	(0.007)	(0.007)
Month of purchase fixed effects	Yes	Yes	Yes	Yes
Household's demographics	No	Yes	Yes	Yes
Child's state fixed effects	No	No	Yes	Yes
Household's county fixed effects	No	No	No	Yes
Household's census tract fixed effects	No	No	No	No
Number of choices	17,268	17,268	17,268	17,268
R^2	0.117	0.124	0.138	0.220

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). Household's demographics include age, education, income, gender, number of children in household, and family size. All control variables are interacted with 7 dummies, one for each brand. Columns 5 and 6 limit the sample to households living in census tracts that contain more than one PSID family.

We then run the same specifications using the subsample of households that purchased a Ford or GM vehicle. As in the intergenerational brand preference results, we do not require that the household's previous vehicle purchase was also a Ford or GM, but we do include a dummy variable that is equal to 1 if the previous vehicle was a Ford or GM. Table IX shows

[46] As noted above, a previous version of this paper contained a more detailed model in which consumers purchased cars multiple times while young and old. Results from this more detailed model also support this conclusion.

within-household brand loyalty results that are analogous to the intergenerational brand loyalty results presented in table 3 of the text.

Table IX: Correlations between household vehicle brand and previous household brand choice among those owning a Ford or GM

	Dependent Variable: Household's Brand			
VARIABLES	(1)	(2)	(3)	(4)
Household's brand = Lagged Household brand	0.369	0.368	0.345	0.243
	(0.019)	(0.019)	(0.019)	(0.020)
Month of purchase fixed effects	Yes	Yes	Yes	Yes
Household's demographics	No	Yes	Yes	Yes
Household's state fixed effects	No	No	Yes	Yes
Household's county fixed effects	No	No	No	Yes
Household's census tract fixed effects	No	No	No	No
Number of choices	9,355	9,355	9,355	9,355
R^2	0.136	0.137	0.153	0.277

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 individual-year-vehicle choice enters the data 7 times, once for each brand (GM, Ford, Chrysler, Toyota, Honda, Other Asian, and European). Household's demographics include age, education, income, gender, number of children in household, and family size. All control variables are interacted with 7 dummies, one for each brand. Columns 5 and 6 limit the sample to households living in census tracts that contain more than one PSID family.