NBER WORKING PAPER SERIES

NONPARAMETRIC ESTIMATION OF DEMAND WITH SWITCHING COSTS: THE CASE OF HABITUAL BRAND LOYALTY

Xinyao Kong
Jean-Pierre H. Dubé
Øystein Daljord

Working Paper 32994
http://www.nber.org/papers/w32994

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
September 2024

We appreciate support from Kilts Center for Marketing and Charles E. Merrill research fund. All authors contributed equally to this work. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Nonparametric Estimation of Demand with Switching Costs: the Case of Habitual Brand Loyalty
Xinyao Kong, Jean-Pierre H. Dubé, and Øystein Daljord
NBER Working Paper No. 32994
September 2024
JEL No. D11, D12, L66, M3

ABSTRACT

We study habitual brand loyalty, one of the earliest empirically-studied forms of switching costs and a classic source of structural state-dependence in consumer demand. Auxiliary instruments and economically-motivated restrictions can tighten nonparametric bounds on the extent of brand loyalty in choice panel data. We also prove that the canonical dynamic discrete-choice model, nested in our nonparametric framework, has "built-in" exclusion restrictions that semiparametrically identify the discount factor, in general, and point identify it for standard parameterizations of switching costs. Case studies of several large consumer goods categories show that brand loyalty accounts for at least 10.8% but no more than 72.2% of the observed choices across categories studied. In some categories, it accounts for over 90% of observed repeat-purchase behavior. Consumers are found to be forward-looking, but more impatient than would be implied by the real rate of interest.

Xinyao Kong
University of Texas at Dallas
800 W Campbell Rd
Richardson, TX 75080
xinyao.kong@utdallas.edu

Jean-Pierre H. Dubé
University of Chicago
Booth School of Business
5807 South Woodlawn Avenue
Chicago, IL 60637
and NBER
jdube@chicagobooth.edu

Øystein Daljord †
University of Chicago
Booth School of Business
5807 South Woodlawn Avenue
Chicago, IL 60637

1 Introduction

High observed repeat-purchase rates of branded goods stimulated one of the earliest empirical literatures on consumer switching costs: *habitual brand loyalty* (HBL).[1] Defined as the treatment effect of past choices on current choices, HBL embodies a form of structural state-dependence attributed to non-economic or psychological *switching costs*.[2] HBL has a number of important economic implications since it moderates the toughness of price competition (e.g., Klemperer, 1995; Farrell and Klemperer, 2007; Dubé et al., 2010), generates substantial barriers to entry into consumer goods industries (Bain, 1956), and likely accounts for some of the $235 billion in estimated annual U.S. intangible brand capital (Corrado et al., 2005).

Researchers typically do not observe switching costs and, instead, infer them from choice data using parametric random utility models. This approach has been used extensively for HBL (e.g., Erdem, 1996; Keane, 1997; Seetharaman et al., 1999; Dubé et al., 2010) as well as switching costs in other markets, such as healthcare (Handel, 2013; Polyakova, 2016; Heiss et al., 2021; Pakes et al., 2021), rewards programs (Hartmann and Viard, 2008), insurance (Honka, 2014), television (Scherbakov, 2016; Moshkin and Shachar, 2002) and pension plans (Illanes, 2017). A potential concern is the extent to which the empirical evidence for switching costs hinges on the parametric assumptions about a given demand model.

We study HBL as a nonparametric treatment effects problem within a multinomial version of Torgovitsky (2019)'s Dynamic Potential Outcomes (DPO) framework. HBL exists if potential outcomes would have been different under different counterfactual past choices. The DPO formulation clarifies the econometric distinction between the structural state-dependence from brand loyalty and unobserved heterogeneity,[3] is not vulnerable to the "initial conditions" problem[4], and is robust to rich patterns of potentially time-varying

[1] See Brown (1953); Cunningham (1956); Kuehn (1958); Bass et al. (1984).

[2] Egel et al. (1968) describe the underlying mechanism for HBL as cognitive dissonance (Brehm, 1956; Festinger, 1957) causing psychological commitments to recently-chosen brands. This distinguishes HBL from the longer-term loyalty due to brand taste formation (Guest, 1955; Bronnenberg et al., 2012). Klemperer (1995) and Farrell and Klemperer (2007) discuss HBL and psychological switching costs as a key example.

[3] See for instance Heckman, 1981; Honore and Kyriazidou, 2000; Honoré and Tamer, 2006; Torgovitsky, 2019; Pakes et al., 2021.

[4] See for instance (Heckman, 1981; Simonov et al., 2020) for parametric solutions and (Honoré and Tamer, 2006) for a partial identification approach.

unobserved heterogeneity.[5] The DPO generates nonparametric bounds on the extent to which observed choices are due to HBL. A researcher can compute these bounds to test for the presence of HBL nonparametrically before estimating a structural model of demand and imposing potentially untestable parametric assumptions.

Methodologically, we extend a partial identification result in Torgovitsky (2019) showing that the unrestricted identified set for HBL in our multinomial setting always contains zero as its lower bound. So, choice data alone cannot disentangle the state dependence from HBL and unobserved heterogeneity in tastes. To tighten the identified set, we first explore restrictions implied by excluded *instruments*, the traditional source of identification discussed in the literature on state-dependent demand (e.g., Chamberlain, 1985; Dubé et al., 2010; Levine and Seiler, 2023; Berry and Compiani, 2023). We also show that the DPO nests popular structural random utility models of discrete-choice demand, including the canonical dynamic discrete-choice model (DDC). As in Torgovitsky (2019), we can therefore explore alternative economically-motivated restrictions that are implied by standard demand models without needing to impose parametric assumptions.

In the spirit of Becker and Murphy (1988), we then test for "rational" HBL in the sense that consumers' choices account for the future lock-in from HBL and the endogeneity of the loyalty state. We prove that the canonical DDC (Rust, 1987) with HBL has built-in exclusion restrictions that satisfy the conditions in Abbring and Daljord (2020) to identify the discount factor. In particular, the flow utility from choosing one brand excludes the prices of substitute brands in the canonical DDC. We then derive a novel set of moment conditions that set identify the discount factor without any parametric assumptions over consumers' flow utilities. For conventional parametric specifications of the switching cost itself, these moments point identify the discount factor. Therefore, our approach does not impose any more structure than would typically be used to estimate the DDC.

We conduct several empirical case studies of mature consumer packaged goods (CPG) markets with established brands using a household shopping panel dataset. Our nonparametric bounds generate compelling evidence for HBL in the brand choices for several of the CPG categories studied. The lower bounds on HBL for the leading brands in these categories are often found to be well above zero, in some cases with at least 17.9% of the current purchases of the top-selling brand arising from HBL. We also find several cases where at least 16.8% of the choice-persistence in the top-selling brand (i.e., repeat purchases) is due

[5] Unlike a random utility model, the DPO does not require any parametric assumptions about the transitory and persistent components of unobserved heterogeneity.

to HBL, and as much as over 90%. In short, the switching costs from HBL appear to be an important driver of demand for the branded consumer goods studied.

