EXPLOITING THE CHOICE-CONSUMPTION MISMATCH: A NEW APPROACH TO DISENTANGLE STATE DEPENDENCE AND HETEROGENEITY

By

K. Sudhir and Nathan Yang

March 2014

COWLES FOUNDATION DISCUSSION PAPER NO. 1941

COWLES FOUNDATION FOR RESEARCH IN ECONOMICS
YALE UNIVERSITY
Box 208281
New Haven, Connecticut 06520-8281

http://cowles.econ.yale.edu/

Exploiting the Choice-Consumption Mismatch: A New Approach to Disentangle State Dependence and Heterogeneity[1]

K. Sudhir[2] Nathan Yang[3]

March 12, 2014

Abstract

This paper offers a new identification strategy for disentangling structural state dependence from unobserved heterogeneity in preferences. Our strategy exploits market environments where there is a choice-consumption mismatch. We first demonstrate the effectiveness of our identification strategy in obtaining unbiased state dependence estimates via Monte Carlo analysis and highlight its superiority relative to the extant choice-set variation based approach. In an empirical application that uses data of repeat transactions from the car rental industry, we find evidence of structural state dependence, but show that state dependence effects may be overstated without exploiting the choice-consumption mismatches that materialize through free upgrades.

Keywords: Consumer dynamics; Heterogeneity; Quasi-experiment econometrics; Service industry; State dependence.

[1] We are grateful to the Wharton Customer Analytics Initiative and an international car rental company for access to the data used for our empirical application.
[2] Yale School of Management, email: k.sudhir@yale.edu.
[3] Yale School of Management, email: nathan.yang@yale.edu.

1 Introduction

Consumer choice shows remarkable stickiness across time. The stickiness may be due to persistent unobserved heterogeneity---preferences that differ across consumers but remain stable with consumers over time; or due to state dependence---a consumer's current choice drives the higher likelihood of the same choice in the future.[4] Disentangling state dependence from heterogeneity has been a major challenge in the literature since Heckman (1981) highlighted the confounding nature of structural state dependence and persistent unobserved heterogeneity. The key takeaway is that not adequately accounting for heterogeneity can exaggerate the estimated level of state dependence. This is not merely an econometric quibble; disentangling these two sources of stickiness in choice across time is important in developing dynamically optimal policies. For example, the optimality of policies pertaining to advertising (e.g., Dube, Hitsch, and Manchanda, 2005; Freimer and Horsky, 2012; Mahajan and Muller, 1986), consumer finance (e.g., Barone, Felici, and Pagnini, 2011; Israel, 2005a, 2005b), federal procurement (e.g., Greenstein, 1993), health (e.g., Arcidiacono, Khwaja, and Ouyang, 2012; Handel, 2013; Iizuka, 2012; Janakiraman et. al., 2008; Naik and Moore, 1996), housing (e.g., Moon and Stotsky, 1993), labor (e.g., Biewen, 2009; Coelli, Green, and Warbuton, 2007; Heckman, 1981; Hyslop, 1999; Prowse, 2012), and pricing (e.g., Che, Sudhir, and Seetharaman, 2007; Cosguner, Chan, and Seetharaman, 2012; Dube et. al., 2008; Dube, Hitsch, and Rossi, 2009, 2010; Pavlidis and Ellickson, 2012) are crucially dependent on whether structural state dependence or heterogeneity drives stickiness in choice.

The literature has thus far relied on a combination of *functional form assumptions about the nature of heterogeneity* and *choice set variation across time* to disentangle unobserved heterogeneity and state dependence. Early on, researchers highlighted the

[4] Some economic mechanisms behind structural state dependence may include consideration set formation, switching costs, and/or learning.

role of functional form assumptions on the structure of unobserved heterogeneity, that permitted them to numerically integrate out the effect of unobserved heterogeneity on choice behavior using simulation-based econometric methods (Arcidiacono, Khwaja, and Ouyang, 2012; Erdem and Sun, 2001; Hyslop, 1999; Iizuka, 2012; Keane, 1997; Prowse, 2012; Seetharaman, 2004), and attribute the residual stickiness in choice behavior to state dependence.[5] Scholars continue to increase the level of flexibility they allow in the functional forms (Dube, Hitsch, and Rossi, 2010; Honore and Kyriazidou, 2000; Moon and Stotsky, 1993), to limit the possibility that a lack of adequate accommodation of heterogeneity does not lead to exaggerated estimates of state dependence. In recent years, researchers in industrial organization and marketing have highlighted the importance of choice set variation over time as an essential ingredient of the disentangling strategy, beyond the functional form assumptions on unobserved heterogeneity. The choice set variation can occur in the form of changes in price (e.g., Dube, Hitsch, and Rossi, 2010), advertising (e.g., Terui, Ban, and Allenby, 2011), availability of alternatives (e.g., Goldfarb, 2006b), or decision context (e.g., Thomadsen, 2012). Some scholars have augmented data to include some forms of observable heterogeneity either in the form of household demographics (e.g., Goldfarb, 2006a; Gupta, Chintagunta, and Wittink, 1997; Paulson, 2011, 2012) or through direct surveys of preferences (e.g., Shin, Misra, and Horsky, 2012), but how much residual unobserved heterogeneity remains beyond these observable controls remains an issue. Thus, despite the large volume of literature on the topic, this identification challenge still remains an open area of research, because existing methods are unable to fully disentangle unobserved heterogeneity from state dependence.

In this paper, we introduce a new identification strategy to disentangle state dependence and unobserved heterogeneity through only revealed preference data via *exclusion*

---

[5] Furthermore, researchers have also uncovered variety seeking in choice as a form of "negative" state dependence (Chintagunta, 1998, 1999; McAlister, 1982) in certain market settings.

restrictions that arise in market environments where a consumer's choice may not match their consumption. Consider the following setting in the context of rental cars; Customers make reservations for a car ahead of time; but when they arrive to pick up the car, the reserved car might be out of stock, and therefore the customer may be offered a free upgrade to a different car at no additional cost. Such upgrades due to inventory shortages are common in many settings (Biyalogorsky et. al., 1999, 2005; Wangenheim and Bayon, 2007), leading to a mismatch between choice and consumption. As in the past literature, choice is affected by preferences and state dependence, but the consumption based on upgrades only affects state dependence; thus providing an exclusion restriction necessary to disentangle state dependence from heterogeneity.

The choice-consumption mismatch can occur in other situations. For instance, free samples may induce customers to consume products they had initially chosen not to try (Bawa and Shoemaker, 2004; Cabral, 2012; Halbheer et. al., 2013; Pauwels and Weiss, 2008; Scott, 1976). Stock-outs in online retail would force customers to consume alternatives if the item they originally clicked on is no longer available (Anupindi, Dada, and Gupta, 1998; Bruno and Vilcassim, 2008; Conlon and Mortimer, 2010, 2013; Diels, Wiebach, and Hildebrandt, 2013; Jing and Lewis, 2011; Musalem et. al., 2010). When customers make purchases with e-commerce retailers, errors in shipped purchases present lead to consumption of products, they were not originally ordered (Collier and Bienstock, 2006a; Collier and Bienstock, 2006b; Gregg and Scott, 2008; Vaidyanathan and Devaraj, 2008). Finally, product recalls force customers to cease the use of originally purchased items in favor of alternatives offered by the firm (Freedman, Kearney, and Lederman, 2012; Haunschild and Rhee, 2004; Marsh, Schroeder, and Mintert, 2006; Van Heerde, Helsen, and Dekimpe, 2007). There are two common characteristics across these examples. First, it is feasible in all of these examples to collect first data on choice before the consumption occurs (e.g., reservations for services, items to be or already checked-out in shopping cart). Second, consumption is shifted in ways that need not be correlated with unobserved preferences.

