A Framework for Empirical Models of Dynamic Demand*

Gautam Gowrisankaran
University of Arizona
HEC Montréal
NBER

Marc Rysman
Department of Economics
Boston University

July 24, 2020

Abstract

This paper provides a framework for modeling dynamic demand. A recent literature estimates the dynamics of consumer demand in a variety of different contexts, including durable goods, storable goods, and subscription services. To make these estimation problems tractable, researchers typically restrict consumer expectations using an *Inclusive Value Sufficiency* (IVS) assumption. Some versions of IVS make assumptions on flow utilities and others on continuation values. A flow utility assumption is more palatable because flow utilities map closely to fundamentals. However, we show that a flow utility assumption requires a further assumption that there are sets of products with common continuation values. We argue that these two sets of assumptions result in different economic implications and hence are suitable in different contexts. We also show that most models employ an assumption on the separability of utility from holdings from utility for new purchases, and we discuss this restriction. Finally, we categorize the assumptions made by various empirical models within our unified framework.

*We wish to thank Igal Hendel, Robin Lee, Aviv Nevo, Pasquale Schiraldi and Philipp Schmidt-Dengler and participants at the 2013 Demand Estimation and Modeling Conference for helpful discussions on these topics. NSF Grants SES-0551348 and SES-0922540 provided support.

1 Introduction

For many classes of goods, consumer behavior is inherently dynamic. A recent set of empirical papers studies demand in a variety of dynamic contexts and has made progress on modeling and computational issues as well as on empirically understanding industry dynamics (see, for instance, Melnikov, 2013; Hendel & Nevo, 2006; Gowrisankaran & Rysman, 2012; Erdem, Imai & Keane, 2003). Generally, these papers have studied one of three classes of products: (1) durable goods, such as consumer electronics, in which products tend to be long-lived and the industry tends to be rapidly evolving; (2) storable goods, such as laundry detergent or ketchup, in which consumers may stockpile during periods of low prices; and (3) subscription services, such as health insurance, banking services, and telecommunication services, where consumers typically have large switching costs. In all three classes of models, consumers behave dynamically, for the different reasons of durability, storability, and switching costs, respectively.

These papers explicitly address the fact that consumer decisions have consequences for the future, and hence they solve dynamic programming models as part of their estimation and/or counterfactual simulations. Each such model requires a description of how consumers form expectations of future industry environments. In the absence of further restrictions, computation of consumer decisions for these models is often an infeasibly complex problem due to a high dimensional state space. The high dimensionality of the state space occurs for at least two reasons. First, there are often many differentiated products in the market. For instance, Gowrisankaran & Rysman (2012) observe more than 90 camcorder models in some months; Schiraldi (2011) sees more than 200 automobile models; and Erdem, Imai & Keane (2003) observe 15 ketchup brands each week. Many cases have substantial entry and exit, implying an even larger state space than with a fixed number of products. Second, even when there are few products and the set of products is stable, there are often multiple product characteristics that evolve over time. For example, Lee (2013) studies the choice between relatively few video game consoles, but each console has several evolving characteristics, notably price and software availability, while Shcherbakov (2016) studies subscription television service, where packages are distinguished by both price and channels, which evolve over time. Ho (2015) studies large banks in China, which vary in evolving quality and price. Moreover, if firms are dynamic oligopolists, price and potentially other characteristics may be determined by a large set of attributes such as the preferences of consumers who have bought each product; each of these attributes would then be a state variable.

To mitigate computational complexity, researchers restrict the sets of information that consumers use to form expectations and make decisions. This in turn restricts the state space of the Bellman equation. Using a simplified state space, one could estimate the parameters of these model with a nested fixed point maximum likelihood algorithm, as proposed by Rust (1987), or with other methods. As the modeling of dynamic demand models moves forward, we believe that it is useful to provide a framework with which to compare the assumptions on dynamic transitions across different papers. This paper makes two related contributions along these lines. First, we provide a unifying framework within which to examine the assumptions on transition dynamics across dynamic demand models. Our model nests storable goods, durable goods, and subscription services in a single model. Second, this paper uses the unifying framework to explain the differences and similarities across different papers in the assumptions used to make the computational problems tractable. The ultimate goal of this paper is to provide guidance for future research on this topic on how to simplify state spaces.

While this paper focuses on state space simplifications, many papers in this literature combine state space simplifications with other techniques that simplify computation further. For instance, Melnikov (2013) and Hendel & Nevo (2006) limit heterogeneity across consumers in their models, which allows them to estimate most of their parameters without imposing the dynamic model. Gowrisankaran & Rysman (2012) compute the Bellman equation jointly with an inversion to recover mean product characteristics (that follows Berry, Levinsohn & Pakes, 1995), which allows them to estimate a model with endogenous prices without fully specifying the firm side. We do not address these simplifications, except to note that they are often equally useful to state space simplifications in making dynamic estimation tractable.

Our main simplifying assumptions are based on what Gowrisankaran & Rysman (2012) call *Inclusive Value Sufficiency* (IVS).[1] The general idea of the IVS assumption is as follows. Rather than track the market in characteristic space, the consumer tracks the market in utility space, which is naturally a scalar variable. The researcher defines an *inclusive value* that captures the expected value of choosing the preferred option from a set of products, and then assumes that the consumer predicts future inclusive values using only current inclusive values, not the full set of product characteristics.[2]

[1] To our knowledge, Melnikov (2013) is the first paper to use this assumption for durable goods, although with a less formal treatment. Note that early versions of Melnikov's paper appeared by 2000. Another early use of this idea is Hendel & Nevo (2006).

[2] Describing the prediction of future inclusive values naturally connotes a concept of expectation and uncertainty. However, the inclusive value approach is also amenable settings with perfect foresight. Indeed, assuming perfect foresight often reduces the computational burden. See Gowrisankaran & Rysman (2012), Sun & Ishihara (2019) and Ho, Rysman & Wang (2020).

Variants of the IVS assumption appear in many empirical papers. We show that there are important differences in the IVS assumptions across papers that change both the nature of the restrictions on economic processes and the level of computational difficulty. This paper divides IVS assumptions into those on flow utilities and those on continuation values. Both assumptions are commonly used. For instance, Hendel & Nevo (2006) define the inclusive value over the flow utilities of current products. Gowrisankaran & Rysman (2012) define the inclusive value to include flow utilities as well as the entire continuation value.[3]

We argue that the two approaches have benefits and drawbacks relative to each other. IVS assumptions on continuation values are generally less attractive than those on flow utilities because continuation values encapsulate future optimizing behavior, and hence are endogenous to the model. However, we show that an IVS assumption on flow utilities requires a further assumption in order to reduce the dimensionality of the state space that encapsulates market information. In particular, a reduction to a dimension of G requires that products can be ex-ante divided into G groups with the same continuation value for all products within each group. Thus, although it is more attractive to make simplifying assumptions on fundamental characteristics, doing so requires a potentially restrictive additional assumption on how product characteristics affect the value of holdings.

An assumption that products have the same future value may be more or less problematic than an assumption on endogenous continuation values, depending on the economic context. For instance, if the researcher would like to use a single inclusive value defined over flow utilities, she must assume that product characteristics do not affect dynamic decision-making, such as the time until an upgrade. The researcher can obtain flexibility by breaking products in multiple groups. However, computational cost goes up quickly in the number of state variables or groups, so this approach is limited. Thus, the specification of the IVS assumption involves a tradeoff between making an assumption on continuation values (which are endogenous), or restricting the way in which product characteristics affect continuation values. We consider a number of papers in this context, and highlight how the choice they make reflects institutional considerations and the goals of the research project.

We also highlight a second assumption that we call *separability of previous purchase*, which has been implicit and perhaps unrecognized in most previous papers, but also contributes to reducing the size of

[3]More recent applications of models using IVS assumptions are Liu, Li & Shen (2020) in the context of automobiles and Ho et al. (2020) studying performance goods, in particular movies in China. Sun & Ishihara (2019) show how to estimate such a model with an MCMC algorithm.

the space. This assumption implies that consumer holdings do not impact the consumer's relative values of products on the market. Without this assumption, the researcher would either have to assume that the consumer tracks different inclusive values for different holdings or would have to assume that these inclusive values are the same across different holdings.

Our paper builds on a substantial recent literature on dynamic demand models. Most of these papers are empirical in nature, and hence specify an assumption on expectations, but do not explicitly consider how their assumption fits within the broader class of possible assumptions. The closest paper to ours is Aguirregabiria & Nevo (2010), which provides a survey of the recent empirical literature on dynamic models. Their paper also addresses the choice of whether or not to incorporate continuation values into IVS in durable and storable goods models. Our paper builds on that work by providing a more formal treatment, by highlighting the role of separability, and by presenting a unified framework for durable goods, storable goods, and subscription services.

The remainder of the paper is divided as follows. Section 2 provides a model. Section 3 discusses assumptions on expectations. Section 4 provides examples of papers that use each assumption. Section 5 concludes.

2 Model

2.1 General framework

We start with a discrete choice framework that nests the features found in storable goods, durable goods, and subscription goods models. Specifically, our framework allows for (1) the consumer to own multiple products; (2) the marginal utility of a product to be affected by the other products that the consumer owns; (3) products to be consumed through usage; (4) attributes of available products to change over time; and (5) the consumer to face switching costs or upgrading costs to buying a new product.

We consider the case of a single consumer.[4] Time is infinite and discrete, indexed by t. At time t, the consumer first chooses whether to add to her holdings by purchasing one of the available products, or to purchase no product. At this time, there are J_t products available for purchase. Each product has a vector of characteristics z_{jt}. The outside option of keeping her current product, which we denote option

[4] Many applications allow variation in consumer preferences across consumers through random coefficients. While we do not discuss aggregation across consumers, our framework can be combined with random coefficients.