Interestingly, our upper bounds are also quite informative with most categories exhibiting less than 62.7% HBL, suggesting that this form of brand loyalty may not be as pervasive as would be implied in parametric choice models that do not allow for heterogeneity in the loyalty coefficient itself. In particular, the assumption that the loyalty or "switching cost" coefficient is homogeneous (i.e., common across consumers) may overstate the true extent to which consumers' choices are driven by HBL. As further evidence that our HBL bounds reflect psychological switching costs, we also rule out learning and search costs, two alternative mechanisms that could also generate the positive state-dependence in choices (e.g., Dubé et al., 2010).

We also reject the standard, myopic random-utility model of choice in favor of the dynamic model with a freely-varying discount factor in each category. The results suggest that while consumers are forward-looking, they assign less weight to their expected future utility (are more impatient) than would be implied by the real rate of interest. These findings are consistent with Miller et al. (2022) and Chevalier and Goolsbee (2009) who also find consumers to be forward-looking, but more impatient than has been assumed in past work using the DDC. Nevertheless, we find that the estimated long-run equilibrium elasticities are considerably larger for the estimated DDC than the estimates from a myopic model that that restricts the discount factor to equal zero. Most estimated cross-price elasticities are between 500 and 1,000 times larger for the DDC than the myopic model. Hence, estimates from the forward-looking demand model imply a more competitive market structure than estimates from the rejected myopic model.

Our findings have important economic implications as HBL is a leading justification for the approximately $180 billion spent annually in the U.S. on sales promotions[6] to steal brand share from competitors.[7] Indeed, the results imply considerably more competition and substitution once forward-looking behavior is accounted for. These results echo the findings of Erdem et al. (2003) and Hendel and Nevo (2006) who also find large differences between the elasticities from the myopic and forward-looking choice model; albeit with an assumed (not estimated) discount factor value.

[6]Sales promotions typically consist of temporary discounts, coupons, in-store displays, loyalty programs and newspaper feature ads geared towards influencing the consumer at the point-of-sale.

[7]Sales promotions represent about 75% of the $251 Billion spent by U.S. companies on advertising in 2020, according to Statista: accessed on 11-12-2023 at `https://www.statista.com/statistics/429036/advertising-expenditure-in-north-america/`.

Our work also contributes to the broader literature cited above that econometrically estimates switching costs using the myopic discrete-choice model with a discount factor implicitly set deterministically to zero. The nonparametric bounds on switching costs provide prima facie evidence for switching costs before estimating a parametric demand model. Our results also enable future work to test for forward-looking behavior semi-parametrically and to base counterfactual analysis on an empirical estimate of the discount factor.

Our approach is most closely related to the recent study of switching costs by Pakes et al. (2021), who propose a moment inequalities estimator that generates semiparametric bounds on a linear HBL parameter in a myopic choice model. Our empirical findings on a tight upper bound for the extent of HBL may be problematic for the assumption of a homogeneous switching cost parameter across consumers. While our DPO evidence is fully nonparametric, it is not amenable to counterfactual policy simulations. It is unclear whether the approach in Pakes et al. (2021) could be adapted to generate bounds on the discount factor in the context of dynamic discrete-choice. Our work is also closely related to Berry and Compiani (2023)'s analysis of the partial identification of state-dependence in entry using instruments. The DPO formulation allows us to accommodate behavioral restrictions that can potentially tighten the bounds on state-dependence even in the absence of instruments, and to obtain a fully non-parametric identification strategy.

Our work is also related to an older literature testing for HBL nonparametrically using within-consumer variation and generally finding inconclusive results (Anderson and Goodman, 1957; Frank, 1962; Bass et al., 1984)). The DPO enables the pooling of time-series data across households without confounding state-dependence and persistent unobserved heterogeneity. The partial identification strategy (Manski, 1995; Molinari, 2020) is also more agnostic as it is fully nonparametric while still allowing for conditioning on marketing variables like prices and promotions.

The rest of the paper is organized as follows. Section 2 defines the canonical discrete-choice formulation of brand purchase behavior. Section 3 discusses the DPO model and the empirical bounds on HBL, along with a set of identifying assumptions to tighten the bounds. Section 4 derives moments from the DDC to identify and estimate the discount factor. Section 5 describes our empirical case studies and results, with policy implications discussed in section 6. We conclude in section 7.

2 Dynamic Discrete-Choice Demand with State-Dependence

We start by presenting the canonical, economic model of dynamic discrete-choice demand (DDC) (Rust, 1987, 1994) with state dependence on the past brand chosen. The state-dependence nests conventional models of consumer switching costs (Klemperer, 1995; Farrell and Klemperer, 2007; Handel, 2013; Polyakova, 2016; Heiss and Winschel, 2008; Pakes et al., 2021) and habitual brand loyalty (Erdem, 1996; Keane, 1997; Dubé et al., 2010).

Each time period $t = 1, ..., \infty$, a consumer makes a discrete purchase decision $d_t \in \{0, 1, \ldots, J\} \equiv \mathcal{D}$ where the choice set $\mathcal{D}$ containing a no-purchase alternative, $j = 0$, and $j = 1, ..., J$ brands. A consumer in state $\mathbf{s}_t \in \mathcal{S}$ obtains choice-specific utility $u_j(\mathbf{s}_t) + \epsilon_{jt}$, $\forall j \in \mathcal{D}$ where we make the conventional normalization, $u_0(\mathbf{s}_t) = 0$. Following the convention in the literature, we assume $\boldsymbol{\epsilon}_t \equiv (\epsilon_{0t}, ..., \epsilon_{Jt})' \sim$ i.i.d. $\mathbf{F}_\epsilon(\boldsymbol{\epsilon})$ is a $J \times 1$ vector of random utility disturbances. For convenience, we assume $\mathbf{F}_\epsilon(\boldsymbol{\epsilon})$ is the Type-I Extreme value distribution. However, our results hold for any absolutely continuous distribution.

The state vector $\mathbf{s}_t = (\mathbf{p}_t, \ell_t)$ contains the vector of current prices, $\mathbf{p}_t \equiv (p_{1t}, ..., p_{Jt})$, and, as in the extant literature on switching costs and HBL,[8] the consumer's current loyalty state, $\ell_t \in \{1, \ldots, J\} \equiv \mathcal{J}$, indicating the previous brand purchased. As in past literature, the loyalty state is not affected when consumers choose not to purchase, $j = 0$. The consumer's loyalty state evolves deterministically in response to her current purchase decision, $d_t \in \mathcal{D}$, and the current loyalty state, ℓ_t. We assume that $\mathbf{p}_t$ evolves according to an exogenous first-order Markov process with L_p discrete support points, $\mathcal{P}$, and transition matrix $\mathbf{G}$. The state space $\mathcal{S} = \mathcal{P} \times \mathcal{J}$ is, therefore, discrete with L support points: $L = |\mathcal{P}| \times J$. We denote the consumer's beliefs about the evolution of the state as the Markov transition distribution $\mathbf{F}_{d_t}(\mathbf{s}_t) = [\Pr(\mathbf{s}_{t+1} = \mathbf{s}_1 | \mathbf{s}_t, d_t), \ldots, \Pr(\mathbf{s}_{t+1} = \mathbf{s}_L | \mathbf{s}_t, d_t)]$.