We begin by providing a heuristic proof of why choice-consumption mismatches help disentangle state dependence and heterogeneity, and why it is superior to the traditional strategy of using choice set variation in combination with rich functional forms to accommodate unobserved heterogeneity. We then demonstrate its effectiveness through a Monte Carlo analysis, where we simulate data consistent with a simple multinomial choice model with both persistent unobserved heterogeneity and structural state dependence, accommodating choice set variation and choice-consumption mismatches. Estimates from our simulated datasets show that choice set variation does help reduce the upward bias, but not as well as the choice-consumption mismatch data. Further unlike choice-consumption mismatches, choice set variation does not completely debias the state dependence parameter.

We then perform an empirical analysis using repeat transactions data from the car rental service industry. Free upgrades driven by inventory shortages are a common occurrence in the industry; therefore this data allows us to exploit mismatch between choice and consumption. Our analysis of the upgrading propensity indicates that upgrades are more likely to occur when the car class a customer has chosen is in short supply---i.e., real time supply conditions at the point of consumption drive the upgrading propensity for a customer independent of customer and rental trip characteristics, providing us an exogenous source of variation in consumption that is independent of customer preferences.

Our estimates of a model of customer car class choice exploiting the choice-consumption mismatch strategy to disentangle state dependence from heterogeneity confirms that structural state dependence is indeed prevalent among consumers. Further, our simulation analysis confirms that the state dependence estimates are exaggerated without the choice-consumption mismatch data. The estimated level of state dependence is higher when we ignore households that have received free upgrades.

We later use the model estimates to perform counterfactual simulations to study the impact of implementing free upgrade policies. We find that due to our estimated level of state dependence an upgrade to a higher margin better class has long-term positive effects on revenue, in that consumers rent from the higher class in the future. To highlight potential confounding effects of unobserved heterogeneity, we show that these increases in revenue are estimated to be markedly larger than what is true when state dependence is inferred based on the sub-sample of households that did not receive upgrades and for whom therefore estimates of state dependence are exaggerated due to the confound with heterogeneity.

2 Related Literature

Functional form assumptions and choice set variation are commonly exploited in research about state dependence (Ackerberg, 2003; Erdem and Keane, 1996; Erdem and Sun, 2001; Keane, 1997; Osborne, 2010; Seetharaman, 2004). However, there remain concerns about the validity of such assumptions. For instance, Paulson (2011) argues that simulation-based estimation procedures rely too heavily on correctly specifying the structure of unobserved heterogeneity. Dube, Hitsch, and Rossi (2010) relax these functional form assumptions and offer a semi-parametric approach to flexibly account for heterogeneity in order to disentangle state dependence and unobserved heterogeneity. To aid in their identification, the authors exploit variation in price discounts as a means to vary choice sets. In a similar manner as price discounts, Goldfarb (2006b) exploits variation in choice sets[11] of online portals due to exogenous

[11] Although Bruno and Vilcassim (2008) do not study long-run effects, variation in retail stock-outs may be applied in a similar manner as Goldfarb (2006b).

changes in availability following denial of service attacks, Handel (2013) uses a change to insurance provision, Thomadsen (2012) uses variation in store choice, and Liu, Derdenger, and Sun (2013) exploit differences in compatibility between various base products and add-ons that affect the choice set for purchasing add-ons.

Paulson (2012) argues that price promotions alone may not induce enough variation in choice sets to facilitate the disentangling of state dependence from heterogeneity. The main issue is that past purchase decisions are always going to be functions of unobserved heterogeneity; to truly disentangle state dependence the variation in choice sets need to be sufficiently large to induce purchases that would not have been made otherwise. Her suggestion is to supplement choice set variation in prices with demographic and/or survey data. For instance, Shin, Misra, and Horsky (2012), and Pavlidis and Ellickson (2012) use supplementary survey response data, while Goldfarb (2006a) and Gupta, Chintagunta, and Wittink (1997) incorporate household-specific heterogeneity using demographic data. Regardless of how well this additional information generates variation in choice sets, the core issue that Paulson (2012) brought up remains, as past decisions are still affected by unobserved heterogeneity. It is this core identification problem that our new exclusion restriction based approach addresses by exploiting mismatches between choice and actual consumption.

3 Identification of State Dependence

3.1 Model and Identification Problem

In this section, we introduce and implement a Monte Carlo simulation exercise to demonstrate the identification power of forced substitution via mismatches between choice and consumption. These simulations are meant to illustrate that mismatches help reduce the positive bias of inferred structural state dependence.

For these simulations, we consider a simple discrete choice model in which customer i chooses to purchase among $j \in \{1,2,...,J\}$ products or services. A customer who chooses product j during transaction t is denoted as $d_{it} = j$. Choosing the baseline option of 1 yields zero utility for the customer. To be consistent with our empirical application, we consider the case here where products are vertically differentiated, and increase in quality such that $\alpha_j > \alpha_{j-1}$.[12] Therefore, a customer receives the following utility from $d_{it} = j$:

$$U_{ijt} = \alpha_j + \beta p_{ijt} + \gamma s_{ijt} + \omega_i + \varepsilon_{it}$$

A customer chooses j if and only if $U_{ijt} > U_{ikt}$ for all $k \neq j$. Persistent unobserved heterogeneity is included in this model via $\omega_i \sim N(0,\sigma_\omega^2)$, ε_{it} is an i.i.d. Type I Extreme Value random variable, and prices are given by p_{ijt}. Structural state dependence is captured by the parameter γ, where $s_{ijt} = 1\{c_{it-1} = j\}$ is a dummy variable indicating whether or not the customer consumed the same product in the previous transaction.

Our primary objective is to obtain as accurate of an estimate for structural state dependence as possible, in the presence of unobserved heterogeneity. It is well known that persistence in behaviors can be caused by unobserved heterogeneity, as past consumption is usually correlated with ω_i. In the typical case, $d_{it} = c_{it}$, then is clear that past brand choice decisions (and therefore consumption) are correlated with unobserved preferences that persist over time as. Therefore, estimates of γ will be confounded by ω_i. To avoid such confounds, one would then need some method of varying d_{it-1} in ways that are independent of unobserved preferences.

[12] Note that the identification arguments we make do not depend on vertical differentiation.

3.2 Identification Based on Choice-Consumption Mismatch

As explained earlier, the choice-consumption mismatch varies d_{it-1} independent of unobserved preferences to help disentangle state dependence from heterogeneity.

Figure 1 Diagram Illustrating Mismatch Between Choice and Consumption

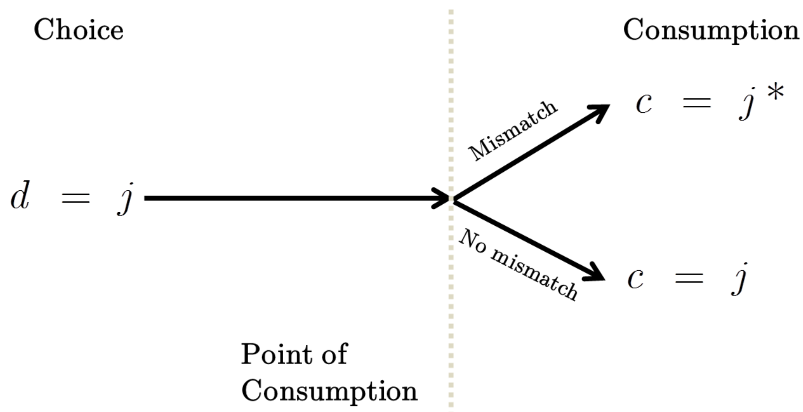

Figure 1 provides a decision diagram that describes potential mismatches between choice and consumption ($d_{it} \neq c_{it}$). Here, a customer who has originally chosen option j may potentially be forced to consume a different product j^*. We denote such an event as $m_{it} = 1$. This mismatch event occurs with a probability of λ that is independent of customer characteristics (e.g., supply driven factors such as inventory shortages).

The assumptions that we need for this identification strategy to be valid are as follows:

$$\varepsilon_{it} \perp c_{it-1}, \omega_i \mid m_{it-1}$$

$$\omega_i \perp c_{it-1} \mid m_{it-1}$$

We now illustrate the conditions for which choice-consumption mismatches serve as an effective exclusion restriction using a simple heuristic proof. When mismatches are often induced by factors exogenous to the customer (as in the examples described in the introduction), the assumption that $\omega_i \perp m_{it-1}$ holds.