0, also has a vector of characteristics, z_{0t}. Following the purchase decision, the consumer chooses how much to consume from her holdings. Based on her purchase and consumption decisions, her holdings stochastically transition to a new level at time $t+1$.

Finally, we partition the products into a set of mutually exclusive groups, $g = 1, \ldots, G$. Let $g(j)$ denote the group of product j and J_g denote the set of products in group g. Ultimately, we will use our groups to reduce the state space by making assumptions that products within a group are similar when applied to the dynamic problem.

We now discuss in turn the state space, consumption decision, flow utility, and optimization.

State space

The consumer's state space is divided into three parts: the market information, her holdings, and her preference shocks. We now detail each of these elements.

First, we let Ω_t denote the market information that the consumer could use for decision-making, which incorporates the set of products available on the market and their characteristics, as well as any other information that might impact future pricing or consumption patterns. We assume that Ω_t follows a Markov process, and that there are a large number of consumers, so that the process is independent of the consumer's own decisions. Note that we do not allow the characteristics of the outside option, z_{t0}, to enter into Ω_t.

Second, we let x_t denote holdings. In principle, x_t could equal the characteristics z_{js} of all products j purchased at times $s < t$, but then the dimensionality of the problem would be large, and would increase with the number of previous purchases. For this reason, researchers typically seek to reduce the computational cost of estimation by describing holdings with a smaller dimensionality. Along these lines, we assume that x_t is an N-dimensional vector. The elements of x_t could denote the amount of a good that the consumer has in inventory, the remaining lifespan of the good, and/or other product characteristics that could affect future consumption and utility.

Finally, the consumer is endowed with a preference shock for each of the J_t available choices as well as for option 0, the option of no purchase. Let $\mathcal{E}_t \equiv (\varepsilon_{0t}, \ldots, \varepsilon_{J_t t})$ denote the $J_t + 1$ vector of preference shock draws and $\mathcal{E}_t$ be distributed according to some distribution with pdf $p(.)$. We assume that $p(.)$ is *i.i.d.* across time periods. At the beginning of time t, the consumer knows the value of $\mathcal{E}_t$ but knows nothing about future shocks (besides their distribution).

A standard assumption is that each ε_{jt} is *i.i.d.* and follows the type 1 extreme value distribution, which leads to convenient closed-form solutions for several of the expectations in this paper. However, these assumptions are not necessary for the results in this paper, so we do not impose them generally.

Consumption and flow utility

We let C_t represent the consumption choice in period t. The dimensionality of C_t is N, which is the same as the dimensionality of holdings. In other words, the consumer chooses how much to consume of each type of product that she holds.

We assume that flow utility is additive in a time-invariant component and the preference shock $\mathcal{E}_t$. Thus, the current flow utility from purchasing $z_{jt}, j \geq 0$ and consuming C_t is $f(\Omega_t, x_t, z_{jt}, C_t) + \varepsilon_{jt}$. Note that characteristics z_{jt} and x_t enter in addition to C_t. That is both because it may matter which product a consumer uses, and also, particularly for services, because a consumer may draw value from holding a product even if she does not consume it (for example, insurance). The flow utility implicitly allows for a disutility of price, which can be considered one of the product characteristics. In durable and storable goods, option 0 will typically not incur a price. For subscription goods, the outside option would correspond to keeping the existing product, which might still then incur a price, e.g., the monthly cost of a cable subscription.

Industry evolution

We now discuss industry evolution. The characteristics in x_{t+1} will change over time in response to purchase z_{jt} and may change over time exogenously or in response to consumption at time t. For instance, the characteristics could be the features of digital cameras, such as megapixels, that stay constant. Alternatively, characteristics could capture the utility service from a car, which would be expected to degrade over time, perhaps stochastically, and degrade further with consumption (which would be miles driven). In a storable goods setting, x_t would be the quantity of the product that is remaining.[5] For instance, when the consumer first purchases the good, it is equal to the package size, such as 64 ounces of detergent. Using the product generates flow utility but reduces x_{t+1}. In the television subscription services market, features may evolve over time exogenously, as the cable system adds or subtracts channels and changes rates. To capture all these effects, we define the stochastic transition function for dynamic holdings as a random variable that is a function of holdings, purchase, and consumption given the purchase of an available product as $x_{t+1} = \psi(x_t, z_{jt}, C)$.[6]

[5] See Gowrisankaran & Rysman (2012), Schiraldi (2011), and Hendel & Nevo (2006) respectively for examples of each of these.

[6] We are implicitly assuming here that ψ does not depend on the state variable Ω after controlling for holdings, purchase, and consumption. To understand when this assumption might be violated, consider a model of fruit as a storable good

For the case of no purchase, $x_{t+1} = \tilde{\psi}(x_t, C)$. Consistent with the definition of $\tilde{\psi}(\cdot)$, the fundamental difference between choosing one of the available products and the outside option is that if the consumer chooses the outside option, its characteristics do not affect her holdings at subsequent periods. For simplicity, we further assume that z_{0t} is the same across periods, and hence we write z_0 from now on. We could extend our model to allow for variation and serial correlation in z_{0t} by allowing for the current z_{0t} to enter as a state. For ease of notation, we combine $\psi(\cdot)$ and $\tilde{\psi}(\cdot)$ into a single case $\psi^j(\cdot)$:

$$\psi^j(x, z_j, C) = \begin{cases} \psi(x, z_j, C) & \texttt{if } j > 0, \\ \tilde{\psi}(x, C) & \texttt{if } j = 0. \end{cases}$$

Optimization

Without further simplifications, we can express the state variables for the consumer as $\{\Omega_t, x_t, \mathcal{E}_t\}$. The consumer chooses consumption and purchase to maximize the present discounted stream of utility. We define β to be the discount factor. Because the problem is stationary, we drop t and use the $'$ symbol to denote next-period values. As in Gowrisankaran & Rysman (2012) and other papers, we define $V(\Omega, x)$ to be the Bellman equation before the realization of $\mathcal{E}$. We can write:

$$V(\Omega, x) = \int \left[\max_{C, j \in \{0,\ldots,J\}} \{ f(\Omega, x, z_j, C) + \beta E \left[V(\Omega', x') \,|\Omega, x, z_j, C \right] + \varepsilon_j \} \right] p(\mathcal{E}) d\mathcal{E}, \tag{1}$$

where $x' = \psi^j(x, z_j, C)$.[7]

In our model, $\mathcal{E}$ affects consumption only through purchase, so optimal consumption does not depend directly on $\mathcal{E}$. We let $C^j(\Omega, x, z_j)$ denote the dynamically optimal consumption decision following the purchase of product j or no purchase, so $(j = 0, \ldots, J)$. The optimal consumption decision is:

$$C^j(\Omega, x, z_j) = \arg\max_C \{ f(\Omega, x, z_j, C) + \beta E \left[V(\Omega', x') \,|\Omega, x, z_j, C \right] + \varepsilon_j \},$$

where $x' = \psi^j(x, z_j, C)$.

For notation sake, it is helpful to define a flow utility that implicitly incorporates optimal consumption. Let $u^j(\Omega, x, z_j) = f(\Omega, x, z_j, C^j(\Omega, x, z_j))$, $j = 0, \ldots, J$. Then, the net flow utility can be defined as

where temperature was a state variable, potentially affecting demand and the state transition. A high temperature might cause more fruit to spoil, implying a different transition than with a low temperature, even conditioning on holdings, purchases, and consumption of fruit. Thus, ψ needs to include the state variable temperature as an argument. One could address this type of situation in our framework by including temperature directly as part of the holdings x.

[7] We write this equation to exploit the assumption of conditional independence as in Rust (1987), which follows from the fact that $\mathcal{E}$ is *i.i.d.* over time.

$u^j(\Omega, x, z_j) + \varepsilon_j$. The associated Bellman equation is:

$$V(\Omega, x) = \max_{j \in \{0, \ldots, J\}} \left\{ u^j(\Omega, x, z_j) + \beta E\left[V(\Omega', x') \,|\Omega, x, z_j\right] + \varepsilon_j \right\}, \tag{2}$$

where $x' = \psi^j(x, z_j, C^j(\Omega, x, z_j))$. In what follows, we do not always write out the dependence of x' on x, Ω and z_j, but it is always implicit.

2.2 Specialization to existing models

Having proposed a general model, we now explain how our work nests some previous works. Here we focus on one paper each for durable goods, storable goods, and subscription services. First, we consider a durable goods paper, Gowrisankaran & Rysman (2012), which models the purchase of digital camcorders. In this paper, the vector z_j describes the set of product characteristics, both observable to the researcher and unobservable. Consumers dynamically track the scalar flow utility of the product. Denote the flow utility of j to be $x(z_j)$. Consumers discard their old product upon purchase of their new product, so x contains the flow utility of only the most recent purchase, and thus $N = 1$. Because consumption does not depreciate the camcorder, there is no meaningful distinction between consumption and flow utility. Hence, Gowrisankaran & Rysman (2012) model flow utility but not consumption directly. We define optimal consumption levels in their setting as $C^j(x, \Omega, z_j) = x(z_j)$ for $j > 0$ and $C^0(x, \Omega, z_0) = x$, with $u^j(x, \Omega, z_j) = x(z_j)$ for $j > 0$ and $u^0(x, \Omega, z_0) = x$. Products do not depreciate and hence $\psi^j(x, z_j, C^j(x, z_j, \Omega)) = x(z_j)$ for $j > 0$ and $\psi^0(x, z_0, C^j(x, z_0, \Omega)) = x$. Finally, as we describe below, the value of purchasing on the market is characterized by a single scalar variable, so all products form a single group, implying that $G = 1$.