The ex-ante value function for the consumer's purchase decision problem is

$$v(\mathbf{s}_t) = \mathbb{E}_\epsilon \left[\max_{j \in \mathcal{D}} \{ v_j(\mathbf{s}_t) + \epsilon_j \} \right] \tag{2.1}$$

where the choice-specific value functions are defined as the solutions to

$$v_j(\mathbf{s}_t) \quad = \quad u_j(\mathbf{s}_t) + \beta \mathbf{F}_j(\mathbf{s}_t) \ln \left\{ \sum_{k \in \mathcal{D}} \exp(\mathbf{v}_k) \right\}, \ \forall j \in \mathcal{D} \tag{2.2}$$

[8]See for instance Farrell and Klemperer, 2007; Erdem, 1996; Seetharaman et al., 1999; Erdem and Sun, 2001; Shum, 2004; Dubé et al., 2010.

and where $\beta \in [0, 1)$ is a discount factor, $\mathbf{v}_j = [v_j(\mathbf{s}_1), \ldots, v_j(\mathbf{s}_L)]'$ is an $L \times 1$ vector, and $\mathbf{F}_j(\mathbf{s}_t)$ is a $1 \times L$ vector. The consumer's utility-maximizing choice probabilities in state s are:

$$\sigma_j(\mathbf{s}) = \frac{\exp(v_j(\mathbf{s}))}{\sum_{k \in \mathcal{D}} \exp(v_k(\mathbf{s}))}, \quad \forall j \in \mathcal{D}. \tag{2.3}$$

We can re-write the choice-specific value functions in (2.2) in the following matrix form:

$$v_j(\mathbf{s}_t) = u_j(\mathbf{s}_t) + \beta \mathbf{F}_j(\mathbf{s}_t)(\mathbf{m} + \mathbf{v}_0),$$

where $\mathbf{v}_0 = [v_0(\mathbf{s}_1), \ldots, v_0(\mathbf{s}_L)]'$ and $\mathbf{m} = [-\ln(\sigma_0(\mathbf{s}_1)), \ldots, -\ln(\sigma_0(\mathbf{s}_L))]'$ are $L \times 1$.

The extant literature typically tests for HBL with variations of the following hypothesis:

$$\mathbb{H}_0 : \mathcal{HBL}_j \equiv v_j(\mathbf{p}, \ell = j) - v_j(\mathbf{p}, \ell = k) > 0 \text{ for some } k \neq j \in \mathcal{J}. \tag{2.4}$$

The alternatives are *variety-seeking* ($\mathcal{HBL} < 0$) and *no state-dependence* ($\mathcal{HBL} = 0$).

HBL creates a dynamic incentive for consumers since choosing brand j today makes it costly to switch to a competing brand in the future. Therefore, in the spirit of Becker and Murphy (1988), $\beta > 0$ implies that the consumer anticipates this future lock-in when making her current choice. When $\beta = 0$, the model collapses to the generic myopic random utility model typically estimated in practice.

3 The Multinomial Dynamic Potential Outcomes Model

Our first goal consists of testing for and quantifying the extent of HBL in a dataset nonparametrically. We use the multinomial analog of the nonparametric, Dynamic Potential Outcomes model (DPO) (Torgovitsky, 2019), defining HBL as the causal effect of a consumer's past choices on current choice. We will show that the DPO nests the DDC with HBL from Section 2 without requiring any parametric assumptions about preferences or the discount function. Furthermore, the DPO model set-identifies the extent of HBL in a brand-choice panel using conventional exclusion restrictions or mild, economically-motivated identifying assumptions, such as stationarity and monotonicity, that are implied by random utility models of demand. As we will show through examples, these restrictions rule out potential outcome paths associated with heterogeneous preferences, thereby tightening the evidence for HBL.

3.1 DPO Model with HBL

Consumer i makes observed brand-choice decisions $Y_{it} \in \mathcal{J} \equiv \{1, ..., J\}$ in time periods $t = 0, 1, ..., T$. Since the extant literature does not test for brand loyalty to the outside good, we exclude the no-purchase alternative from this discussion. After the initial period, the consumer receives the loyalty *treatment* $Y_{i(t-1)}$ in period t so that brand loyalty is defined as the causal effect of last period's brand-choice on the current period's brand choice outcome. As explained above, this definition of loyalty conforms with the extant literature on brand choice. Conceptually, the DPO framework herein can accommodate more sophisticated definitions of the treatment that includes the entire choice history. More important, any evidence herein for our one-lag form of brand loyalty does not rule-out such higher-order forms of loyalty from co-existing in the data (i.e., there is no specification bias unlike in a parametric choice model that assumes the wrong form of loyalty).

In each period t, there exists a set of unobserved potential outcomes, $U_{it}(1), ..., U_{it}(J) \in \mathcal{J}$, corresponding to the brand choices that would have occurred under each of the loyalty states $Y_{i(t-1)} \in \mathcal{J}$. The observed brand-choice outcomes, $\boldsymbol{Y}_i \equiv (Y_{i0}, Y_{i1}, ..., Y_{iT}) \in \mathcal{Y} \equiv \{1, ..., J\}^{T+1}$ are related to potential outcomes $\boldsymbol{U}_i \equiv (Y_{i0}, \boldsymbol{U}_{i1}, ..., \boldsymbol{U}_{iT})$ where $\boldsymbol{U}_{it} \equiv (U_{it}(1), ..., U_{it}(J))$ for $t = 1, 2, ..., T$ through the DPO model

$$Y_{it} = \sum_{j=1}^{J} \mathbb{1}[Y_{i(t-1)} = j] U_{it}(j), \ \forall t \geq 1 \tag{3.1}$$

where the initial condition Y_{i0} is observed but not modeled, avoiding the *initial conditions* problem (Heckman, 1981).

In contrast with past research that defines HBL through the observed outcomes, $\boldsymbol{Y}_i$, the DPO defines HBL through the unobserved potential outcomes $\boldsymbol{U}_i$. The DPO allows for state dependence through the probability that different loyalty states cause different potential outcomes. With no further restrictions, the DPO exhibits HBL when we observe positive state-dependence for a brand j: $U_{it}(j) = j, \ U_{it}(k) \neq j$ for some $k \neq j \in \mathcal{J}$. The DPO exhibits variety-seeking when we observe negative state-dependence for a brand j: $U_{it}(j) \neq j, \ U_{it}(k) = j$, for some $k \neq j \in \mathcal{J}$. The loyalty state has no causal effect in the DPO when we observe the same potential outcomes for a brand j in different loyalty states: $U_{it}(j) = U_{it}(k)$ for all $j \neq k \in \mathcal{J}$, such as when brand choice reflects (persistent) tastes and consumers are *always-takers* of brand j. The central challenge in the DPO framework is that the potential outcomes are only observed in the observed loyalty state. Therefore, at best, we can learn about the distribution of HBL in subsets of the population.