Based on the model we have described, we can write lagged consumption in light of choice-consumption mismatches as follows:

$$c_{it-1} = (1 - m_{it-1}) d_{it-1}(\alpha, \beta, \gamma, \omega_i) + m_{it-1} j^*$$

It then becomes clear that as the probability of a mismatch increases, the degree to which ω_i confounds the expected consumption measure approaches zero. Consequently, the requirement that $\omega_i \perp c_{it-1} \mid m_{it-1}$ is likely to be satisfied with large values of λ.

Researchers have in the past disentangled structural state dependence from unobserved heterogeneity using choice set variation. Using a similar model as before, we now explore the identification power of such variation in the. The difference now is that instead of a potential mismatch between choice and consumption, there is a probability, which we denote as λ, that a customer's choice set changes. For our exposition, we frame these choice set changes around price discounts. In the event that a customer faces a change in the choice set, the new price for j is $p^*_{ijt} = \delta p_{ijt}$, where $\delta \in (0,1)$ is the fraction of the original price that the customer would have had to pay. With this new choice set, the customer then makes decision d^*_{it}, instead of d_{it}. When the customer does not encounter a choice set change, the price remains at p_{ijt}. Based on the model we have described, we can write lagged consumption in light of price discounts as follows:

$$c_{it-1} = (1 - m_{it-1}) d_{it-1}(\alpha, \beta, \gamma, \omega_i) + m_{it-1} d^*_{it-1}(\alpha, \beta, \gamma, \omega_i, \delta)$$

Notice that even when the probability of a change in consumption set is large via frequent price discounting, lagged consumption remains a function of unobserved preferences. Hence while choice set variation can reduce the bias, it can almost never truly debias the state dependence estimate.

3.3 Monte Carlo Analysis

We now illustrate using a simulation the bias reduction benefits of the choice-consumption mismatch strategy for identifying state dependence.

For our first set of simulations, we consider a scenario with 1,000 customers who make 5 repeat purchases each, and are potentially faced with choice-consumption mismatches. Each customer can choose between three products, $j \in \{1,2,3\}$, where product 1 is the baseline option that yields zero utility. In terms of the other parameterizations, we set the intercepts as $\alpha_2 = 0.1$ and $\alpha_3 = 0.8$ respectively. Here, product 3 is of a higher quality than product 2. Price sensitivity is set at $\beta = -0.3$. State dependence effects are set at $\gamma = 0.6$. For the variance of unobserved heterogeneity, we set $\sigma_\omega = 5$. We try different values for the mismatch probability, namely $\lambda \in \{0.25, 0.5, 0.75\}$. For the prices of products 2 and 3, we draw them from a truncated Normal distribution with means 0.2 and 0.9 respectively.

With each parameterization, we forward simulate the sequence of choices (d_{it}) and actual consumption (c_{it}) for each customer, which serve as the simulated datasets for our subsequent estimations. To implement the choice-consumption mismatches, we try to mimic an environment in which customers are given free upgrades. Therefore, with probability λ, customers who had originally chosen the lower two options, 1 and 2, may be upgraded for free to option 3 instead (i.e., $j^* = 3$).

For our next set of simulations, we consider again a scenario with 1,000 customers who make 5 transactions each and face the possibility of facing a new choice set with probability λ. We set the same parameters as before. In these simulations, we now have the additional parameter, which is the price discount set at $1-\delta=0.25$. This price discount is applied to product 3.

Table 1 Estimates for Structural State Dependence Using Simulated Data

	Choice-consumption mismatch					Choice set variation				
λ	Estimate	SE	95% CI		% bias	Estimate	SE	95% CI		% bias
25%	1.249	0.007	1.235	1.262	108%	1.559	0.012	1.537	1.582	160%
50%	0.819	0.006	0.807	0.831	37%	1.363	0.013	1.338	1.388	127%
75%	0.606	0.007	0.593	0.619	1%	1.322	0.015	1.293	1.352	120%

We can then estimate the model parameters using each of the simulated datasets. To estimate this discrete choice model, we use simulated maximum likelihood. Table 1 provides us the estimates of structural state dependence from each of the simulated datasets. The first three columns provide us the results from simulations that exploit the choice-consumption mismatches, while the latter three columns provide us the results from simulations that exploit choice set variation. Recall that the data was generated with the state dependence parameter $\beta=0.6$, we wish to determine how effective the choice-consumption mismatch and choice set variation are at eliminating the bias.

We first look at the bias reduction from increasing the mismatch probability, as suggested earlier in our discussion about identification. Confirming the intuition behind our assertion, we see that the estimates approach the true value as λ increases. Most importantly, the bias is virtually eliminated when customers face a high probability of choice-consumption mismatch. Furthermore, the true value of state dependence lies within the 95% confidence interval for the estimates. In our simulations with variation in choice sets, the bias reduction associated with changes in the choice set is markedly less than in our simulations with the choice-consumption mismatch; the confidence interval does not include the true parameter value even with high probability of choice set variation.

To summarize, this Monte Carlo analysis demonstrates the benefit of exploiting the choice-consumption mismatch in disentangling state dependence and heterogeneity, the greater the frequency with which such mismatches occur, the greater the potential to reduce the bias in estimates of state dependence due to the confound with unobserved heterogeneity. In fact, unlike choice set variation that does not completely eliminate bias, the mismatch approach has the potential to completely debias the state dependence estimate.

4 Empirical Application: Car Rental Industry

4.1 Data Description

Our setting is the car rental industry, in which we utilize a sample of data from an international car rental company on repeat transactions of customers from 2011 to 2012. Repeat customers are identified in the data via their loyalty program membership.

As shown in Table 2, about 20% of the users rented 2 times, while about 6% and 2% rented 3 and 4 times respectively. The remaining 3% of users rented 5 or more times. As our empirical analysis of state dependence will be based on the car class choice among travelers, we focus on the subset of customers that have booked with the car rental company at least twice over the course of 2 years. This leaves us with nearly 100,000 transactions. As is standard in the choice literature, we assume here that customers who rent only once and customers who rent multiple times are not different in terms of their unobserved preferences towards car class alternatives.

For each transaction, we can identify which car class was booked, driven, and paid for. Classes are vertically differentiated, so the higher the class, the higher the quality of the car rental.[14] In the event that a user drives a higher class than was originally booked, and pays for the higher class, we would classify that transaction as being an *upsell*. About 2% of the sample contains such upsell transactions. In the event that a user drives a higher class than was originally booked, but pays the same amount as for the class that was originally booked, we would classify that transaction as being an *upgrade*. Upgrades occur in about 51% of the sample. This high upgrade probability suggests that the empirical application using car rental data will benefit from our new identification strategy that exploits the choice-consumption mismatch. Based on the previously reported simulation, we know the choice-consumption mismatch data is more effective in debiasing state dependence estimate when the proportion of mismatches is high.

Table 2 Distribution of the Number of Transactions Across Users

Transactions	Frequency	Percent	Cumulative
1	219,491	69.53	69.53
2	58,186	18.43	87.96
3	19,554	6.19	94.15
4	6,988	2.21	96.36
5	2,790	0.88	97.25
6	1,440	0.46	97.7
7	931	0.29	98
8	752	0.24	98.24
9	504	0.16	98.4
10	560	0.18	98.57
11	418	0.13	98.71
12	324	0.1	98.81
13	286	0.09	98.9
14	294	0.09	98.99
15	195	0.06	99.05
16	224	0.07	99.12
17	187	0.06	99.18
18	162	0.05	99.24
19	95	0.03	99.27
20	240	0.08	99.34

[14] This assertion is based on insights obtained during a conference call with the car rental company's executives facilitated by Wharton's Customer Analytics Initiative on October 4, 2013. Due to a confidentiality agreement with the car rental company, we are unable to disclose exactly which exact models belong to each car class. Note that in the data there is actually a 26th class. This class is assigned to car models that belong to a range of different classes. Given the potential inaccuracies of this particular class label, we exclude all transactions involving class 26.