Hendel & Nevo (2006) model a storable good, laundry detergent. Again, let z_j describe the set of product characteristics. Hendel & Nevo keep track of one dynamic holdings variable, the quantity of detergent. Thus, x is the amount of detergent remaining, and $N = 1$. Consumers track only the quantity of the product they have left, and the benefits of any other product characteristics arrive "up front" as a form of static payoff at the time of purchase. Let $x(z_j)$ now denote the size, or quantity, of product j with $x(z_0) = 0$ (i.e., the outside option has size 0). In Hendel & Nevo, a group is defined by the product size. They observe four product sizes, so $G = 4$.

Consumption in Hendel & Nevo (2006) represents the quantity of detergent used. Because the brand purchased is not relevant at the consumption stage, we can denote optimal consumption $C^j(x, \Omega, z_j) =$

$\overline{C}(x + x(z_j), \Omega)$ for some function $\overline{C}(\cdot, \cdot)$, because z_j enters consumption in a particular way.[8] Consumption here has dynamic consequences. For instance, the individual will consume less when prices are high (all else equal) because there is a greater probability that prices will be high in the future and that she will then need to limit consumption or purchase a large quantity at a high price. The stock of laundry detergent is reduced by consumption and hence $\psi^j(x, z_j, C^j(x+x(z_j), \Omega)) = x + x(z_j) - C^j(x+x(z_j), \Omega)$.

Finally, Shcherbakov (2016) models a subscription service, cable television. Using our notation, Shcherbakov has three products: cable, satellite, and no TV. Consumers hold one of these three products at a time, discarding their previous choice when they buy a new one. Each consumer is faced with 2 choices, plus the outside option of keeping her current product.[9] Holdings x can take on one of three values, 1, 2, and 3, corresponding to no television, cable, and satellite respectively, and so $N = 1$. Similarly, there are three available products, although one cannot buy the same product that one already holds; one would instead buy the outside option to continue to hold this product. Both satellite and cable have two characteristics, quality and price, each of which evolve according to Markov processes, implying that the quality and price of both products form the aggregate state space Ω. As with Gowrisankaran & Rysman (2012), consumption does not cause the product to depreciate. Thus, we define the consumption level of each product to be its quality. The utility level of purchasing a product is the consumption level minus the switching cost; the utility level from no purchase is the consumption level minus the recurring monthly service fee. Let $x(z_j)$ now denote the product number of the product z_j. Then, $\psi^j(x, z_j) = j$ if $j > 0$ and $\psi^0(x, z_0) = x$.

3 Simplifying the state space using value functions

Without any further assumptions, our model developed in Section 2.1 will have a state space that includes every element in (Ω, x). The industry state, Ω, will typically have very large dimensionality, rendering the problem intractable. This section provides our first main result on simplifying the state space, which considers simplifications based on the value function.

A common way of simplifying the state space is to make assumptions on the *inclusive value* which is the expected maximum utility of all the options within a group.[10] In this case, a group will indicate a

[8] Hendel & Nevo also model an *i.i.d.* shock to consumption (in addition to a shock to purchase). For simplicity, we ignore this feature. Because the shock is *i.i.d.*, it does not fundamentally complicate the state space simplification here.

[9] Consumers in Shcherbakov's model can also choose the tier of cable service. We ignore this choice for simplicity.

[10] Exceptions which we discuss below are Erdem et al. (2003) and Song & Chintagunta (2003) who directly model the

set of products for which the current attributes provide a sufficient statistic with which to consider the future attributes.

With this in mind, we make two sets of definitions. First, we define the value-based inclusive value for group $g = 1, \ldots, G$, as:

$$\delta_g^v(\Omega, x) = \int \left[\max_{j \in J_g} \{ u(\Omega, x, z_j) + \beta E[V(\Omega', x') | \Omega, x, z_j] + \varepsilon_j \} \right] p(\mathcal{E}) d\mathcal{E}, \tag{3}$$

where $x' = \psi^j(x, z_j, C^j(\Omega, x, z_j))$. Let $\Delta^v(\Omega, x)$ be the G vector of values $\delta_g^v(\Omega, x)$.[11] The inclusive value gives the expected value of the opportunity to choose among the options in the group. Second, we define $\hat{\varepsilon}_g^v$ to be the difference between the realized value of choosing from group $\delta_g^v(\Omega, x)$ and the expected value. Denote by $\hat{\mathcal{E}}^v$ the vector $(\varepsilon_0, \hat{\varepsilon}_1^v, \ldots, \hat{\varepsilon}_G^v)$, which is distributed with some pdf $p_{\hat{\varepsilon}^v}(\cdot)$. Thus, the actual value the consumer realizes from choosing from group g is $\delta_g^v(\Omega, x) + \hat{\varepsilon}_g^v$.

With these definitions, we can rewrite the Bellman equation in terms of the inclusive values and the value of no purchase:

$$V(\Omega, x) = \int \max \left[\max_{g \in \{1, \ldots, G\}} \left\{ \delta_g^v(\Omega, x) + \hat{\varepsilon}_g^v \right\}, u(\Omega, x, z_0) + \beta E[V(\Omega', x') | \Omega, x, z_0] + \varepsilon_0 \right] \hat{p}(\hat{\mathcal{E}}^v) d\hat{\mathcal{E}}^v. \tag{4}$$

Our approach to simplifying the state space is to propose assumptions that allow us to redefine the first argument of the value function to be a function of $\Delta^v(x)$ rather than of Ω. This will reduce the dimensionality of this argument of the state space from the dimensionality of Ω to the number of groups G interacted with the number of elements in x, which we denote $o(x)$. We require both assumptions on consumption and utility and on the link between the current Δ^v and future values of Δ^v.

We take up the consumption and utility assumption first. The focus of this paper is on purchase behavior rather than consumption, so we simply make the assumption on consumption and utility directly:

Assumption 1 *Consumption (Cons-V): Consider two states* (Ω, x) *and* $(\tilde{\Omega}, x)$*. If* $\Delta^v(\Omega, x) = \Delta^v(\tilde{\Omega}, x)$ *then* $u^j(\Omega, x, z_j) = u^j(\tilde{\Omega}, x, z_j)$ *and* $C^j(\Omega, x, z_j) = C^j(\tilde{\Omega}, x, z_j)$.

Assumption Cons-V ensures that there cannot be elements of the state space that affect consumption or utility that do not otherwise affect purchase behavior. For instance, in the case of laundry detergent, the weather might affect how much clothing a consumer wears and thus affect how often a consumer does laundry, and through this affect the demand for laundry detergent. However, our model rules out this

reference utility of particular products.

[11] In both cases, the "v" superscript indicates that the inclusive value defined in equation (3) is based on value functions.

effect.[12] Using Assumption Cons-V, we define $\hat{u}(\Delta^v, x, z_j)$ to be the mean utility obtained with choice z_j when faced with holdings x.

We now turn to how consumers predict the continuation value from any given state. Because current inclusive values are sufficient to make current purchase choices, knowing future inclusive values would be sufficient to know the continuation value.[13] However, a consumer can use the entire state space Ω to predict future inclusive values. In order to simplify the state space, we assume that consumers predict future inclusive values based only on current inclusive values. Formally, let $p_{\Delta^v}(\cdot|\cdot)$ refer to a conditional pdf of Δ^v. Our assumption is:

Assumption 2 *Inclusive Value Sufficiency (IVS-V): Consider two states* (Ω, x) *and* $(\tilde{\Omega}, x)$*. If* $\Delta^v(\Omega, x) = \Delta^v(\tilde{\Omega}, x)$ *then* $p_{\Delta^v}(\Delta^v(\Omega', x')|\Omega, x) = p_{\Delta^v}(\Delta^v(\tilde{\Omega}', x')|\tilde{\Omega}, x)$*.*

Assumption IVS-V states that consumers use only the current inclusive values to form expectations of future inclusive values, not the entire state space. The assumption is restrictive: two consumers in markets with different market structures that generate the same inclusive values (such as one market with lower prices and less products) will have the same expectations about future values of purchase. However, the assumption may be reasonable in a variety of economics contexts. It could represent a statement about the evolution of the underlying market structure, but it could also be interpreted stating that consumers make inferences about purchase timing in a boundedly rational way.[14] Rather than track the entire state space, we assume consumers track only simple summary variables of the state space.

With IVS-V, we can replace Ω in the value function with $\Delta^v(x)$. We formalize:

Proposition 1 *Consider two states* (Ω, x) *and* $(\tilde{\Omega}, x)$*. Let Assumptions Cons-V and IVS-V hold. Then* $\Delta^v(\Omega, x) = \Delta^v(\tilde{\Omega}, x)$ *implies that* $V(\Omega, x) = V(\tilde{\Omega}, x)$*. Thus, the state space can be simplified from* (Ω, x) *to* (Δ^v, x)*.*

Proof Our approach is to prove the proposition for the case of finite horizons and then take appropriate limits to address the case of infinite horizons. To ease notation, we omit Euler's constant γ which enters

[12] Extending a model to have such a variable would be straightforward, though it would require Assumption Cons-V to be conditioned on that variable.

[13] Melnikov (2013) makes a similar point in a simplified setting. This point is related to the static concept of an inclusive value, where McFadden (1978) shows that the inclusive values for "nests" of products is sufficient to determine which nest to purchase from.

[14] In this sense, our research is related to that of Krusell & Smith (1998), who make simplifying assumptions on the state space that summarizes the market in order to simplify their computational problem.

utility every period. Consider first a model where the market ends at period $T > t$ and define $V_t^T(\Omega, x)$ to be the value function in period t when the market ends at period T. Because the inclusive value is a function of the value function, the inclusive value must be indexed in a similar way. We denote the inclusive value for group g with holdings x in period t when the final period is T as $\delta_{gt}^{vT}(\Omega, x)$, and the vector of inclusive values as $\Delta_t^{vT}(\Omega, x)$. We prove the proposition by induction.