An additional challenge in the DPO, and a classic challenge in the literature on state-dependence in choices, is the separate identification of state-dependence and heterogeneity (Frank, 1962; Heckman, 1981; Erdem, 1996; Keane, 1997; Dubé et al., 2010). The DPO also accounts for unobserved heterogeneity by allowing for variation in the potential outcomes $(U_{it}(1), ..., U_{it}(J))$ in the population. We provide examples below in section 3.3.1.

3.2 The DPO Nests the Canonical DDC

We now show that the DPO nests the myopic and dynamic random utility models from section 2. Below we omit prices (or other marketing variables) from the state. In the canonical DDC, the potential outcomes in any given loyalty state are the choices generated by utility maximization

$$U_{it}(y) = \underset{k \in \mathcal{J}}{\operatorname{argmax}} \{v_k(y) + \epsilon_{kt}\}, \ \forall y \in \mathcal{J} \tag{3.2}$$

where the observed treatment is $y = Y_{i(t-1)} \in \mathcal{J}$. As before, we can relate the observed outcomes $\boldsymbol{Y}_i$ to the potential outcomes as follows:

$$Y_{it} = \sum_{y \in \mathcal{J}} \mathbb{1}[Y_{i(t-1)} = y] U_{it}(y). \tag{3.3}$$

Recall that the choice model in section 2 includes a no-purchase alternative, $j = 0$, to which consumers do not form any loyalty. For the empirical DPO analysis, we condition on purchase (i.e., exclude non-purchase) to reduce the computational burden. However, it is straightforward to reformulate (3.3) to accommodate no-purchase in our DDC analysis as follows:

$$\begin{aligned} U_{it}(y) &= \underset{k \in \mathcal{J}}{\operatorname{argmax}} \{v_k(y) + \epsilon_{kt}\}, \ \forall y \in \mathcal{D} \\ Y_{it} &= \sum_{y \in \mathcal{J}} \mathbb{1}[Y_{i(t-1)} = y] U_{it}(y) + \mathbb{1}[Y_{i(t-1)} = 0] U_{it}(\ell_{t-1}) \\ \ell_t &= \mathbb{1}[Y_{i(t-1)} = 0]\ell_{t-1} + \mathbb{1}[Y_{i(t-1)} \neq 0] Y_{i(t-1)} \end{aligned} \tag{3.4}$$

where y is a counterfactual loyalty state, and ℓ_t is the realized loyalty state.

While the DPO model is agnostic about the parametric form of the potential outcomes function $U_{it}(\cdot)$, the random utility model of demand imposes several parametric restrictions including utility maximization behavior, the additive separability of the random utility disturbances, $\boldsymbol{\epsilon}$, a parametric distributional assumption for the utility disturbances, $\mathbf{F}_{\epsilon}(\boldsymbol{\epsilon})$, and, in the case of the DDC, assumes geometric discounting and a parametric distribution for consumers' beliefs about the state transition process. The DPO can therefore be a useful starting point to document nonparametric evidence for HBL before estimating a parametric DDC to conduct counterfactual analysis.

3.3 Identification of HBL in the DPO model

The empirical goal consists of testing for a causal effect of past brand choice on current decision-making. We therefore define our target parameter vector as follows:

$$\mathcal{HBL}_t(P) \in \{\mathcal{HBL}_{1t}(P), ..., \mathcal{HBL}_{Jt}(P)\}\,, \tag{3.5}$$

where

$$\mathcal{HBL}_{jt}(P) = \mathbb{P}_P[\exists k \neq j \in \mathcal{J} \ \ s.t.\ U_{it}(j) = j, U_{it}(k) \neq j] \tag{3.6}$$

where $\mathbb{P}_P$ is the probability of an event under the probability mass function over the possible potential outcome paths P, and the entire potential outcome path is $\boldsymbol{U}_i \equiv (Y_{i0}, \boldsymbol{U}_{i1}, ..., \boldsymbol{U}_{iT})$, with support $\mathcal{U} \equiv \{1, .., J\}^{JT+1}$. Formally, $\mathcal{HBL}_{jt}$ measures the population probability on date t that there exists at least one brand $k \neq j$ such that a consumer who would choose brand j if they had previously chosen it would choose something else if they had instead previously chosen k. In our brand choice setting, the $\mathcal{HBL}$ parameter likely also depends on demand-shifters at the point of sale, like prices and in-store display advertising. It is straightforward to define versions of the parameter, $\mathcal{HBL}|X$, that depends on such demand-shifters, X, to ensure our results are not spurious.

Measuring HBL under the DPO model (3.1) differs conceptually from the more familiar structural random utility framework, where HBL consists of a preference parameter. Our model primitive now consists of the probability mass function P, where $P(\boldsymbol{\mu})$ represents the probability of observing potential outcome path $\boldsymbol{\mu}$ in the consumer population. Let $\mathcal{P}$

denote the set of all probability mass functions P where

$$P : \mathcal{U} \to [0,1] \quad \text{s.t.} \sum_{\boldsymbol{\mu} \in \mathcal{U}} P(\boldsymbol{\mu}) = 1. \tag{3.7}$$

The parameter space, $\mathcal{P}^\dagger$, consists of the subset of $\mathcal{P}$ that satisfies the researcher's prior assumptions:

$$\mathcal{P}^\dagger = \{P \in \mathcal{P} : \rho(P) \geq 0\}, \tag{3.8}$$

where $\rho : \mathcal{P} \to \mathbb{R}^{d_\rho}$ is a function representing assumptions on P, d_ρ is the dimensionality of the restrictions, and the inequality is interpreted component-wise. $\rho(P)$ denotes additional prior restrictions, discussed below in more detail, that we will impose to tighten identification. In the empirical application of the DPO, we will also need to address the fact that the dimension of $\mathcal{U}$ increases exponentially in J and T.

We can now define the *identified set*, $\mathcal{P}^\star$, as the observationally equivalent subset of the parameter space, $\mathcal{P}^\dagger$, for which the DPO can rationalize the observed brand choices. Let $\mathbb{P}[\boldsymbol{Y}_i = \boldsymbol{y}]$ denote the observed probability of $\boldsymbol{Y}_i$. Then the observational equivalence requires that, for every $\boldsymbol{y} \equiv (y_0, y_1, .., y_T) \in \mathcal{Y}$,

$$\mathbb{P}[\boldsymbol{Y}_i = \boldsymbol{y}] = \mathbb{P}_P[Y_{i0} = y_0, U_{it}(y_{t-1}) = y_t \text{ all } t \geq 1] = \sum_{\boldsymbol{\mu} \in \mathcal{U}_{oeq}(\boldsymbol{y})} P(\boldsymbol{\mu}), \tag{3.9}$$

where $\mathcal{U}_{oeq}$ is the set of all potential outcome sequences that can generate the observed choices $\boldsymbol{y} = (y_0, y_1, .., y_T)$ through $\mu_0 = y_0$ and $\mu_t(y_{t-1}) = y_t$ for all $t \geq 1$.