Figure 2 Distribution of Car Class Choices

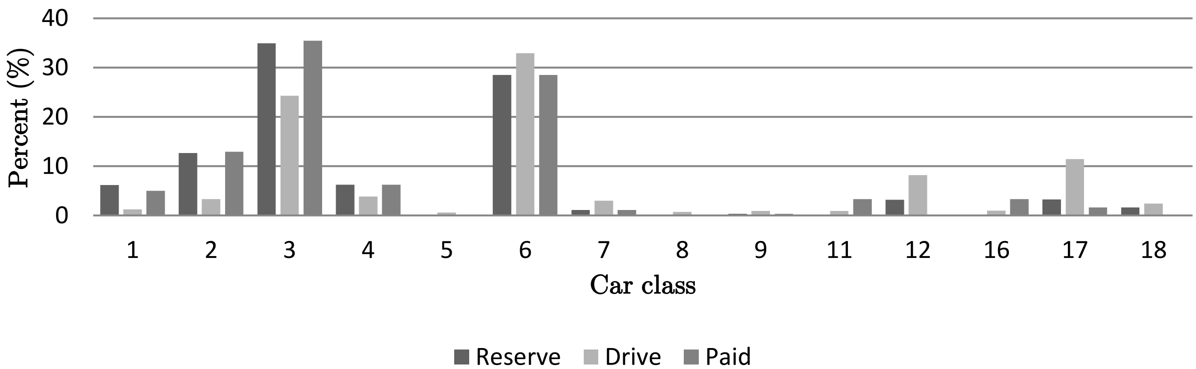

Figure 2 displays the distribution of car class choices across transactions.[16] From this histogram we see that users are primarily booking and paying for lower class cars (i.e., below 6). However, in classes 1-4, which constitutes a significant fraction of the overall transactions, a large fraction of customers do not end up driving the same car they reserved. It appears that classes 6, 12 and 17 are the most commonly used cars for providing free upgrades (i.e., they constitute about 32, 15 and 18 percent of the cars that users drive upon receiving upgrades). Class 3 has a higher proportion of people paying for it than that reserved, suggesting this class is used by the firm for upsell to those who book in classes 1 and 2 (i.e., about 45 percent of customers who originally booked classes 1 and 2 are upsold to class 3).

Table 3 Summary Statistics for Trip Characteristics

	Full sample		Upgrade		No upgrade	
Variable	Mean	Std. Dev.	Mean	Std. Dev.	Mean	Std. Dev.
Airport	0.429	0.495	0.422	0.494	0.437	0.496
Phone reserve	0.114	0.318	0.117	0.321	0.111	0.314
Business	0.382	0.486	0.405	0.491	0.357	0.479
Weekend	0.473	0.499	0.465	0.499	0.482	0.500
Duration	4.229	6.467	4.254	6.855	4.203	6.034
# transactions	2.386	3.804	2.739	4.693	2.016	2.512
Price	205.040	240.574	188.786	223.656	222.033	255.975
Age	52.308	11.828	51.923	11.812	52.710	11.831
Tier	1.970	1.132	2.138	1.214	1.795	1.011
Observations	96209		49174		47035	

[16] For visual clarity, note that the figure does not display the percentage of transactions that involve car classes 10, 13, 14, 15, 19, 20, 21, 22, 23, 24, and 25 as they each constitute less than 1%.

Other trip characteristics that we incorporate in our analysis include whether the car is rented from an airport location, is booked over the phone, is for business purposes, and/or is a weekend rental. We see also know the duration of each rental. From Table 3, about 40% of the transactions occur via airport rental locations, 11% are booked via phone, 38% are for business purposes, and 47% occur on the weekend. The typical car rental length is about 4 days. A user spends on average about $205 per transaction. The average tier of a customer is about 2, where 1 is the lowest tier and 7 is the highest.[17]

We now provide a comparison of summary statistics across users based on whether or not they received upgrades. This comparison serves to demonstrate that the observable user-trip characteristics are similar across the two sub-samples. In general, the mean and standard deviation looks quite similar across the sub-samples. The only noticeable difference is in prices, in which upgraded customers appear to be paying $30 less than customers who did not receive free upgrades.

4.2 Empirical Patterns of Upgrades

Upgrades generate choice-consumption mismatches by forcing users to experience classes that are different (and higher) than the classes originally booked, but without any additional cost. For our identification approach, we rely on the assumption that these mismatches are exogenous to consumer preferences. Based on the market environment, we suggested that these upgrades are driven by supply considerations such as inventory. It is also possible that upgrades are linked to elite status and other consumer/trip characteristics. To the extent we are able to control for such observable consumer/trip characteristics in the upgrading propensity, the supply side instruments related to inventory would serve to provide the necessary exclusion restrictions for the choice-consumption mismatch strategy to work.

[17] Higher tiers are considered to be more “elite.” Based on information provided by Wharton’s Customer Analytics Initiative, tier level membership is based on the number of rental transactions, number of rental days, a monthly or annual fee, or some combination of all three. However, it was not disclosed by the car rental company as to the exact membership requirements and benefits for each level.

Table 4 Summary Statistics for Inventory Conditions

			Percentile						
Variable	Mean	Std. Dev	1%	25%	50%	75%	99%	Min	Max
# check-out	1.128	0.453	1	1	1	1	3	1	9
Net supply	0.003	0.463	-1	0	0	0	1	-8	5

We focus on three variables that may be used to proxy for stock-outs. As the data itself does not contain inventory information, we have to infer general demand-supply conditions using the available information.[18] Table 4 provides summary statistics for the supply-side proxies we use.

The first variable we consider is the total number of check-outs for the current reserved transaction class at a particular location within the hour of rental. This measure gives us an idea about the demand for specific car classes at each rental location. With this measure, one hypothesis we first test is whether upgrade propensity increases with the demand for cars. The intuition is that if demand is high for the car class that is booked, then the chance that this booked class is no longer available is high, and thus, a greater likelihood of receiving a free upgrade. Figure 3 confirms that there is indeed a disproportionately larger amount of transactions with upgrades as the demand is high (i.e., 2 or more check-outs versus only 1 check-out).

Figure 3 Percentage Difference Between the Number of Transactions With and Without Upgrades

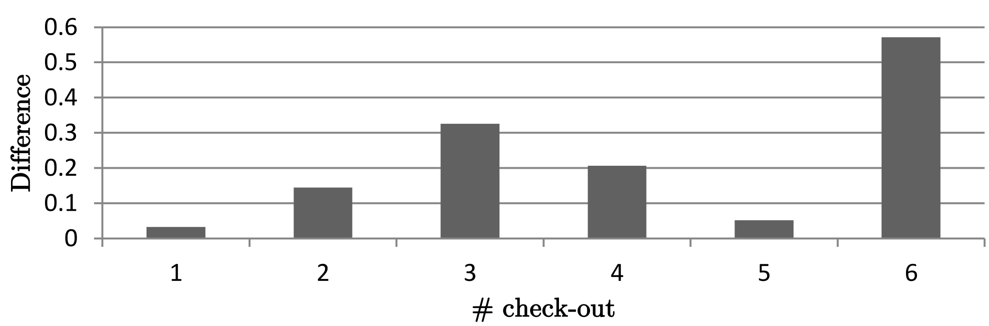

[18] The car rental company was unable to provide us data on (real-time) inventories when we requested such information.

The second variable we consider is the total number of check-ins net of the total number of check-outs at a particular location at the time of a transaction. As the number of check-ins help proxy for the number of cars returned, and the number of check-outs proxy for the number of cars demanded, the net difference of these variables may be interpreted as the net supply (or flow) of available cars. Our second hypothesis is to test whether or not upgrade propensity decreases with this measure. If the net supply is high, then the stock-out probability is low, thereby reducing the likelihood of free upgrades. Figure 4 confirms our intuition, since the percentage difference between the number of transactions with and without upgrades diminishes as net supply increases (i.e., negative net supply versus positive net supply).