First, we consider the base case. In period T, we can write (4) as:

$$V_T^T(\Omega, x) = \int \max\left[\max_{g\in\{1,\ldots,G\}}\left\{\delta_{gT}^{vT}(\Omega, x) + \hat{\varepsilon}_{gT}^v\right\}, u^0(\Omega, x, z_0) + \varepsilon_{0T}\right]\hat{p}^v(\hat{\mathcal{E}}^v)d\hat{\mathcal{E}}^v. \tag{5}$$

In (5), $\delta_{gT}^{vT}(\Omega, x)$ is defined as in (3) but without the continuation value, because we are in the last period. Similarly, the value of no purchase is just $u^0(\Omega, x, z_0)$, without the continuation value.

The right-hand side of (5) has two elements that depend on Ω, and we show that each element does not change value when we substitute Ω with $\tilde{\Omega}$. The first part of the right-hand side does not change value when we substitute $\tilde{\Omega}$ for Ω by the assumption of the proposition that $\Delta^v(\Omega, x) = \Delta^v(\tilde{\Omega}, x)$, which implies $\Delta_T^{vT}(\Omega, x) = \Delta_T^{vT}(\tilde{\Omega}, x)$. The second part of the right-hand side does not change because $u^0(\Omega, x, z_0) = u^0(\tilde{\Omega}, x, z_0)$ as a result of the assumption of the proposition that $\Delta^v(\Omega, x) = \Delta^v(\tilde{\Omega}, x)$ and Assumption Cons-V. Thus, $V_T^T(\Omega, x) = V_T^T(\tilde{\Omega}, x)$.

Now, we consider the inductive step. For some $t < T$, assume that for all Ω, $\tilde{\Omega}$, and x, $\Delta_t^{vT}(\Omega, x) = \Delta_t^{vT}(\tilde{\Omega}, x)$ implies that $V_t^T(\Omega, x) = V_t^T(\tilde{\Omega}, x)$. We wish to show that $\Delta_{t-1}^{vT}(\Omega, x) = \Delta_{t-1}^{vT}(\tilde{\Omega}, x)$ implies $V_{t-1}^T(\Omega, x) = V_{t-1}^T(\tilde{\Omega}, x)$.

The Bellman equation here is:

$$V_{t-1}^T(\Omega, x) =$$

$$\int \max\left[\max_{g\in\{1,\ldots,G\}}\left\{\delta_{gt-1}^{vT}(\Omega, x) + \hat{\varepsilon}_{gt-1}^v\right\}, u^0(\Omega, x, z_0) + \beta E\left[V_t^T(\Omega', x')|\Omega, x, z_0\right] + \varepsilon_{0t-1}\right]\hat{p}(\hat{\mathcal{E}}^v)d\hat{\mathcal{E}}^v.$$

The right-hand side has three elements that depend on Ω, and we show that each element does not change value when we substitute Ω with $\tilde{\Omega}$. The first term on the right-hand side does not change value because an assumption of the proposition is that $\Delta^v(\Omega, x) = \Delta^v(\tilde{\Omega}, x)$. This assumption along with Assumption Cons-V further implies $u^0(\Omega, x, z_0) = u^0(\tilde{\Omega}, x, z_0)$.

We now turn to the third element, $E\left[V_t^T(\Omega', x')|\Omega, x, z_0\right]$. The expectation over Ω' and x' depends on the transition probability $P[(\Omega', x')|\Omega, x, z_0]$. By the assumptions in the setup of the model, the distributions of Ω and x are conditionally independent, so we can write:

$$P[(\Omega', x')|\Omega, x, z_0] = P[\Omega'|\Omega, x, z_0]P[x'|\Omega, x, z_0].$$

First, we take up the distribution of x'. Because the transition function for x in this case is $x' = \psi^j(x, z_0, C)$ and C is the same under Ω and Ω' by Assumption Cons-V, we have that $P[x'|\Omega, x, z_0] = P[x'|\tilde{\Omega}, x, z_0]$. Next, rather than focus on $P[\Omega'|\Omega, x, z_0]$, we focus on $P[\Delta_t^{vT}(\Omega', x)|\Omega, x, z_0]$. We have that $P[\Delta_t^{vT}(\Omega', x)|\Omega, x, z_0] = P[\Delta_t^{vT}(\tilde{\Omega}', x)|\tilde{\Omega}, x, z_0]$ by the assumption of the proposition that $\Delta^v(\Omega, x) = \Delta^v(\tilde{\Omega}, x)$ and IVS-V. Thus, the distribution of Δ_t^{vT} and x in period t is the same under Ω_{t-1} and $\tilde{\Omega}_{t-1}$. Thus, $E[V_t^T(\Omega', x')|\Omega, x, z_0] = E[V_t^T(\tilde{\Omega}', x')|\tilde{\Omega}, x, z_0]$ by the assumption of the inductive step.

As a result, because every element of $V_{t-1}^T(\Omega, x)$ is unchanged, we have that $V_{t-1}^T(\Omega, x) = V_{t-1}^T(\tilde{\Omega}, x)$. This proves the finite horizon case.

The infinite horizon case holds because

$$V(\Omega, x) = \lim_{T\to\infty} V^T(\Omega, x) = \lim_{T\to\infty} V^T(\tilde{\Omega}, x) = V(\tilde{\Omega}, x).$$

The limit exists and the first and third equalities hold because of discounting and the fact that the characteristics are bounded. The middle equality holds as a result of the preceding inductive proof. Note that we drop the subscript t when taking the limit as T goes to ∞ because the problem becomes stationary. ■

The proof relies heavily on the simplifying assumptions we have made about the value of purchase. The challenge of the proof comes mostly from the fact that no inclusive value assumption applies to the value conditional on purchasing the outside option. An alternative modeling approach might be to apply the simplifying assumptions to the outside option as well. For instance, we could think of the outside option as a separate group, characterized by its own inclusive value, and subject the inclusive value to IVS-V. However, this approach would be unattractive because the resulting value function might no longer be consistent with optimization. For instance, Proposition 2 in Gowrisankaran & Rysman (2012), which shows that there always exists a set of product characteristics that can rationalize IVS-V, regardless of the functional form of the transition process for the inclusive values, would not apply if we treated the outside option as an inclusive value.

Proposition 1 allows us to express V and u as a function of Δ^v instead of Ω. We obtain:

$$V(\Delta^v, x) = \int \max \tag{6}$$

$$\left[\max_{g\in\{1,\dots,G\}}\left\{\delta_g^v(\Delta_g^v,x)+\hat{\varepsilon}_g^v\right\}, u\left(\Delta^v,x,z_0\right)+\beta E\left[V\left(\Delta^{v\prime},x'\right)|\Delta^v,x,z_0\right]+\varepsilon_0\right]\hat{p}(\hat{\mathcal{E}}^v)d\hat{\mathcal{E}}^v.$$

In the value case, Δ^v includes the entire future flow value from a purchase, rather than being a simple function of observables $z_1,\dots,z_J$. Hence, while (6) provides a relatively tractable approach to computing V, it does not tell us how we might match Δ^v to observable information or how we might evaluate the the density of the transition from (Δ^v,x) to $(\Delta^{v\prime},x')$.

Thus, we proceed by defining Δ^v in terms of observable variables and explaining how papers have computed this transition density. Specifically, we can use the fact that we have simplified V to write Δ^v as an implicit function of the value function, holdings, and current product characteristics. This expression is then:

$$\delta_g^v(\Omega,x)=\int\left[\max_{j\in J_g}\left\{u\left(\Delta^v,x,z_j\right)+\beta E\left[V\left(\Delta^{v\prime},x'\right)|\Delta^v,x\right]+\varepsilon_j\right\}\right]p(\mathcal{E})d\mathcal{E}. \tag{7}$$

Note that (6) and (7) together implicitly define $V(\cdot,\cdot)$ and $\delta_g^v(\cdot,\cdot)$. Gowrisankaran & Rysman (2012) solve for these functions by simultaneously iterating on both equations until convergence.

Researchers typically fit a simple AR(1) functional form for the process for δ_g^v. For instance, Gowrisankaran & Rysman (2012) estimate

$$\delta_g^v(\Omega',x')=\gamma_0+\gamma_1\delta_g^v(\Omega,x)+\nu_t \tag{8}$$

as their base specification. This specification has been extended in many ways. For instance, Gowrisankaran & Rysman (2012) experiment with including the number of products J_t as an explanatory variable in (8). Lee (2013) defines multiple inclusive values and allows inclusive values to affect the transitions of other inclusive values.

Note that Gowrisankaran & Rysman (2012) explore IVS-V and the AR(1) functional form assumption in their context and show that it fits the empirical data well and that relaxations of the assumption do not have a strong effect on results, perhaps because in their data this assumption captures the important fact that Δ^v rises substantially over time in their setting. In some applications, we may observe a state variable that does not fit naturally into a product characteristic that affects product utility. For instance, a consumer may observe an indicator for whether prices will fall in the future, but has no effect today. In that case, we could change the notation for the value function and inclusive value to condition on this state variable separately from the rest of Ω and x and the results still hold.