We now turn to the testing of HBL and the corresponding identified set. Recall that we only observe the potential outcome in the observed state, $U_{it}(Y_{i(t-1)})$, so that $\{U_{it}(y)\}_{y \neq Y_{i(t-1)} \in \mathcal{J}}$ is unobserved. The following proposition clarifies that the $\mathcal{HBL}_{jt}$ parameters (3.6) are not point-identified, but are set-identified. The proposition extends Torgovitsky (2019)'s Proposition 1 to the case of multinomial choice.

Proposition 1. *Suppose that $\mathcal{P}^\dagger = \mathcal{P}$. If $\theta = \mathcal{HBL}_{jt}(P)$, then*

$$\begin{aligned} \Theta^\star &= \big[0, \mathbb{P}[Y_{i(t-1)} = j, Y_{it} = j] + \mathbb{P}[Y_{i(t-1)} \neq j]\big] \quad \text{when } |\mathcal{J}| \geq 3, \\ \Theta^\star &= \big[0, \mathbb{P}[Y_{i(t-1)} = j, Y_{it} = j] + \mathbb{P}[Y_{i(t-1)} \neq j, Y_{it} \neq j]\big] \quad \text{when } |\mathcal{J}| = 2. \end{aligned} \tag{3.10}$$

Proof. See Appendix A.1. □

Proposition 1 confirms that the data alone (i.e., $\mathcal{P}^\dagger = \mathcal{P}$) cannot reject pure heterogeneity in tastes as the driver of observed choices. In other words, the data alone can never reject the hypothesis that consumer behavior consists of "always-taking" whereby consumers' choices are independent of the state. Nevertheless, the data alone can tighten the upper bound on $\mathcal{HBL}$ depending on the amount of serial dependence in choices – switching behavior. We illustrate these results in the simple pedagogical example below. In the next section, we discuss the identifying information contained in behavioral restrictions and excluded instruments.

3.3.1 A Pedagogical Example

Assume binary choice, $J = 2$, and a population with only two consumers, $i \in 1, 2$, who make an initial choice in period $t = 0$ and are observed to make two subsequent choices in periods $t = 1$ and $t = 2$. Suppose we observe the choice sequences $\boldsymbol{Y}_1 = (0, 0, 0)$ and $\boldsymbol{Y}_2 = (0, 1, 1)$ for consumers 1 and 2, respectively. Recall that we do not observe each consumer's true potential outcome path, $\boldsymbol{U}_i$. Instead, we can only make inferences about the distribution of $\boldsymbol{U}_i$ by examining the set of candidate potential outcome paths we can exclude based on the observed choices, $\boldsymbol{Y}_i$.

Let $\boldsymbol{\mu} = (\mu_0, (\mu_1(0), \mu_1(1)), (\mu_2(0), \mu_2(1))) \in \mathcal{U}$ denote a candidate potential outcomes path that starts with an initial choice μ_0 and exhibits pairs of potential outcomes $(\mu_t(0), \mu_t(1))$ in each period $t = 1, 2$. We use the shorthand $(\cdot, (\cdot, \cdot), (\cdot, \cdot))$ for the *set* of potential outcome paths where each $\cdot$ can take any value in $\{0, 1\}$. In other words, $(\cdot, (\cdot, \cdot), (\cdot, \cdot))$ is the set $\mathcal{U}$ and, for instance, $(0, (\cdot, \cdot), (1, \cdot))$ is the set of paths where $\mu_0 = 0$ and $\mu_2(0) = 1$. In our example, there are $2^5 = 32$ candidate potential outcomes paths in $\mathcal{U}$, since each $\cdot$ can take either 0 or 1.

Without additional restrictions, the identified set corresponding to our data is

$$\begin{aligned} \mathcal{P}^\star_{oeq} = \{P : \quad & \sum_{\boldsymbol{\mu} \in \mathcal{U}} P(\boldsymbol{\mu}) = 1, \quad P(\boldsymbol{\mu}) \in [0, 1], \\ & P[\boldsymbol{\mu} \in (0, (0, \cdot), (0, \cdot))] = \mathbb{P}[\boldsymbol{Y}_i = (0, 0, 0)] = \frac{1}{2}, \\ & P[\boldsymbol{\mu} \in (0, (1, \cdot), (\cdot, 1))] = \mathbb{P}[\boldsymbol{Y}_i = (0, 1, 1)] = \frac{1}{2}\} \end{aligned} \tag{3.11}$$

The first row is required for P to be a probability mass function on $\mathcal{U}$. In the second and third rows, the first equality reflects observational equivalence: the set of potential outcomes paths on the left-hand side that are consistent with the observed choices on the right-hand side. In the second line, for instance, the consumer chooses $j = 0$ in the initial period and each of the subsequent two periods, so that the observed state is also 0 in both subsequent periods. We do not observe the potential outcomes in state 1 since the consumer is never in state 1 in the data. We use $\cdot$ to denote the potential outcomes in the unobserved loyalty states each period. The second equality in each of rows two and three reflects the empirical frequency of the choice path, which is 50% in each case. Even though only 8 of the 32 candidate potential outcome paths can be consistent with the data,[9] we cannot rule out pure heterogeneity. For instance, $\boldsymbol{\mu} \in (0, (0, 0), (0, \cdot))$ is in the identified set above and is consistent with consumer $i = 1$ being an *always-taker* of $j = 0$ in period $t = 1$. That is, we cannot rule out that consumer $i = 1$ would have chosen $j = 0$ regardless of the state.

We can now define the target parameters $\mathcal{HBL}_t(P)$ for any P in the identified set. For period 1, we have

$$\mathcal{HBL}_1(P) = P[\boldsymbol{\mu} \in (\cdot, (0, 1), (\cdot, \cdot))] = P[\boldsymbol{\mu} = (0, (0, 1), (0, 0))] + P[\boldsymbol{\mu} = (0, (0, 1), (0, 1))]$$

where the first equality is the definition of $\mathcal{HBL}_1(P)$ and the second equality reflects the fact that only two of the eight potential outcome paths consistent with $\boldsymbol{Y}_1$ and $\boldsymbol{Y}_2$ exhibit HBL in period 1. We can bound $\mathcal{HBL}_1(P)$ within the closed interval $[0, 1/2]$:

$$0 \leq P[\boldsymbol{\mu} = (0, (0, 1), (0, 0))] + P[\boldsymbol{\mu} = (0, (0, 1), (0, 1))] \leq \mathbb{P}[\boldsymbol{Y}_i = (0, 0, 0)] = 1/2.$$

where the first inequality reflects the constraint that probabilities need to be non-negative and the second inequality reflects the observational equivalence constraint in the second line of (3.11). But, in period 2 we cannot reject that both consumers exhibited HBL and, hence, $\mathcal{HBL}_2 \in [0, 1]$. This is because both consumers repeat-purchase the same brand in period 2, and we cannot rule out that they would have purchased the other brand had they been loyal to it instead. Averaging over time, our estimate of $\mathcal{HBL}$ for the total sample is $\hat{\mathcal{HBL}} = \frac{1}{2}(\hat{\mathcal{HBL}}_1 + \hat{\mathcal{HBL}}_2) \in [0, \frac{3}{4}]$.