Figure 4 Percentage Difference Between the Number of Transactions With and Without Upgrades

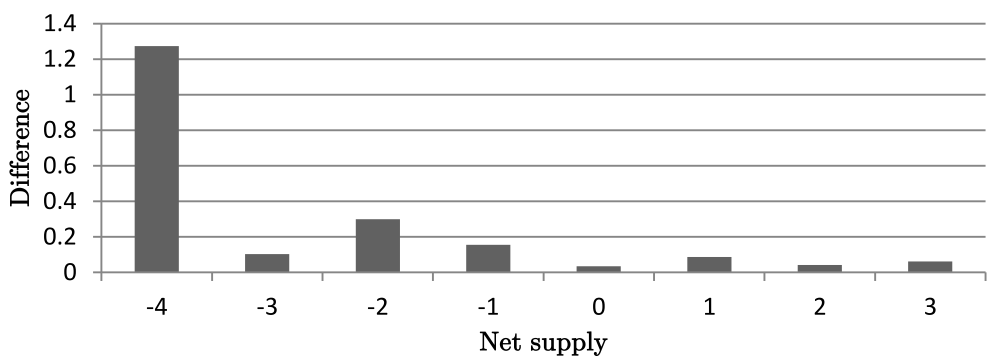

Using these supply-side measures, we estimate two different probit specifications with user-level random effects. Table 5 presents the main upgrade patterns in our data. The first column highlights our analysis using the proxy for demand. First note that upgrades are correlated with trip/user characteristics. For instance, a user is less likely to receive an upgrade at an airport, or on a weekend. Older customers, as well as those paying a higher price are also less likely to receive a free upgrade. In contrast, business users, high volume users, and those that belong to a higher tier are more likely to receive a free upgrade.

Table 5 Probit Specification for Upgrade Propensity

	Upgrade		Upgrade	
	Estimate	SE	Estimate	SE
# check-out	0.0510***	(0.00941)		
Net supply			-0.0221*	(0.00861)
Airport	-0.0810***	(0.00954)	-0.0729***	(0.00940)
Phone reserve	-0.0249	(0.0142)	-0.0259	(0.0142)
Business	0.176***	(0.00909)	0.179***	(0.00906)
Weekend	-0.0140	(0.00877)	-0.0151	(0.00876)
Duration	0.0423***	(0.00174)	0.0422***	(0.00174)
# transactions	0.0215***	(0.00181)	0.0219***	(0.00180)
Price	-0.00154***	(0.0000605)	-0.00153***	(0.0000603)
Age	-0.00231***	(0.000374)	-0.00236***	(0.000374)
Tier	0.109***	(0.00408)	0.109***	(0.00408)
Constant	-0.211	(0.160)	-0.144	(0.161)
Random effects	Yes		Yes	
Observations	96209		96209	

Most importantly, we see that upgrade propensity increases with demand. Analogously, the second column confirms a negative relationship between upgrade propensity and net supply. Even after targeting strategies based on user/trip type are controlled for, we provide empirical evidence that highlights a relationship between supply-side conditions and free upgrades.[19] These results motivate further the idea that choice-consumption mismatches (through upgrades) are likely to be driven by "exogenous" factors.

4.3 Model

This section presents the random utility logit model with endogeneity and structural state dependence that we use in our empirical application. The model contains two stages. First, customers choose the car class they wish to rent in the reservation stage. After making the reservation, customers reach the point of consumption stage, at which point the car class they end up driving may or may not be the same as the class originally chosen.

[19] Note that we also tried specifications with upsells as the dependent variable. In these specifications, we find no empirical relationship between upselling propensity and supply-side conditions. The main drivers behind observed upsells are the user-trip characteristics.

4.3.1 Reservation Stage

In the reservation stage, each customer i decides on which car class to rent at the beginning of each transaction t; we denote the decision to choose car class j as $d_{it} = j \in \{1,2,...,\mathrm{J}\}$. Customers decide on classes that yield the highest utility, where utility is defined as:

$$U_{ijt} = \alpha_j + \beta X_{it} + \gamma s_{ijt} + \omega_{ij} + \varepsilon_{ijt}$$

Customers make their decisions based on trip characteristics, represented by the vector X_{it}. Furthermore, as higher car classes are of higher quality, we include a car class intercept α_j, which we assume gets larger as j increases. There may be unobserved and persistent factors as to why some car classes are inherently preferred by some customers, which we model using random effects $\omega_{ij} \sim N(0,\sigma^2)$. The error term ε_{ijt} follows an i.i.d. Type I Extreme Value distribution.

State dependence is captured by the state variable $s_{ijt} = 1\{c_{it-1} = j\}$, which is an indicator for whether in the previous transaction, the user actually drove class j in the previous transaction.

4.3.2 Point of Consumption Stage

Each transaction is completed at the point of consumption, which is when customers pick up the car keys at the sales desk. Upon the customer's arrival to the point of consumption, the customer may end up driving a different class than the one originally booked in the reservation stage for two reasons. First, the customer may receive a free car class upgrade to class $j^{UG} > j$, which we indicate with $m_{it}^{UG} = 1$. Second, the customer may accept an upsold class $j^{US} > j$, which we indicate with $m_{it}^{US} = 1$. Therefore, the customer's past consumption can be expressed in a similar manner as our earlier Monte Carlo analysis:

$$c_{it-1} = (1 - m_{it-1}^{UG})(1 - m_{it-1}^{US})d_{it-1} + m_{it-1}^{UG} j^{UG} + m_{it-1}^{US} j^{US}$$

Based on this specification, it is clear that $c_{it-1} \neq d_{it-1}$ is possible. This specification suggests potential endogeneity in the past consumption c_{it-1}. Elements that are endogenous include m_{it-1}^{UG} and m_{it-1}^{US}. To address this endogeneity issue, we employ a limited information maximum likelihood approach along the lines of Villas-Boas and Winer (1999).

The first source of endogeneity comes from upgrades, as the description of our data reveals that they may be targeted. One assumption we make here is that once customers receive a free class upgrade option, we assume that they accept doing so allows them to drive a higher quality car without paying a higher price. Therefore, we focus on modeling the firm's decision about whether or not to provide the free upgrade. Here, the latent payoff to the firm for providing an upgrade is defined as:

$$\Pi_{it} = \psi Z_{it} + \eta_{it}$$

In addition to the user-trip characteristics that enter into a customer's utility, the latent payoff from initiating an upgrade incorporates the total number of check-outs and net supply. Both user characteristics and supply-side conditions are then included in the vector Z_{it}. The error term here is denoted by η_{it}, which we assume to follow an i.i.d. Type I Extreme Value distribution. We denote the distribution for η_{it} that rationalizes $m_{it}^{UG} = 1$ as $f(\eta_{it})$.

The second source of endogeneity comes from upsells. Based on the institutional details from our empirical setting, we assume that some customers are presented with opportunities to be upsold. Sales representatives may induce customers to switch and pay for a higher class via some price discount, which we represent as μX_{it}. As certain types of customers appear more likely to receive and accept upsells, we allow the price discount benefit to be a function of observable user-trip characteristics. We now discuss how the distribution that rationalizes $m_{it}^{US} = 1$ can be recovered. Note that at the reservation stage, class j was chosen over class j^{US} as $U_{ijt} > U_{ij^{US}t}$; but at the point of consumption, the sales agent's marketing efforts may lead to $U_{ij^{US}t}^{US} > U_{ijt}$ where:

$$U_{ij^{US}t}^{US} = \mu X_{it} + \alpha_{j^{US}} + \beta X_{it} + \gamma s_{ij^{US}t} + \omega_{ij^{US}} + \varepsilon_{ij^{US}t}$$

Notice here that the main difference between $U_{ij^{US}t}^{US}$ and $U_{ij^{US}t}$ is the term μX_{it}. Thus, $m_{it}^{US} = 1$ would hold provided that the condition $U_{ij^{US}t}^{US} > U_{ijt}$ is satisfied:

$$U_{ij^{US}t}^{US} - U_{ijt} = \mu X_{it} + (\alpha_{j^{US}} - \alpha_j) + \gamma(s_{ij^{US}t-1} - s_{ijt-1}) + (\omega_{ij^{US}} - \omega_{ij}) + (\varepsilon_{ij^{US}t} - \varepsilon_{ijt}) > 0$$