Finally, note that our results do not impose the familiar assumption that $p_\varepsilon(.)$ follow the Type 1 extreme value distribution. However, with this assumption, as is well known from papers such as

McFadden (1981) and Rust (1987), we can write the Bellman equation as:

$$V(\Delta^v(x), x) = \ln\left(\left(\sum_{g=1}^{G} \exp\left(\delta_g^v(\Delta^v(x), x)\right)\right) + \exp\left(u\left(\Delta^v(x), x, z_0\right) + \beta E\left[V\left(\Delta^{v\prime}(x'), x'\right)|\Delta(x), x, z_0\right]\right)\right),$$

where

$$\delta_g^v(\Delta^v(x), x) = \ln\left(\exp\left(\sum_{j\in J_g} u(\Delta^v(x), x, z_j) + \beta E\left[EV\left(\Delta^{v\prime}(x'), x'\right)|\Delta^v(x), x, z_j\right]\right)\right).$$

These closed-form solutions are useful for shortening computation time. Most papers solve the value function via successive approximations. Important exceptions are Conlon (2012), who uses the MPEC method advocated by Judd & Su (2012) and Dube, Fox & Su (2012) in a related model, and Sun & Ishihara (2019), who use an MCMC method.

On one hand, the approach in this section is attractive because it models a complex dynamic problem with a relatively small state space. On the other hand, it makes assumptions that are strong and may not map very easily to fundamentals. Notably, Δ^v depends on endogenous decision-making in the present and future, so imposing a functional form assumption such as AR(1) might be unattractive, even if it fits well. An assumption that is directly on fundamental attributes would be more interpretable and hence preferable.

4 Simplifying the state space using flow utilities

This section provides our second main result on simplifying the state space, which considers simplifications based on flow utilities rather than the value function. In this section, we use three assumptions. Our first two assumptions have close analogs to the two assumptions in Section 3, but are based on flow utilities instead of value functions. Our proof requires a new assumption, due to the fact that the other assumptions are based only on flow utilities.

We start with two sets of definitions, analogous to our definitions in Section 3. First, for group $g = 1, \ldots, G$, the flow-based inclusive value is:

$$\delta_g^f(\Omega, x) = \int \max_{j\in J_g} \left\{u(\Omega, x, z_j) + \varepsilon_j\right\} p(\mathcal{E}) d\mathcal{E}.^{15} \tag{9}$$

Similarly to $\Delta^v(\Omega, x)$, we define $\Delta^f(\Omega, x)$ to be the G vector of values $\delta_g^f(\Omega, x)$. Second, we define $\hat{\varepsilon}_g^f$ to be the difference between the realized value of choosing from group $\delta_g^f(\Omega, x)$ and the expected value. Let

[15] The "f" superscript indicates that the inclusive value is based on flow utilities.

$\hat{\mathcal{E}}^f$ denote the vector $(\varepsilon_0, \hat{\varepsilon}_1^f, \ldots, \hat{\varepsilon}_G^f)$, which is distributed with some pdf $p_{\hat{\varepsilon}^f}(\cdot)$. Thus, the actual value the consumer realizes from choosing from group g is $\delta_g^f(\Omega, x) + \hat{\varepsilon}_g^f$.

Using these definitions, we now offer our assumption on consumption and utility, analogous to Section 3:

Assumption 3 *Consumption (Cons-F): Consider two states* (Ω, x) *and* $(\tilde{\Omega}, x)$. *If* $\Delta^f(\Omega, x) = \Delta^f(\tilde{\Omega}, x)$, *then* $C^j(\Omega, x, z_j) = C^j(\tilde{\Omega}, x, z_j)$ *and* $u(\Omega, x, z_j) = u(\tilde{\Omega}, x, z_j)$.

This assumption implies that state variables that directly affect consumption and utility must do so exclusively through the inclusive value. We next offer our assumption on the link between inclusive values from period to period:

Assumption 4 *Inclusive Value Sufficiency (IVS-F): Consider two states* (Ω, x) *and* $(\tilde{\Omega}, x)$. *If* $\Delta^f(\Omega, x) = \Delta^f(\tilde{\Omega}, x)$ *then* $p_{\Delta^f}(\Delta^f(\Omega', x')|\Omega, x) = p_{\Delta^f}(\Delta^f(\tilde{\Omega}', x')|\tilde{\Omega}, x)$.

However, for the flow utility case, our proof requires an additional assumption that we did not impose in the value approach: that purchases can affect the continuation value only in a way that is constant for all products within a group. Without this assumption, knowing the flow inclusive value for each group will not be sufficient to determine the consumer's choice of group, because the consumer will need to know which continuation value from within the group the consumer will obtain. Our assumption is:

Assumption 5 *Constant Continuation Value within a Group (CCV):*
If $g(k) = g(j)$ *then* $\psi^j(\Omega, x, z_j) = \psi^k(\Omega, x, z_k), \ \forall \ \Omega, x$.

Thus, holdings in the next period depend on the group from which the product was purchased, rather than any other product characteristics. As a result, the purchase choice affects the value function only through its group.

Under CCV, we can write the expectation of the value function as $E\left[V\left(\Omega', x'\right)|\Omega, x, g\right]$, with g replacing z_j, which allows us to write the maximum utility within group g as:

$$\max_{j \in J_g} \left\{u(\Omega, x, z_j) + \varepsilon_j\right\} + \beta E\left[V\left(\Omega', x'\right)|\Omega, x, g\right], \tag{10}$$

for $g \in \{1, \ldots, G\}$. In (10), CCV allows us to separate the maximization of products within a group from the continuation value, because the continuation value is constant within groups. We can then rewrite the Bellman equation as an outer choice over groups and an inner choice over products within the groups, plus the choice of the outside outside option. Incorporating our above two definitions of Δ^f and $\hat{\varepsilon}^g$, we obtain:

$$V(\Omega, x) = \int \max \tag{11}$$
$$\left(\max_{g=\{1,\ldots,G\}} \left\{ \delta_g^f(\Omega, x) + E\left[V\left(\Omega', x'\right)|\Omega, x, g\right] + \hat{\varepsilon}_g^f \right\}, u(\Omega, x, z_0) + \beta E\left[V(\Omega', x')|\Omega, x, 0\right] + \varepsilon_0 \right) \hat{p}(\hat{\mathcal{E}^f}) d\hat{\mathcal{E}^f}.$$

An important difference between (11) and (4)—which is the analog for the value function case—is that in (11), even when the consumer buys an inside product, we must still directly model the future value in the Bellman equation.

We now offer our main result of this section:

Proposition 2 *Let Assumptions Cons-F, IVS-F and CCV hold. Consider two states (Ω, x) and $(\tilde{\Omega}, x)$. Then $\Delta^f(\Omega, x) = \Delta^f(\tilde{\Omega}, x)$, implies that $V(\Omega, x) = V(\tilde{\Omega}, x)$. Thus, the state space can be simplified from (Ω, x) to $(\Delta^f(x), x)$.*

Proof As in Proposition 1, we prove the proposition for the case of finite horizons and then take appropriate limits to address the case of infinite horizons. As before, define $V_t^T(\Omega, x)$ to be the value function in period t when the market ends at period T. Index the inclusive value similarly, to be $\delta_{gt}^{fT}(\Omega, x)$, and the vector of inclusive values to be $\Delta_t^{fT}(\Omega, x)$. We prove the proposition by induction.

First, we consider the base case. In period T, we can write (4) as:

$$V_T^T(\Omega, x) = \int \max \left[\max_{g \in \{1,\ldots,G\}} \left\{ \delta_{gT}^{fT}(\Omega, x) + \hat{\varepsilon}_{gT}^f \right\}, u(\Omega, x, z_0) + \varepsilon_{0T} \right] \hat{p}(\hat{\mathcal{E}^f}) d\hat{\mathcal{E}^f}.$$

The right-hand side has two elements that depend on Ω, and we show that each element does not change value when we substitute Ω with $\tilde{\Omega}$. For the first element, we have the assumption of the proposition that $\Delta^f(\Omega, x) = \Delta^f(\tilde{\Omega}, x)$, which implies $\Delta_T^{fT}(\Omega, x) = \Delta_T^{fT}(\tilde{\Omega}, x)$. For the second element, the assumption of the proposition that $\Delta^f(\Omega, x) = \Delta^f(\tilde{\Omega}, x)$ and Assumption Cons-F imply $u(\Omega, x, z_0) = u(\tilde{\Omega}, x, z_0)$. Thus, $V_T^T(\Omega, x) = V_T^T(\tilde{\Omega}, x)$.

Now, we consider the inductive step. For some $t < T$, assume that for all Ω, $\tilde{\Omega}$, and x, $\Delta_t^{fT}(\Omega, x) = \Delta_t^{fT}(\tilde{\Omega}, x)$ implies that $V_t^T(\Omega, x) = V_t^T(\tilde{\Omega}, x)$. We wish to show that $\Delta_{t-1}^{fT}(\Omega, x) = \Delta_{t-1}^{fT}(\tilde{\Omega}, x)$ implies $V_{t-1}^T(\Omega, x) = V_{t-1}^T(\tilde{\Omega}, x)$.

Writing out the Bellman equation, we have:

$$V_{t-1}^T(\Omega, x) = \int \max \left[\max_{g \in \{1,\ldots,G\}} \left\{ \delta_{gt-1}^{fT}(\Omega, x) + \beta E\left[V_t^T(\Omega', x')|\Omega, x, g\right] + \hat{\varepsilon}_{gt-1}^f \right\}, \right.$$

$$u(\Omega, x, z_0) + \beta E\left[V_t^T(\Omega', x')|\Omega, x, 0\right] + \varepsilon_{0t-1}\Big]\hat{p}(\hat{\mathcal{E}^f})d\hat{\mathcal{E}^f}.$$

The right-hand side has four elements that depend on Ω, and we show that each element does not change value when we substitute Ω with $\tilde{\Omega}$. The first term on the right-hand side does not change value because an assumption of the proposition is that $\Delta^f(\Omega, x) = \Delta^f(\tilde{\Omega}, x)$. This assumption along with Assumption Cons-F further implies $u(\Omega, x, z_0) = u(\tilde{\Omega}, x, z_0)$.