As established in Proposition 1, the pure empirical evidence in the unrestricted data always fail to reject "no HBL." However, the serial dependence in the data – consumer 1's repeat-choice of $j = 0$ in $t = 1$ and consumer 2's repeat-choice of $j = 1$ in $t = 2$ – tighten the upper bound.

[9] There are 4 potential outcomes paths in the set $(0, (0, \cdot), (0, \cdot))$ and 4 in the set $(0, (1, \cdot), (\cdot, 1))$.

3.4 Identifying Assumptions

We now explore two sources of nonparametric restrictions, the function ρ in (3.8), to tighten the $\mathcal{HBL}$ bounds: (1) instruments and (2) behavioral restrictions implied by economic models of choice. These restrictions can also be adapted to the case where the DPO model conditions on marketing variables, such as prices or point-of-sale promotions. See Appendix B.2.3 for technical details.

3.4.1 Instruments

As in Berry and Compiani (2023), we can use a generalized instrumental variable (GIV) to restrict the potential outcome paths that can rationalize the data and tighten the $\mathcal{HBL}$ bounds.[10] Following Chamberlain (1985), we use lagged demand-shifting variables as instruments for the endogenous loyalty state. Under the assumption these variables do not have any carry-over effects, they do not affect current potential outcomes conditional on the current state - HBL and the current values of demand shifters.[11] We use the following nonparametric exclusion restriction:

Assumption IV: Consider $\boldsymbol{U}_{it} = (U_{it}(1), ..., U_{it}(J))$, the actual potential outcomes for period t as before, some other candidate potential outcomes path $\boldsymbol{\mu} \equiv (\mu(1), ..., \mu(J)) \in \{1, ..., J\}^J$, and an observed covariate $\boldsymbol{X}_{it}$. The IV restriction is satisfied if for every $P \in \mathcal{P}^\dagger$, all $t \geq 1$, every $\boldsymbol{\mu}$, and every $\boldsymbol{x}^1, \boldsymbol{x}^0 \in \mathcal{X}$ such that $\mathbb{P}[\boldsymbol{X}_{it} = \boldsymbol{x}^1] \neq 0$ and $\mathbb{P}[\boldsymbol{X}_{it} = \boldsymbol{x}^1, \boldsymbol{X}_{i(t-1)} = \boldsymbol{x}^0] \neq 0$,

$$\mathbb{P}_P[\boldsymbol{U}_{it} = \boldsymbol{\mu} | \boldsymbol{X}_{it} = \boldsymbol{x}^1] = \mathbb{P}_P[\boldsymbol{U}_{it} = \boldsymbol{\mu} | \boldsymbol{X}_{it} = \boldsymbol{x}^1, \boldsymbol{X}_{i(t-1)} = \boldsymbol{x}^0].$$

In other words, conditional on the current demand-shifters, lagged demand-shifters do not change the potential outcomes. Such an instrument could be constructed with prices or promotions, for instance.

Microeconomic theory provides yet another natural candidate instrument restriction

[10] Berry and Compiani (2023)'s analysis of monopoly entry into a market with persistent heterogeneity and state-dependence is analogous to the binary version of our brand-choice model; although their analysis also extends to multinomial action spaces and oligopoly market structures.

[11] Erdem and Sun (2001) and Dubé et al. (2009a) implement parametric versions of such tests in reduced-form by conditioning on lagged prices in choice models that exclude the loyalty state. Levine and Seiler (2023) explore an interesting polar case where a product is exogenously dropped temporarily from the choice set, causing its lagged price to increase temporarily to infinite.

based on the downward-sloping nature of demand in the current price level, at least for the myopic version of the model. More formally, prices satisfy the Monotone Instrumental Variable restriction introduced by Manski and Pepper (2000, 2009):

Assumption MIV: Consider $U_{it}(y)$, the potential outcome in loyalty state y as before, and an observed covariate $\boldsymbol{X}_{it} = (X_{ijt}, \boldsymbol{X}_{i,-j,t})$, where X_{ijt} takes values in a partially ordered set. The MIV restriction is satisfied if for every $P \in \mathcal{P}^{\dagger}$, $\mathbb{P}_P[U_{it}(y) = j | X_{ijt} = x^1, \boldsymbol{X}_{i,-j,t} = \boldsymbol{x}^0]$ is weakly decreasing in x^1 for every $\boldsymbol{x}^0$, each j, every y, and every $t \geq 1$.

In the case of prices, the MIV assumption implies that the potential demand for brand j in any given loyalty state is decreasing in its current own price, holding competitor prices constant. Unlike our IV restriction above, the MIV restriction on price could be rejected if consumers are forward-looking and form beliefs that generate non-monotonicities in prices. For instance, a price reduction on an older vintage of a product might cause the consumer to delay purchase in anticipation of a subsequent clearance sale, violating the MIV restriction.

Example: Suppose we adapt the simple pedagogical example from section 3.3.1 to include the relative price of brand 1 versus brand 0 as either high or low, $p \in H, L$. The data contain each consumer's choice history $\boldsymbol{Y} = \{\boldsymbol{Y}_1 = (0,0,0)$ and $\boldsymbol{Y}_2 = (0,1,1)\}$ and price history $\boldsymbol{p} = \{\boldsymbol{p}_1 = (H, L)$ and $\boldsymbol{p}_2 = (L, L)\}$. So, consumer 1 faces a relatively high price in period 1 and a relatively low price in period 2.