4.3.3 Econometric Specification

With the consumer choice model, along with the data generating processes for upgrading and upselling decisions, we can now specify the likelihood for structural estimation. The likelihood function is therefore written as:

$$\begin{aligned} &L(\{\alpha_j\}_{\forall j}, \beta, \gamma, \sigma, \psi, \mu) \\ &= \prod_{t=1}^{T}\prod_{j=1}^{J}\prod_{\forall j^{US} > j} f(\eta_{it}) \int \{ f(\varepsilon_{ijt}, \varepsilon_{ij^{US}t} \mid \omega_{ij}, \omega_{ij^{US}}) g(d_{it} \mid \eta_{it}, \varepsilon_{ijt}, \varepsilon_{ij^{US}t}, \omega_{ij}) \} d\omega_{ij} d\omega_{ij^{US}} \end{aligned}$$

The term $f(\varepsilon_{ijt}, \varepsilon_{ij^{US}t} \mid \omega_{ij}, \omega_{ij^{US}})$ is the joint probability density function implied by the upsell acceptance decisions by customers. This joint probability density function is conditional on unobserved heterogeneity as a customer ultimately decides whether or not to accept the upsell. Finally, the car class choice decision is captured by $g(d_{it} \mid \eta_{it}, \varepsilon_{ijt}, \varepsilon_{ij^{US}t}, \omega_{ij})$, which can be written as:

$$
\begin{aligned}
& g(d_{it} \mid \eta_{it}, \varepsilon_{ijt}, \varepsilon_{ij^{US}t}, \omega_{ij}) \\
& = \prod_i \frac{\exp(\alpha_j + \beta X_{it} + \gamma s_{ijt} + \omega_{ij} + \varepsilon_{ijt})}{\sum_k \exp(\alpha_k + \beta X_{it} + \gamma s_{ikt} + \omega_{ik} + \varepsilon_{ikt})}
\end{aligned}
$$

To estimate the likelihood, we turn to simulated maximum likelihood (SML), which allows us to integrate out the unobserved heterogeneity terms.

4.4 Main Estimates

Given the model above, we consider two different specifications. To highlight the importance of variation in past upgrades, we compare the state dependence estimates across two samples: (1) the entire sample of transactions and (2) sub-sample of observations that exclude customers who received two or more free upgrades previously.

Table 6 Key Estimates from the Structural Model

	Full sample		Sub-sample	
	Estimate	SE	Estimate	SE
State dependence (γ)	0.620***	(0.144)	2.249***	(1.021)
Unobserved heterogeneity (σ)	0.932***	(0.120)	0.935***	(0.225)
Controls	Yes		Yes	
Random effects	Yes		Yes	
Observations	96209		47035	

Table 6 highlights the estimated state dependence and heterogeneity parameters.[20] In both cases, unobserved heterogeneity is present and the estimated variance for unobserved heterogeneity is similar. However, the structural state dependence effects are exaggerated when we exclude customers who received two or more free upgrades. These empirical results are consistent with our earlier Monte Carlo analysis, as inferred state dependence decreases (towards the true value) with the frequency of choice-consumption mismatches.

4.5 Economic Value of a Free Upgrade Policy

In this section, we evaluate the effectiveness of free upgrades as a promotional tool. The presence of state dependence implies that policies such as free upgrades or samples may have carry-over effects over time. Furthermore, we investigate the extent to which our evaluation of free upgrade policies is affected by biases in inferred state dependence.

For this analysis, we pick the most frequently booked class 3, and offer free upgrades to all customers who pick that class. Upgraded customers then have the opportunity to drive a class that is one level up, so the upgraded class would be 4. Given this promotion policy, we simulate the customer car class choice behavior in subsequent purchases. Combined with average prices for each car class, the simulated decisions under the various scenarios are then used to construct simulated revenues across classes.

With the counterfactual upgrade policy, we then compare the revenues without the free upgrades, to the revenues with free upgrades. Intuitively, one would expect the introduction of free upgrades would increase the revenue for class 4, while at the same time, decrease the revenue for class 3.

We then repeat this analysis using fitted model based on the sub-sample of observations which exclude customers who received upgrades in the past. Note that for comparability

[20] We refer the reader to the Appendix for a full set of estimates.

between the simulations based on full sample and sub-sample estimates, we use the same number of customers when performing these simulations. Table 7 highlights the main findings from these counterfactual simulations. The first two columns compare the revenue across scenarios without and with free upgrades using the fitted model, while the latter two columns compare the revenue across scenarios without and with free upgrades using the fitted based on the sub-sample that excludes customers that received two or more upgrades in the past. Although we do not have data on cost, policies that shift customers towards the higher classes are presumed to be profitable, as margins are most likely larger for the higher classes. Therefore, a free upgrade campaign may be profitable via its ability to induce inertial choices towards more profitable car classes.

Table 7 Economic Value of Providing Free Upgrades

	Full sample		Sub-sample	
Class	No upgrade	Upgrade	No upgrade	Upgrade
1	$17,202	$17,871	$17,477	$19,123
2	$75,556	$78,344	$68,544	$73,646
3	$243,450	$220,760	$240,300	$178,240
4	$46,213	$56,387	$41,935	$83,896
5	$1,117	$1,158	$698	$744
6	$223,830	$230,850	$253,940	$266,280
7	$9,176	$9,521	$6,744	$7,217
8	$849	$874	$643	$694
9	$5,743	$5,967	$3,306	$3,555
11	$1,405	$1,460	$1,201	$1,295
12	$39,665	$41,078	$32,770	$34,624
13	$825	$851	$41	$42
14	$585	$613	$15	$16
16	$1,516	$1,569	$1,193	$1,286
17	$27,586	$28,537	$28,565	$30,002
18	$18,844	$19,517	$18,470	$19,790
19	$1,617	$1,672	$1,403	$1,498
20	$12,462	$12,931	$10,371	$11,115
21	$2,572	$2,666	$1,292	$1,391
22	$807	$843	$642	$697
23	$635	$657	$179	$191
24	$730	$757	$870	$942
Total	$732,384	$734,884	$730,599	$736,283

As expected, revenue increases for classes 4 after the free upgrade policy, while revenues decrease for classes 3. Furthermore, Figure 5 shows that the free upgrade policy may have positive effects that carry into non-promoted classes. The reason we see such patterns is that by upgrading customers who originally picked 3 to car class 4, the policy effectively lowers the latent utility for 3 via the state dependence effect. Although most customers will be drawn to the upgraded class 4 in subsequent transactions due to state dependence, there remains a subset of them who will choose alternative classes in light of the lowered utility from consuming 3. For instance, one subset may consist of customers who originally picked class 3, but switch into classes 1 and 2 after being upgraded class 4. Alternatively, another subset may include those who switch into car classes even higher than 4, such as 6. The car rental company would benefit more from the latter subset of customers, as opposed to the former group. Notice however that the extent to which the free upgrade policy spills into other classes approaches zero as these classes move further away from 3.