We now turn to the two continuation values. The expectation over Ω' and x' depends on the transition probability $P[(\Omega', x')|\Omega, x, z_j]$. By the assumptions in the setup of the model, the distributions of Ω and x are conditionally independent, so we can write:

$$P[(\Omega', x')|\Omega, x, z_j] = P[\Omega'|\Omega, x, z_0]P[x'|\Omega, x, z_j],$$

for $j = 0, \ldots, J_t$. The transition function for x is $x' = \psi^j(x, z_j, C)$ and C is the same under Ω and Ω' by Assumption Cons-F, so we have that $P[x'|\Omega, x, z_j] = P[x'|\tilde{\Omega}, x, z_j]$ for $j = 0, \ldots, J_t$. Next, rather than focus on $P[\Omega'|\Omega, x, z_j]$, we focus on $P[\Delta_t^{fT}(\Omega', x')|\Omega, x, z_j]$. We have that $P[\Delta_t^{fT}(\Omega', x')|\Omega, x, z_j] = P[\Delta_t^{fT}(\tilde{\Omega}', x')|\tilde{\Omega}, x, z_j]$ by the assumption of the proposition that $\Delta_{t-1}^{fT}(\Omega, x) = \Delta_{t-1}^{fT}(\tilde{\Omega}, x)$ and IVS-F. Thus, the distribution of Δ_t^{fT} and x in period t is the same under Ω_{t-1} and $\tilde{\Omega}_{t-1}$. Thus, $E[V_t^T(\Omega', x')|\Omega, x, z_j] = E[V_t^T(\tilde{\Omega}', x')|\tilde{\Omega}, x, z_j]$ by the assumption of the inductive step.

As a result, because very element of $V_{t-1}^T(\Omega, x)$ is unchanged, we have that $V_{t-1}^T(\Omega, x) = V_{t-1}^T(\tilde{\Omega}, x)$. This proves the finite horizon case.

The infinite horizon case holds because

$$V(\Omega, x) = \lim_{T\to\infty} V^T(\Omega, x) = \lim_{T\to\infty} V^T(\tilde{\Omega}, x) = V(\tilde{\Omega}, x).$$

The limit exists and the first and third equalities hold because of discounting and the fact that the characteristics are bounded. The middle equality holds as a result of the preceding inductive proof. Note that we drop the subscript t when taking the limit as T goes to ∞ because the problem becomes stationary. ■

Similarly to the value case, Proposition 2 allows us to express V, δ_g^f and u as being a function of $\Delta^v(x)$ instead of Ω. We obtain:

$$V(\Delta^f(x), x) = \int \max \max_{g\in\{1,\ldots,G\}} \left\{\delta_g^f(\Delta^f(x), x) + E[V(\Delta^{f\prime}(x'), x')|\Delta^f(x), x, g] + \hat{\varepsilon}_g^f\right\}, \tag{12}$$
$$u\left(\Delta^f(x), x, z_0\right) + \beta E\left[V\left(\Delta^{f\prime}(x'), x'\right)|\Delta^f(x), x, 0\right] + \varepsilon_0\hat{p}(\hat{\mathcal{E}}^f)d\hat{\mathcal{E}}^f,$$

where

$$\delta_g^f\left(\Delta^f(x), x\right) = \int \left[\max_{j\in J_g}\left\{u\left(\Delta^f(x), x, z_j\right) + \varepsilon_j\right\}\right] p(\mathcal{E}) d\mathcal{E}.^{16}$$

Thus, we can reduce the first part of the state space from Ω to $\Delta^f(x)$. In practice, researchers again typically fit a simple AR(1) functional form for the process for Δ^f similarly to the value function case case.

As in the previous subsection, the assumption that ε_j follows the Type 1 extreme value distribution leads to the Bellman equation satisfying:

$$V(\Delta^f(x), x) =$$

$$\ln\left(\sum_{g=1}^{G} \exp\left\{\delta_g^f(\Delta^f(x), x) + \beta E[V(\Delta^{f\prime}(x'), x')|\Delta^f(x), x, g]\right\} + \right.$$

$$\left. \exp\left(u\left(\Delta^f(x), x, z_0\right) + \beta E\left[V\left(\Delta^{f\prime}(x'), x'\right)|\Delta^f(x), x, 0\right]\right)\right),$$

where

$$\delta_g^f(\Delta^f, x) = \ln\left(\sum_{j\in J_g} \exp\left(u\left(\Delta^f(x), x, z_j\right)\right)\right).$$

These solutions make numerical calculations relatively quick.

An advantage of the flow approach is that this functional form assumption is applied to an object that is defined only over flow utilities rather than flow utilities and continuation values. A disadvantage of the flow approach is that it requires that the dynamics are only over the choice across groups, not the individual products.

5 Separability of previous purchases

We have thus far have focused on inclusive values that condition on a customer's holdings x. Hence, our inclusive values have been indexed by holdings, i.e. $\Delta^v(x)$ and $\Delta^f(x)$. If x has a substantial number of elements, this might still yield a large dimensional state space in the absence of further simplifications.

As an example, consider Lee (2013), who models consumers choosing between game consoles from three different companies. Lee allows for consumers to hold any subset of consoles, implying that consumer holdings can take on 2^3 values. Lee further allows for the value of buying a console to depend

[16] As with the value case, we abuse notation slightly here because δ_g^f is a function of $(\delta_1^f, \ldots, \delta_{g-1}^f, \delta_{g+1}^f, \ldots, \delta_G^f)$ rather than of Δ^f.

on which other consoles the individual holds, but does not allow consumers who own one console (e.g., Microsoft) to repurchase a new version of the same console. Thus, Lee's consumers need to keep track of the inclusive value for each set of holdings except the empty set, yielding $2^3 - 1 = 7$ inclusive values.

To address these challenges, many papers simplify the state space further by assuming that consumer holdings do not affect the relative value of different new purchases. We denote this feature *separability of previous purchase (SPP)*. SPP is natural in many contexts. For instance, if the consumer discards her current product upon purchasing a new product, then current holdings would not affect her relative valuations of the available new products. We show below that we can exploit SPP to define an inclusive value that is a function of available products only, suppressing the dependence on holdings. The great majority of papers of which we are aware make a separability of previous purchases assumption, and derive large computational benefits from this assumption. Lee (2013) does not make such an assumption, which allows him to model the possibility that holder of console A obtains less utility from purchasing console B than she would if she did not own any console, e.g. because she values only B for games that she could not play on A.

The exact assumption required—and hence the economic implications of the assumption—differs for the flow and value cases. We employ the following assumption for the value case:

Assumption 6 *Separability of Previous Purchases (SPP-V)*

(a) $u(\Delta^v(x), x, z_j) = u_g^x(\Delta^v(x), x) + u^z(z_j)$,

(b) $\psi^j(x, z_j, C) = \psi^j(z_j)$ *if* $j > 0$.

The first part of Assumption SPP-V states that the current utility of a purchase is additive in two parts: (1) a part that is not affected by the purchase, $u_g^x(\Delta^v(x), x)$ (though it can still be affected by the purchase group g), and (2) a part that is not affected by holdings $u^z(z_j)$.

To understand this, suppose we are modeling camcorders and that all available camcorders are in the same group. The first part of Assumption SPP-V could occur if consumers discarded their existing camcorder upon purchase of a new one, in which case $u_g^x(\Delta^v(x), x) = 0$; if consumers used a second camcorder in a way that was unaffected by the purchase of a new camcorder, in which case $u_g^x(\Delta^v(x), x)$ would be a function of only x; or if consumers sold their existing camcorder upon purchase of a new one,[17] in which case $u_g^x(\Delta^v(x), x)$ would be the sales price.

[17] See Schiraldi (2011) for an example.

Suppose we divide camcorders into two groups, one with cameras with zoom and one without. Assumption SPP-V is then also consistent with camcorders with zoom adding relatively more utility than camcorders without zoom, based on the level x of holdings at the time of purchase. In this way, SPP-V allows holdings to affect the relative utility of different groups at the time of purchase, though not of different products. The reason that we can allow an interaction between holdings and the group of purchase is because we track different a separate inclusive values for each different group, which implies any flow utility term that is common across products within a group can be added to utility in the Bellman equation without adding dimensionality to the state space.

The second part of Assumption SPP-V implies that when a consumer purchases, her current holdings do not affect future holdings. It is not consistent with the purchase affecting the future state differently across products or even across groups. Together, these two parts of SPP-V imply that current holdings do not affect any of the elements of the inclusive value.

It might appear that switching costs would be difficult to model under SPP-V, because switching costs imply that holdings affect a consumer's relative utility of different goods. In practice, this has not been a major constraint. In many contexts, there may be relatively few goods, and so each good can be in a separate group. Examples are Shcherbakov (2016), who studies the choice between cable and satellite, and Ho (2015) who studies the choice among major banks in China. Even with many goods, clever redefinitions of products can allow researchers to separate the inclusive value from holdings in the value approach, such as in Schiraldi (2011).

Using SPP-V, the relative values of different purchase options are not a function of the current holdings. Thus, we have the following lemma.