With no restrictions, we again find $\mathcal{HBL}_1 \in [0, \frac{1}{2}]$ and $\mathcal{HBL}_2 \in [0,1]$ so that $\hat{\mathcal{HBL}} \in [0, \frac{3}{4}]$. One implication of imposing IV is:

$$P[\boldsymbol{\mu} \in (\cdot, (\cdot, \cdot), (0,0) | \boldsymbol{p}_1] = P[\boldsymbol{\mu} \in (\cdot, (\cdot, \cdot), (0,0) | \boldsymbol{p}_2].$$

Conditional on period-2 prices, the probability of potential outcomes (0,0) in period 2 is independent of lagged prices. The right-hand side is 0 since it is inconsistent with either consumer's observed choices and prices. Under IV, it follows that the left-hand side is also 0, so that consumer $i = 1$ cannot have been an always-taker on product $j = 0$ in period $t = 2$. Another implication of IV is:

$$P[\boldsymbol{\mu} \in (\cdot, (\cdot, \cdot), (1,1) | \boldsymbol{p}_1] = P[\boldsymbol{\mu} \in (\cdot, (\cdot, \cdot), (1,1) | \boldsymbol{p}_2].$$

Now the left-hand side is 0 since it is inconsistent with either consumer's data. Under IV, it follows that the right-hand side is also 0, so that consumer $i = 2$ cannot have been an always-taker in period $t = 2$. Therefore, under IV, we can now point-identify $\mathcal{HBL}$ in period $t = 2$: $\mathcal{HBL}_2 = 1$.

3.4.2 Stationarity

As in Torgovitsky (2019), we define stationarity, a standard assumption for parametric models of HBL, over multiple periods as follows:

Assumption ST: Let m be a non-negative integer that the researcher chooses for assuming multi-period stationarity. Define $\boldsymbol{U}_{it} = (U_{it}(1), ..., U_{it}(J))$ for all $t \geq 1$. Define $\boldsymbol{\mu} \equiv (\boldsymbol{\mu}_1, ..., \boldsymbol{\mu}_m) \in \{1, ..., J\}^{J(m+1)}$, where $\boldsymbol{\mu}_m = (\mu_m(1), ..., \mu_m(J))$. Then for any $P \in \mathcal{P}^\dagger$, every $\boldsymbol{\mu} \in \{1, ..., J\}^{J(m+1)}$, and every $t, t' \geq m + 1$,

$$\mathbb{P}_P[\boldsymbol{U}_{i(t-m)} = \boldsymbol{\mu}_1, ..., \boldsymbol{U}_{it} = \boldsymbol{\mu}_m] = \mathbb{P}_P[\boldsymbol{U}_{i(t'-m)} = \boldsymbol{\mu}_1, ..., \boldsymbol{U}_{it'} = \boldsymbol{\mu}_m].$$

A simple form of stationarity with $m = 0$ ensures the time-invariance of the joint distribution of single-period potential outcomes $(U_{it}(1), ..., U_{it}(J))$ in the population.

The following proposition establishes that assumption ST is implied by a standard DDC with stationary flow utilities (preferences), covariates and utility shocks.

Proposition 2. *Suppose that every $P \in \mathcal{P}^\dagger$ is consistent with $U_{it}(y)$ being generated through (3.2) by the DDC model in section 2. Denote the observed state as $\boldsymbol{s} \equiv (\boldsymbol{x}, y)$ where x represents all covariates and y the loyalty state. Let $\mathbf{F}_{\boldsymbol{X}}(\boldsymbol{x}_{t_1+\tau}, ..., \boldsymbol{x}_{t_n+\tau})$ be the cumulative distribution function of $\{\boldsymbol{X}_t\}$ at times $t_1 + \tau$,...,$t_n + \tau$. For a chosen non-negative integer r, if $\mathbf{F}_{\boldsymbol{X}}(\boldsymbol{x}_{t-r}, ..., \boldsymbol{x}_t) = \mathbf{F}_{\boldsymbol{X}}(\boldsymbol{x}_{t'-r}, ..., \boldsymbol{x}_{t'})$ for all $t, t' \geq r + 1$, Assumption ST is satisfied with $m = r$.*

Proof. See Appendix A.1. □

This definition of ST implicitly restricts the underlying distribution of any demand-shifting covariates. If we also assume a stationary distribution over these covariates (e.g., ergodic Markov),[12] then ST mechanically holds once the chain of choices and covariates has reached their stationary distribution. However, it is straightforward to adapt the ST restriction and compute bounds on $\mathcal{HBL}$ conditional on covariates, like prices and in-store promotions, so that ST imposes no implicit restrictions on the covariates (see Appendix

[12] CPG prices can be Markovian (Eichenbaum et al., 2011; Dekimpe et al., 1999).

B.2.3 for technical details). This conditional HBL also serves as a robustness check that our empirical bounds are not picking up spurious brand loyalty from unmodeled non-stationarity in covariates.

After conditioning on covariates, ST is implied by most applications of discrete-choice demand due to the commonly-assumed two-component structure of unobserved heterogeneity: (1) a transitory component with i.i.d. “random utility" shock and (2) a time-invariant, perfectly-correlated serial dependent component with random coefficients.

Under ST, it is not possible for mechanisms that generate a non-stationary treatment effect of past choices on current choices to be the source of HBL. Therefore, under ST, any evidence for HBL cannot be attributed to alternative, non-stationary sources of persistence like Bayesian learning (e.g., about the brand and/or product quality as in Erdem and Keane (1996)).

Example: Consider a slightly modified version of our two-consumer, three-period example in Section 3.3.1, where we now observe outcomes $\boldsymbol{Y}_1 = (0,0,1)$ and $\boldsymbol{Y}_2 = (0,1,1)$. The observational equivalence condition requires that $\mathbb{P}[\boldsymbol{Y}_i = (0,0,1)] = P[\boldsymbol{\mu} \in (0,(0,\cdot),(1,\cdot))] = 1/2$ and $\mathbb{P}[\boldsymbol{Y}_i = (0,1,1)] = P[\boldsymbol{\mu} \in (0,(1,\cdot),(\cdot,1))] = 1/2$.

If we let $m = 0$, Assumption ST imposes the following additional restrictions that for every $\tilde{u}_0, \tilde{u}_1 \in \{0,1\}$,

$$P[\boldsymbol{\mu} \in (\cdot,(\tilde{u}_0,\tilde{u}_1),(\cdot,\cdot))] = P[\boldsymbol{\mu} \in (\cdot,(\cdot,\cdot),(\tilde{u}_0,\tilde{u}_1))] \tag{3.12}$$

In practice, ST is testable and in our example it would not hold if we only observed the consumer with choice sequence $\boldsymbol{Y}_1$. However, ST does hold in our simple two-consumer panel because it adds identifying information across consumers. Starting with the right-hand side of (3.12) when $\tilde{u}_0 = 0$ and $\tilde{u}_1 = 0$, we have $P[\boldsymbol{\mu} \in (\cdot,(\cdot,\cdot),(0,0))] = 0$ because this potential outcomes pattern is inconsistent with either $\boldsymbol{Y}_1$ or $\boldsymbol{Y}_2$. It follows that on the left-hand-side, which is only consistent with $\boldsymbol{Y}_1$, we have $P[\boldsymbol{\mu} \in (\cdot,(0,0),(\cdot,\cdot))] = 0$. Therefore, under Assumption ST, we can only rationalize $\boldsymbol{Y}_1$ for $\boldsymbol{\mu} \in (0,(0,1),(1,\cdot))$, giving us: $\mathbb{P}[\boldsymbol{Y}_i = (0,0,1)] = P[\boldsymbol{\mu} \in (0,(0,1),(1,\cdot))] = 1/2$. In sum, Assumption ST rules out that consumers 1 is an always-taker in period $t = 1$.