Figure 5 Change in Revenues Across Classes After Free Upgrade Policy

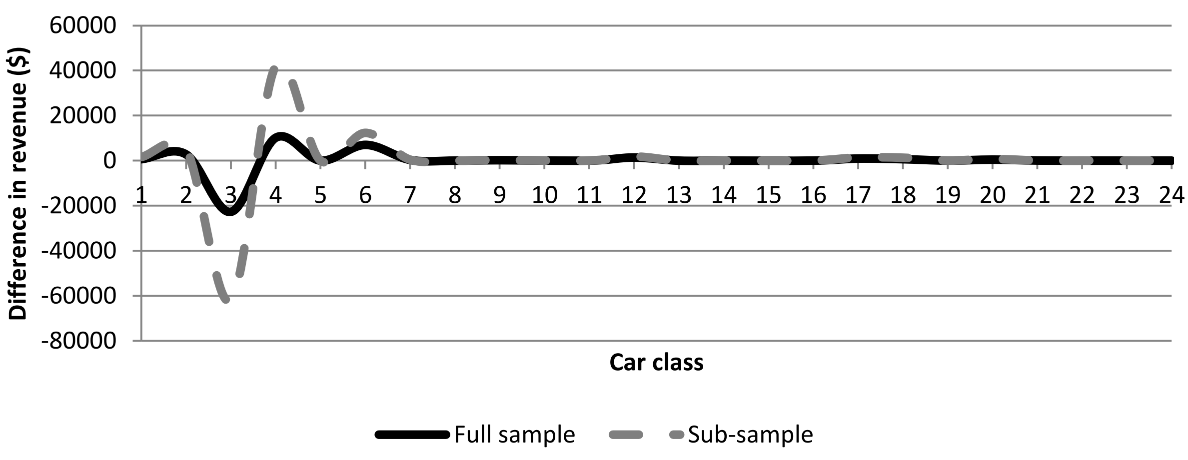

When we compare these results with those generated using the fitted model based on the sub-sample, we see that the economic benefit of free upgrades is larger in terms of revenue share gains for the higher end class 4. The increase in revenue for the upgraded class is noticeably larger than that obtained from our analysis using the full sample. This finding leads us to believe that the exclusion of choice-consumption mismatch data may result in overly optimistic assessments about the tangible benefits of free upgrade campaigns. Ultimately, these overoptimistic forecasts would lead us to pursue more promotional campaigns (that are costly) than truly warranted.

5 Conclusion

We introduce a new empirical strategy for identifying structural state dependence that exploits mismatches between choice and consumption. These mismatches help us (partially) break the correlation between past consumption and unobserved preferences, and will ultimately facilitate more optimal dynamic marketing strategies. In our Monte Carlo analysis, we demonstrate that in simulated datasets where free upgrades are frequently offered to customers, the bias in inferred state dependence can be reduced almost entirely. In contrast, existing approaches using choice set variation via price discounts is not very effective in eliminating the bias.

To apply our identification method, we estimate state dependence using data on repeat transactions from the car rental service industry. Free upgrades happen very frequently in the data, and are correlated with supply-side conditions pertaining to inventory. Such institutional features provide us an ideal environment to study and exploit mismatches between choice and consumption.

Two main results emerge from this empirical analysis. First, we confirm the presence of state dependence in a simple multinomial choice model that allows for unobserved customer-level random effects. Second, we show that inferred state dependence may be overstated if variation in past free upgrades is ignored. The second result allows us to conclude that unobserved heterogeneity is a relevant issue, and that free upgrades can serve to reduce the positive bias in inferred state dependence; thereby confirming our results from Monte Carlo analysis that state dependence is exaggerated in the absence of exclusion restrictions obtained through mismatches between choice and consumption.

Counterfactual analysis using the estimated model illustrate that the estimated level of state dependence has significant marginal effects on subsequent purchasing decisions. Furthermore, the same analysis using a sub-sample of observations that exclude users who received upgrades yields overstated effects, confirming the managerial importance of correctly disentangling state dependence and heterogeneity. Finally, we show that free upgrade campaigns can have long-run benefits; such campaigns shift purchases towards upgraded higher-end cars higher margins over the long term. But when choice-consumption mismatches are omitted in estimation of state dependence, the projections of increase in revenue shares of promoted higher-end classes are overstated.

From a practical standpoint, our new method for disentangling state dependence and unobserved heterogeneity can be applied to a variety of settings for which researchers can record as data, stated choices and actual consumption. For example, if we are using data from the service industry, we would need to know which option is reserved, and which option is actually experienced at the point of consumption. If instead we are using data from online retail, we would record which items are purchased, in addition to which items are actually delivered. Furthermore, our identification approach opens the door to experimentation strategies for managers as a means to more accurately estimate demand systems with state dependence by randomly selecting customers for free service upgrades or product switches upon shipment. Ultimately, the more accurate inferences about state dependence will not only improve dynamic advertising, marketing mix, pricing, promotion, and targeting strategies, but also provide more accurate predictions of the rate of returns from such strategies.

References

Ackerberg, D. 2003. Advertising, Learning, and Consumer Choice in Experience Good Markets: An Empirical Examination. *International Economic Review* 44, 1007-1040.

Anupindi, R., Dada, M., and Gupta, S. 1998. Estimation of Consumer Demand with Stock-Out Based Substitution: An Application to Vending Machine Products. *Marketing Science* 17, 408-423.

Arcidiacono, P., Khwaja, A., and Ouyang, L. 2012. Habit Persistence and Teen Sex: Could Increased Access to Contraception Have Unintended Consequences for Teen Pregnancies? *Journal of Business & Economic Statistics* 30, 312-325.

Barone, G., Felici, R., and Pagnini, M. 2011. Switching costs in local credit markets. *International Journal of Industrial Organization* 29, 694-704.

Bawa, K., and Shoemaker, R. 2004. The Effects of Free Sample Promotions on Incremental Brand Sales. *Marketing Science* 23, 345-363.

Biewen, M. 2009. Measuring State Dependence in Individual Poverty Histories when there is Feedback to Employment Status and Household Composition. *Journal of Applied Econometrics* 24, 1095-1116.

Biyalogorsky et. al. 1999. Research Note: Overselling with Opportunistic Cancellations. *Marketing Science* 18, 605-610.

Biyalogorksy et. al. 2005. The Economics of Service Upgrades. *Journal of Service Research* 7, 234-244.

Bruno, H., and Vilcassim, N. 2008. Structural Estimation with Varying Product Availability. *Marketing Science* 27, 1126-1131.

Cabral, L. 2012. Lock in and switch: Asymmetric information and new product diffusion. *Quantitative Marketing and Economics* 10, 375-392.

Che, H., Sudhir, K., and Seetharaman, P.B. 2007. Bounded Rationality in Pricing Under State-Dependent Demand: Do Firms Look Ahead, and if So, How Far? *Journal of Marketing Research* 44, 434-449.

Chintagunta, P. 1998. Inertia and Variety Seeking in a Model of Brand-Purchase Timing. *Marketing Science* 17, 253-270.

Chintagunta, P. 1999. Variety Seeking, Purchase Timing, and the “Lighting Bolt” Brand Choice Model. *Management Science* 45, 486-498.

Coelli, M., Green, D., and Warburton, W. Breaking the cycle? The effect of education on welfare receipt among children of welfare recipients. *Journal of Public Economics* 91, 1369-1398.

Collier, J. and Bienstock, C. 2006a. How Do Customers Judge Quality in an E-tailer. *MIT Sloan Management Review* 48, 35-40.

Collier, J. and Bienstock, C. 2006b. Measuring Service Quality in E-Retailing. *Journal of Service Research* 8, 260-275.

Conlon, C. and Mortimer, J. 2010. Effects of Product Availability: Experimental Evidence. NBER Working Paper No 16506.

Conlon, C. and Mortimer, J. 2013. Demand Estimation Under Incomplete Product Availability. *American Economic Journal: Microeconomics* 5, 1-30.

Cosguner, K., Chan, T., and Seetharaman, P.B. 2012. A Structural Econometric Model of Dynamic Manufacturer Pricing: A Case Study of the Cola Market. Working paper, Washington University of St. Louis.

Diels, J., Wiebach, N., and Hildebrandt, L. 2013. The impact of promotions on consumer choices and preferences in out-of-stock situations. *Journal of Retailing and Consumer Services* 20, 587-598.

Dube et. al. 2008. Category Pricing with State-Dependent Utility. *Marketing Science* 27, 417-429.

Dube, J-P., Hitsch, G., and Manchanda, P. 2005. An Empirical Model of Advertising Dynamics. *Quantitative Marketing and Economics* 3, 107-144.

Dube, J-P., Hitsch, G., and Rossi, P. 2009. Do Switching Costs Make Markets Less Competitive? *Journal of Marketing Research* 46, 435-445.

Dube, J-P., Hitsch, G., and Rossi, P. 2010. State dependence and alternative explanations for consumer inertia. *RAND Journal of Economics* 41, 417-445.

Erdem, T., and Keane, M. 1996. Decision-making Under Uncertainty: Capturing Dynamic Brand Choice Processes in Turbulent Consumer Goods Markets. *Marketing Science* 15, 1-20.