Lemma 1 *Under Cons-V, IVS-V and SPP-V, we can write $\delta_g^v(x) = \hat{\delta}_g^v + u_g^x(x)$, where:*

$$\hat{\delta}_g^v = \int \left[\max_{j \in J_g} \left\{ u^z(z_j) + \beta E\left[EV(x') \,|z_j\right] + \varepsilon_j \right\} \right] p(\mathcal{E}) d\mathcal{E}. \tag{13}$$

Proof Under SPP-V, we can impose $u(\Omega, x, z_j) = u_g^x(\Omega, x) + u^z(z_j)$ in Equation 3. The term $u_g^x(\Omega, x)$ is constant across products within the group, and the elements of u_g^x do not affect transitions by the second part of SPP-V, so u_g^x can be taken out of the maximization on the right-hand side of the definition of $\delta_g^v(\Omega, x)$. Rewriting:

$$\delta_g^v(\Omega, x) = u_g^x(\Omega, x) + \int \left[\max_{j \in J_g} \left\{ u^z(z_j) + \beta E\left[V(\Omega', x') \,|\Omega, x, z_j\right] + \varepsilon_j \right\} \right] p(\mathcal{E}) d\mathcal{E}.$$

Defining $\hat{\delta}_g^v(\Omega)$ with the second part of the equation completes the proof. ■

Similarly to $\Delta^v(\Omega, x)$, define $\hat{\Delta}^v(\Omega)$ to be the G vector of values $\hat{\delta}_g^v(\Omega)$. As a result of Lemma 1, we can work with $\hat{\Delta}^v(\Omega)$, which is of dimensionality G, rather than $\Delta^v(\Omega, x)$ which is of dimensionality $G \times o(x)$. We can then restate Proposition 1 for the SPP case:

Corollary 2 *Consider two states* (Ω) *and* $(\tilde{\Omega})$. *Let Assumptions Cons-V, IVS-V and SPP-V hold. Then* $\hat{\Delta}^v(\Omega) = \hat{\Delta}^v(\tilde{\Omega})$ *implies that* $V(\Omega) = V(\tilde{\Omega})$. *Thus, the state space can be simplified from* (Ω, x) *to* $(\hat{\Delta}^v, x)$.

Proof We omit the proof for brevity, as it follows closely the proof of Proposition 1. Note that under Assumption Cons-V, we have that $\hat{\Delta}^v(\Omega) = \hat{\Delta}^v(\tilde{\Omega})$ implies $u_g^x(\Omega, x) = u_g^x(\tilde{\Omega}, x)$ for all g. ■

Because $\{\hat{\Delta}^v(\Omega), x\}$ is of lower dimensionality than $\{\Delta^v(\Omega, x), x\}$, Corollary 2 provides a reduction in dimensionality, and thus a computational advantage over Proposition 1. For instance, if Lee (2013) were to adopt SPP, then consumers would need to keep of an inclusive value for each of the three consoles rather than for each combination of consoles, reducing the dimensionality of the state space for each consumer from seven to three. With more than three potential products, SPP would generate serious numerical advantages, though at the expense of more assumptions.

It appears difficult to address storable goods under SPP-V, because in a storable goods context, current holdings typically affect future holdings. Fortunately, the flow case requires an assumption that is weaker than SPP-V. That is because in order to remove x from $\hat{\delta}_g^f()$, we must eliminate it only from current utility. It is not necessary to remove holdings from the continuation when a consumer purchases, because that is already handled by CCV. The formal assumption is:

Assumption 7 *Separability of Previous Purchases (SPP-F):* $u(\Omega, x, z_j) = u_g^x(\Omega, x) + u^z(z_j)$.

Similar to Assumption SPP-V and Lemma 1, Assumption SPP-F leads to the following result:

Lemma 3 *Under SPP-F,* $\delta_g^f(\Omega, x) = \hat{\delta}_g^f(\Omega) + u_g^x(\Omega, x)$, *where:*

$$\hat{\delta}_g^f(\Omega) = \int \max_{j \in J_g} \{u^z(z_j) + \varepsilon_j\}\, p(\mathcal{E}) d\mathcal{E}.$$

Proof Under SPP-F, we can impose $u(\Omega, x, z_j) = u_g^x(\Omega, x) + u^z(z_j)$ in Equation 9. The term $u_g^x(\Omega, x)$ is constant across products within the group, so u_g^x can be taken out of the maximization on the right-hand side of the definition of $\delta_g^f(\Omega, x)$. Rewriting:

$$\delta_g^f(\Omega, x) = u_g^x(\Omega, x) + \int \left[\max_{j \in J_g} \{ u^z(z_j) + \varepsilon_j \} \right] p(\mathcal{E}) d\mathcal{E}.$$

Defining $\hat{\delta}_g^f(\Omega)$ with the second part of the equation completes the proof. ■

For instance, the storable goods example, storage costs would be captured in $u_g^x(\Omega, x)$ and utility from new purchases would be in δ_g^f. As in the value case, Assumption SPP-F allows us to define the inclusive value over Ω and not x, which leads to a reduction in dimensionality from $G \times o(x)$ to G.

There is a further benefit, which is that the inclusive value is defined only over exogenous flow utilities. This arises because the endogenous consumption choice is separated into $u_g^x(\Omega, x)$. Models that incorporate endogenous consumption can implement Assumptions CCV and SPP-F in order to define the inclusive value over exogenous items. A leading example of such a paper is Hendel & Nevo (2006). Note that most papers do not explicitly impose SPP-V or SPP-F. Rather, the SPP assumption is implied by the modeling framework. That is the case for Hendel & Nevo (2006), and arises from their model setup that the current purchase can affect the value function only through the size of the product.[18]

Define $\hat{\Delta}^v(\Omega)$ to be the G vector of values $\hat{\delta}_g^v(\Omega)$. We can then restate Proposition 2 for the SPP case:

Corollary 4 *Consider two states* (Ω) *and* $(\tilde{\Omega})$*. Let Assumptions Cons-F, IVS-F and SPP-F hold. Then* $\hat{\Delta}^f(\Omega) = \hat{\Delta}^f(\tilde{\Omega})$ *implies that* $V(\Omega) = V(\tilde{\Omega})$*. Thus, the state space can be simplified from* (Ω, x) *to* $(\hat{\Delta}^f, x)$*.*

Proof We omit the proof for brevity, as it follows closely the proof of Proposition 2. Note that under Assumption Cons-F, we have that $\hat{\Delta}^f(\Omega) = \hat{\Delta}^f(\tilde{\Omega})$ implies $u_g^x(\Omega, x) = u_g^x(\tilde{\Omega}, x)$ for all g. ■

Melnikov (2013) and Hendel & Nevo (2006) make a further assumption that there are no random coefficients in the model, particularly on variables that do not have dynamic content. This assumption allows for another simplification beyond the use of IVS-F, which is that they can solve their model in two stages, so it is much faster to solve.[19] Another computational issue is that with aggregate data, such as in

[18] Thus, our approach represents the assumption of Hendel & Nevo (2006) that the current purchase affects the value function only through the size of the product in two separate assumptions: The first is the effect of current purchase on the continuation value (IVS-F) and the second is the effect on current utility (SPP-F).

[19] Imposing no random coefficients in Gowrisankaran & Rysman (2012) also makes the solution very quick, because under this assumption, they must solve the model only once.

Melnikov (2013) and Gowrisankaran & Rysman (2012), we typically must further introduce an inversion of market shares to compute mean utilities. Without random coefficients, Melnikov (2013) can do this step analytically as in Berry (1994). With random coefficients, Gowrisankaran & Rysman (2012) use a fixed point algorithm in the spirit of Berry, Levinsohn & Pakes (1995), which must be solved alongside the algorithm for solving Equation 3. The issues of whether to model consumer heterogeneity or whether to use aggregate data as opposed to household level data are outside the scope of this paper. We mention them here to clarify what assumptions are the sources of simplified estimation in this literature.

6 Applications of the results to empirical work

Most papers do not make all of these assumptions explicitly. In this section, we write several examples from the literature in our notation. We focus on whether they use IVS-V or CCV and IVS-F, as well as SPP-V or SPP-F. We break up our discussion into durable goods, storable goods, and subscription service goods.

6.1 Durable goods

The durable goods papers we discuss do not model consumption (C) so we ignore it here. We have already discussed Gowrisankaran & Rysman (2012). To finish the discussion, recall that they use a single partition, $G = 1$, and that $x = u^z(z_j)$. Because x is a continuous variable, it varies within the (single) group of products. Thus, their model does not satisfy CCV, and so they use IVS-V. The state space has two scalar variables, x and δ^v. They assume products are infinitely durable, so x is constant between purchases.

Schiraldi (2011) is similar. He models the second-hand automobile market, and so is particularly interested in how quality affects holding times. Thus, CCV is unattractive in his context. He makes a similar set of assumptions to Gowrisankaran & Rysman (2012), but allows x to evolve stochastically (downwards) over time. He uses IVS-V. Interestingly, Schiraldi finds it convenient to treat the durable goods problem with transaction costs as a sequence of one-period rental problems where the consumer can avoid the transaction cost by continuing to rent the good he already holds. Thus, from our perspective, Schiraldi rewrites the durable goods model with transaction costs as a subscription services model with switching costs. As discussed above, switching costs can violate SPP-V, but he rewrites the value function in a way that allows him to address this issue.

Melnikov (2013) models demand for printers. He uses a single partition of goods, $G = 1$. He assumes consumers purchase only once in their lives. Thus, he can interpret $u^z(z_j)$ as the present discounted value of the future stream of utility, and so there is no further continuation value. In his model, the value of no purchase $u^x(x, \Omega, z_0)$ is a constant to be estimated and $u^z(z_j)$ captures all future value. Naturally, assuming the continuation value is zero implies that it is constant across product choices, and so the model satisfies Assumption CCV. He uses Assumption IVS-F. Carranza (2010) uses a similar dynamic model for individual consumers to Melnikov, though with the addition of a random coefficient on the constant term.

Lee (2013) studies video game consoles. While he models only three products (consoles), they have several evolving characteristics, especially the complementary value of video games, so he still finds it useful to make an IVS assumption. His paper is an example of how IVS is valuable not only when we observe consumers choosing between many products, but also when choosing between a few products with several characteristics. He models each product as a separate group. Thus, each partition contains only one product, so he easily satisfies Assumption CCV. Thus, he uses Assumption IVS-F. As pointed out earlier, Lee does not use SPP, so his number of inclusive values is greater than his number of groups.