Erdem, T., and Sun, B. 2001. Testing for Choice Dynamics in Panel Data. *Journal of Economics & Business Statistics* 19, 142-152.

Freedman, S., Kearney, M., and Lederman, M. 2012. Product Recalls, Imperfect Information, and Spillover Effects: Lessons from the Consumer Response to the 2007 Toy Recalls. *Review of Economics and Statistics* 94, 499-516.

Freimer, M., and Horsky, D. 2012. Periodic Advertising Pulsing in a Competitive Market. *Marketing Science* 31, 637-648.

Freimer, M., and Horsky, D. 2008. Try It, You Will Like It – Does Consumer Learning Lead to Competitive Price Promotions? *Marketing Science* 27, 796-810.

Fink, A., and Reiners, T. 2006. Modeling and solving the short-term car rental logistics problem. *Transportation Research Part E* 42, 272-292.

Giarda, E. 2013. Persistency of financial distress amongst Italian households: Evidence from dynamic models for binary panel data. *Journal of Banking & Finance* 37, 3425-3434.

Goldfarb, A. 2006a. State dependence at internet portals. *Journal of Economics and Management Strategy* 15, 317-352.

Goldfarb, A. 2006b. The medium-term effects of unavailability. *Quantitative Marketing and Economics* 4, 143-171.

Greenstein, S. 1993. Did Installed Base Give an Incumbent any (Measurable) Advantages in Federal Computer Procurement? *RAND Journal of Economics* 24, 19-39.

Gregg, D. and Scott, J. 2008. A Typology of Complaints About eBay Sellers. *Communications of the ACM* 51, 69-74.

Gupta, S., Chintagunta, P., and Wittink, D. 1997. Household heterogeneity and state dependence in a model of purchase strings: Empirical results and managerial implications. *International Journal of Research in Marketing* 14, 341-357.

Halbheer et. al. 2013. Choosing a Digital Content Strategy: How Much Should be Free? *International Journal of Research in Marketing*, forthcoming.

Handel, B. 2013. Adverse Selection and Inertia in Health Insurance Markets: When Nudging Hurts. *American Economic Review* 103, 2643-2682.

Haunschild, P., and Rhee, M. 2004. The Role of Volition in Organizational Learning: The Case of Automotive Product Recalls. *Management Science* 50, 1545-1560.

Heckman, J.J. "Heterogeneity and State Dependence." In Sherwin Rose, ed., *NBER Studies in Labor Markets*. University of Chicago Press, 1981.

Israel, M. 2005a. Tenure Dependence in Consumer-Firm Relationships: An Empirical Analysis of Consumer Departures from Automobile Insurance Firms. *RAND Journal of Economics* 36, 165-192.

Israel, M. 2005b. Services as Experience Goods: An Empirical Examination of Consumer Learning in Automobile Insurance. *American Economic Review* 95, 1444-1463.

Jing, X., and Lewis, M. 2011. Stockouts in Online Retailing. *Journal of Marketing Research* 48, 342-354.

Keane, M. 1997. Modeling Heterogeneity and State Dependence in Choice Behavior. *Journal of Business & Economic Statistics* 15, 310-327.

Liu, X., Derdenger, T., and Sun, B. 2013. An Empirical Analysis of Consumer Purchase Behavior of Base Products and Add-ons Given Compatibility Constraints. Working paper, CMU.

Mahajan, V., and Muller, E. 1986. Advertising Pulsing Policies for Generating Awareness for New Products. *Marketing Science* 5, 89-106.

Marsh, T., Schroeder, T., and Mintert, J. 2006. Impacts of meat product recalls on consumer demand in the USA. *Applied Economics* 36, 897-909.

McAlister, L. 1982. A dynamic attribute satiation model for choices made across time. *Journal of Consumer Research* 9, 141-150.

Naik, N., and Moore, M. 1996. Habit Formation and Intertemporal Substitution in Individual Food Consumption. *Review of Economics and Statistics* 78 321-328.

Osborne, M. 2011. Consumer learning, switching costs, and heterogeneity: A structural examination. *Quantitative Marketing and Economics* 9, 25-70.

Paulson, K. 2012. Identification of structural state dependence in brand choice panel data. Working paper, UCSD.

Paulson, K. 2011. Using Conditioning Covariates to Identify the Dynamic Binary Choice Model. Working paper, UCSD.

Pauwels, K., and Weiss, A. 2008. Moving from Free to Fee: How Online Firms Market to Change Their Business Model Successfully. *Journal of Marketing* 72, 14-31.

Pavlidis, P., and Ellickson, P. 2012. Switching Costs and Market Power Under Umbrella Branding. Working paper, University of Rochester.

Prowse, V. 2012. Modeling Employment Dynamics with State Dependence and Unobserved Heterogeneity. *Journal of Business & Economic Statistics* 30, 411-431.

Scott, C. 1976. The Effects of Trial and Incentives on Repeat Purchase Behavior. *Journal of Marketing Research* 13, 263-269.

Seetharaman, P.B. 2004. Modeling Multiple Sources of State Dependence in Random Utility Models: A Distributed Lag Approach. *Marketing Science* 23, 263-271.

Shin, S., Misra, S., and Horsky, D. 2012. Disentangling Preferences and Learning in Brand Choice Models. *Marketing Science* 31, 115-137.

Terui, N., Ban, M., and Allenby, G. 2011. The Effect of Media Advertising on Brand Consideration and Choice. *Marketing Science* 30, 74-91.

Thomadsen, R. 2012. The Impact of Switching Stores on State Dependence in Brand Choice. Working paper, UCLA.

Vaidyanathan, G. and Devaraj, S. 2008. The role of quality in e-procurement performance: An empirical analysis. *Journal of Operations Management* 36, 407-425.

Van Heerde, H., Helsen, K., and Dekimpe, M. 2007. The Impact of Product-Harm Crisis on Marketing Effectiveness. *Marketing Science* 26, 230-245.

Villas-Boas, J.M., and Winer, R. 1999. Endogeneity in Brand Choice Models. *Management Science* 45, 1324-1338.

Wangenheim, F., and Bayon, T. 2007. Behavioral Consequences of Overbooking Service Capacity. *Journal of Marketing* 71, 36-47.

A Additional Details about Estimates

Table 8 Complete Set of Model Estimates from the Empirical Application

	Full sample		Sub-sample	
	Estimate	SE	Estimate	SE
State dependence	0.620	0.144	2.249	1.021
Variance for unobserved heterogeneity	0.932	0.120	0.935	0.225
Customer car class decision				
Airport	0.631	0.018	0.500	3.480
Phone reserve	0.118	0.007	0.231	1.656
Business	0.041	0.121	0.146	3.732
Weekend	0.924	0.184	0.722	4.565
Duration	0.200	2.042	-0.034	2.179
# transactions	-0.003	8.857	0.000	1.168
Price	-0.012	0.417	-0.051	0.110
Age	0.178	0.380	-0.111	2.604
Tier	0.620	0.144	2.249	1.021
Upgrade decision				
Airport	-0.932	0.120	-0.935	0.225
Phone reserve	-0.031	0.003	-0.037	0.000
Business	0.917	0.000	0.900	0.001
Weekend	-0.316	0.002	-0.254	0.000
Duration	0.703	0.000	0.631	0.001
# transactions	0.341	0.010	0.040	0.003
Price	0.837	0.825	-0.875	0.000
Age	-0.820	0.269	-0.255	0.028
Tier	0.884	0.003	0.603	0.001
# check-out	0.954	0.504	1.406	0.139
Net supply	-5.695	1.348	-5.933	0.093
Upsell decision				
Airport	0.006	0.020	0.007	0.036
Phone reserve	-0.450	0.042	-0.458	0.076
Business	-0.190	0.021	-0.211	0.039
Weekend	-0.027	0.019	-0.040	0.034
Duration	0.001	0.002	-0.001	0.003
# transactions	0.014	0.002	0.041	0.003
Price	0.000	0.000	0.000	0.000
Age	-0.005	0.001	-0.006	0.001
Tier	0.023	0.009	0.057	0.016