Gowrisankaran, Park & Rysman (2014) estimate demand for DVD players. Their partition has only one element, containing all of the products ($G = 1$). Products are infinitely durable and the characteristics of purchased products do not change over time. The paper assumes that consumer holdings depends only on the number of products purchased in the past, so product characteristics do not have dynamic content. As in Hendel & Nevo (2006), consumers obtain all of the future benefits from product characteristics "up front." In our notation, Gowrisankaran, Park & Rysman (2014) assume $N = 1$ and x is the number of products purchased. When purchasing, the consumer gets the holding value of $x + 1$ products and the "up front" value of $u^z(z_j)$. The ψ^j function is straightforward: $x' = x + d$, where d is an indicator for purchase in the current period. Naturally, an indicator for whether a consumer purchases is constant in the choice of product, so Gowrisankaran, Park & Rysman (2014) satisfy CCV. As a result, they use Assumption IVS-F.

6.2 Storable goods

We have covered Hendel & Nevo (2006) but to complete the discussion, recall that they discretize quantity into four levels, and partition the products into these four groups ($G = 4$). The state variable for consumer holdings is the totally quantity held, so Hendel & Nevo (2006) assume that the brand has no dynamic impact, but delivers up-front utility. Because the quantity purchased is the only variable that affects the continuation value and it is constant with each group of products (because groups are defined by size), this model satisfies CCV. As a result, Hendel & Nevo (2006) use Assumption IVS-F.

6.3 Subscription services

Shcherbakov (2016) studies the choice between cable and satellite television service. With only two products, he uses $G = 2$. He uses Assumption IVS-F and tracks two inclusive values. Ho (2015) takes a similar approach to banks in China, where he amalgamates banks into only four choices, which then are treated as separate elements of the partition. He also uses Assumption IVS-F. In contrast, Nosal (2012) models Medicare Advantage plans, where consumers have many choices. Both papers allow consumers to choose between the products, and a number of plan or bank characteristics affect the future flow utility from that product. Thus, their models do not satisfy Assumption CCV and they use Assumption IVS-V.

Overall, we see that researchers make these decisions based on the economic phenomena they wish to capture and computational implications. Researchers always have flexibility in this choice. For instance, Gowrisankaran & Rysman (2012) could assume that consumers track only an indicator for whether they hold a product or not, which would enable the use of Assumption IVS-F instead of IVS-V. Similarly, Hendel & Nevo (2006) could assume that consumers track only a single inclusive value for the market rather than one for each size available, but then because size has dynamic content, the model would not satisfy Assumption CCV. While this would reduce the state space, they would then require Assumption IVS-V instead of IVS-F. Alternatively, Hendel & Nevo (2006) could model utility so that more characteristics affect the continuation value, and then use IVS-F. Overall, modeling the effect of all product characteristics on holding time is important to Gowrisankaran & Rysman (2012), but Hendel & Nevo (2006) can answer their questions just using quantity for dynamic optimization, which leads them to make their particular assumptions.

6.4 Alternative approaches

Erdem, Imai & Keane (2003) estimate a dynamic model of demand for storable good, ketchup, that does not rely on IVS. They develop an approach that is quite different than ours. They seek to model the evolution of the characteristics of each product separately, rather than model the evolution of an aggregate summary statistic like the inclusive value. In order to make the problem tractable, they take advantage of specific institutional features of the ketchup market. First, they observe 15 products in the market with no entry or exit in their data set, and they observe variation in only one characteristic, price. Furthermore, the products are a set of only four brands, with each brand exhibiting one product that dominates sales, the 24 oz bottle. Thus, they treat the prices of the four 24 oz bottles as the state space, and assume that all of the other products price as a fixed ratio of those four prices, plus some white noise. The consumer is able to form predictions on the evolution of the entire market by tracking only four prices.

Erdem, Imai & Keane (2003) is attractive in that it does not require a simplification like IVS. However, it requires other serious restrictions, and these restrictions are highly specific to one particular market. Common features like product entry and exit, markets with many products or many varying characteristics would substantially complicate this approach, whereas an IVS approach can handle these issues quite naturally.[20]

Similarly, Nair (2007) studies consumer demand for video games. In his approach, consumers treat each video game as a separate market so the product characteristics except price are fixed. He then models consumers as forming expectations over price directly. The state space for each consumer and product is then one dimensional. This substantially lowers the computational burden because the dimensionality of the state space grows linearly instead of exponentially in the number of products.

7 Conclusion

Recent research on dynamic demand addresses contexts with many products and product characteristics, for which numerical programming limitations require simplifying assumptions. A number of recent papers use some version of *Inclusive Value Sufficiency*, which restricts consumers to make predictions based only

[20]Like Erdem, Imai & Keane (2003), Song & Chintagunta (2003) models dynamic demand based on underlying product characteristics rather than inclusive values. In the context of digital cameras, they model the market as having only three products, representing the three main brands.

on inclusive values for purchase over some set of goods. Different papers implement it in different ways, and in particular may or may not define the inclusive value over endogenous continuation values.

Our paper provides a unifying model that nests durable goods, storable goods, and subscription services. We contrast these assumptions in our general model. We show that the inclusive value can be defined over exogenous characteristics only if the choice among those characteristics does not have dynamic content. We also explore the role of separability assumptions on the value of holdings and new purchases in reducing the state space. These issues are relevant for studies of storable goods, durable goods, and subscription services.

References

Aguirregabiria, V. & Nevo, A. (2010). Recent developments in empirical IO: Dynamic demand and dynamic games. *Proceedings of the Econometric Society World Congress, forthcoming.*

Berry, S. (1994). Estimating discrete choice models of product differentiation. *RAND Journal of Economics*, *25*, 242–262.

Berry, S., Levinsohn, J., & Pakes, A. (1995). Automobile prices in market equilibrium. *Econometrica*, *63*, 841–890.

Carranza, J. E. (2010). Product innovation and adoption in market equilibrium: The case of digital cameras. *International Journal of Industrial Organization*, *28*, 604–618.

Conlon, C. T. (2012). A dynamic model of prices and margins in the LCD TV industry. Unpublished Manuscript.

Dube, J.-P. H., Fox, J. T., & Su, C.-L. (2012). Improving the numerical performance of static and dynamic aggregate discrete choice random coefficient demand estimation. *Econometrica*, *80*, 2231–2267.

Erdem, T., Imai, S., & Keane, M. P. (2003). Brand and quantity choice dynamics under uncertainty. *Quantitative Marketing and Economics*, *1*, 5–64.

Gowrisankaran, G., Park, M., & Rysman, M. (2014). Measuring network effects in a dynamic environment. Unpublished Manuscript, Boston University.

Gowrisankaran, G. & Rysman, M. (2012). Dynamics of consumer demand for new durable goods. *Journal of Political Economy*, *120*, 1173–1219.

Hendel, I. & Nevo, A. (2006). Measuring the implications of sales and consumer stockpiling behavior. *Econometrica*, *74*, 1637–1673.

Ho, C.-Y. (2015). Switching cost and the deposit demand in China. *International Economic Review*, *56*, 723–749.

Ho, C.-Y., Rysman, M., & Wang, Y. (2020). Demand for performance goods: Import quotas in the Chinese movie market. Unpublished manuscript, Boston University.

Judd, K. L. & Su, C.-L. (2012). Constrained optimization approaches to estimation of structural models. *Econometrica*, *80*, 2213–2230.

Krusell, P. & Smith, A. A. (1998). Income and wealth heterogeneity in the macroeconomy. *Journal of Political Economics*, *106*, 867–896.

Lee, R. (2013). Vertical integration and exclusivity in platform and two-sided markets. *American Economic Review*, *103*, 2960–3000.

Liu, Y., Li, S., & Shen, C. (2020). The dynamic efficiency in resource allocation: Evidence from vehicle license lotteries in Beijing. SSRN Working Paper #3547563.

McFadden, D. L. (1978). Modeling the choice of residential location. In K. A., F. Lundqvist, L. Snickars, & J. Weibull (Eds.), *Spatial Interaction Theory and Planning Models* (pp. 75–96). Amsterdam: North Holland.

McFadden, D. L. (1981). Econometric models of probabilistic choice. In D. L. McFadden & C. Manski (Eds.), *Structural Analysis of Discrete Choice with Econometric Implications*. Cambridge: MIT Press.

Melnikov, O. (2013). Demand for differentiated durable products: The case of the U.S. computer printer market. *Economic Inquiry*, *51*, 12771298.

Nair, H. (2007). Intertemporal price discrimination with forward-looking consumers: Application to the US market for console video games. *Quantitative Marketing and Economics*, *5*, 239–292.

Nosal, K. (2012). Estimating switching costs for Medicare Advantage plans. Unpublished Manuscript, University of Mannheim.

Rust, J. (1987). Optimal replacement of GMC bus engines: An empirical model of Harold Zurcher. *Econometrica*, *55*, 999–1033.

Schiraldi, P. (2011). Autobmobile replacement: A dynamic structural approach. *RAND Journal of Economics*, *42*, 266–291.

Shcherbakov, O. (2016). Measuring consumer switching costs in the television industry. *RAND Journal of Economics*, *47*, 366–393.

Song, I. & Chintagunta, P. (2003). A micromodel of new product adoption with heterogeneous and forward-looking consumers: An application to the digital camera category. *Quantitative Marketing and Economics*, *1*, 371–407.

Sun, Y. & Ishihara, M. (2019). A computationally efficient fixed point approach to dynamic structural demand estimation. *Journal of Econometrics*, *208*, 563–584.