Turning to the case when $\tilde{u}_0 = 0$ and $\tilde{u}_1 = 1$ in (3.12), the left-hand-side potential outcomes pattern $(\cdot,(0,1),(\cdot,\cdot))$ is only consistent with $\boldsymbol{Y}_1$ and so $P[\boldsymbol{\mu} \in (\cdot,(0,1),(\cdot,\cdot))] = P[\boldsymbol{\mu} \in (0,(0,1),(1,\cdot))] = \mathbb{P}[\boldsymbol{Y}_i = (0,0,1)] = 1/2$, where the first equality comes from observational equivalence and the second equality comes from the previous step, where Assumption ST ruled out heterogeneity for $\boldsymbol{Y}_1$ in period $t = 1$. By Assumption ST, the right-hand-side $P[\boldsymbol{\mu} \in (\cdot, (\cdot, \cdot), (0, 1))] = 1/2$. Since the right-hand-side is only consistent with $\boldsymbol{Y}_2$, it implies that $P[\boldsymbol{\mu} \in (0, (1, \cdot), (0, 1))] = 1/2$ and so $\boldsymbol{Y}_2$ is only rationalized with $\boldsymbol{\mu} \in (0, (1, \cdot), (0, 1))$. Therefore, Assumption ST also rules out heterogeneity for $\boldsymbol{Y}_2$ in period $t = 2$.

Restrictions (3.12) when $(\tilde{u}_0, \tilde{u}_1) = (1, 0)$ or $(1, 1)$, together with restrictions reflecting the two cases above, and observational equivalence, imply that $P[\boldsymbol{\mu} \in (0, (0, 1), (1, \tilde{u}))] = P[\boldsymbol{\mu} \in (0, (1, \tilde{u}), (0, 1))]$ for $\tilde{u} = 0, 1$. The combination of the data and Assumption ST therefore generate the identified set

$$
\begin{aligned}
\mathcal{P}^{\star}_{ST} = \{P: \quad & \sum_{\boldsymbol{\mu} \in \mathcal{U}} P(\boldsymbol{\mu}) = 1, \quad P(\boldsymbol{\mu}) \in [0, 1], \\
& P[\boldsymbol{\mu} \in (0, (0, 1), (1, \cdot))] = \mathbb{P}[\boldsymbol{Y}_i = (0, 0, 1)] = \frac{1}{2}, \\
& P[\boldsymbol{\mu} \in (0, (1, \cdot), (0, 1))] = \mathbb{P}[\boldsymbol{Y}_i = (0, 1, 1)] = \frac{1}{2}, \\
& P[\boldsymbol{\mu} = (0, (1, \tilde{u}), (0, 1))] = P[\boldsymbol{\mu} = (0, (0, 1), (1, \tilde{u}))] \quad \text{for } \tilde{u} = 0, 1\}.
\end{aligned}
$$

Furthermore, we have $\mathcal{HBL}_1(P) = P[\boldsymbol{\mu} \in (\cdot, (0, 1), (\cdot, \cdot))] = P[\boldsymbol{\mu} \in (\cdot, (0, 1), (1, \cdot))] = 1/2$ and $\mathcal{HBL}_2(P) = P[\boldsymbol{\mu} \in (\cdot, (\cdot, \cdot), (0, 1))] = P[\boldsymbol{\mu} \in (\cdot, (1, \cdot)), (0, 1)] = 1/2$ for all $P \in \mathcal{P}^{\star}_{ST}$. Therefore, $\mathcal{H}\hat{\mathcal{B}}\mathcal{L}_1$ and $\mathcal{H}\hat{\mathcal{B}}\mathcal{L}_2$ are point identified under stationarity in this example and 50% of the population exhibits HBL in each period. Overall, we have $\mathcal{H}\hat{\mathcal{B}}\mathcal{L} = \frac{1}{2}$.

Intuitively, Assumption ST tightens the identified set by restricting the flexibility of population heterogeneity in the data, thereby bounding the degree of HBL from below. In the example above, Assumption ST excluded heterogeneity of the form $(\cdot, (0, 0), (\cdot, \cdot))$ in $t = 1$ using observational equivalence in $t = 2$, effectively bounding the $\mathcal{HBL}$ parameters above 0.

3.4.3 Monotone Treatment Selection (MTS)

As in Torgovitsky (2019), we also consider Manski and Pepper (2000)'s Monotone Treatment Selection restriction:

Assumption MTS: For $\forall y, \tilde{y} \in \{1, ..., J\}$ and all $t \geq 2$ s.t. $\mathbb{P}[Y_{i(t-1)} = j | Y_{i(t-2)} = \tilde{y}] \in (0, 1)$, every $P \in \mathcal{P}^\dagger$ satisfies

$$\mathbb{P}_P[U_{it}(y) = j | Y_{i(t-1)} = j, Y_{i(t-2)} = \tilde{y}] \geq \mathbb{P}_P[U_{it}(y) = j | Y_{i(t-1)} \neq j, Y_{i(t-2)} = \tilde{y}]. \tag{3.13}$$

Assumption MTS implies that a consumer who purchased brand j in period $t-1$ is more likely to choose brand j in all potential loyalty states during period t than some other consumer who does not choose j in period $t-1$. Conditioning on $Y_{i(t-2)}$ ensures that $Y_{i(t-1)}$ are comparable. Intuitively, the MTS assumption implies that consumers who self-selected into loyalty state j have a weakly higher propensity to purchase j in the following period than those who didn't select into loyalty state j. In the myopic choice model, this propensity to purchase might, for example, come from a persistent unobserved preference for the brand or from positively serially-correlated utility shocks.

The following proposition establishes that weakly positive serial correlation in the choice probabilities is sufficient for the potential outcomes implied by the DDC model in section 2 to satisfy MTS.

Proposition 3. *Suppose that every $P \in \mathcal{P}^\dagger$ is consistent with $U_{it}(y)$ being generated through (3.2) by the DDC model in section 2. Assumption MTS is satisfied if the following condition is satisfied for each $(y, \tilde{y}) \in \mathcal{J}^2$:*

$$\mathbb{E}_{(\mathbf{p}_t, \mathbf{p}_{t-1})}\left[\sigma_j(\mathbf{p}_t, y)\sigma_j(\mathbf{p}_{t-1}, \tilde{y}) | Y_{i(t-2)} = \tilde{y}\right] \geq \mathbb{E}_{\mathbf{p}_t}\left[\sigma_j(\mathbf{p}_t, y) | Y_{i(t-2)} = \tilde{y}\right] \cdot \mathbb{E}_{\mathbf{p}_{t-1}}\left[\sigma_j(\mathbf{p}_{t-1}, \tilde{y}) | Y_{i(t-2)} = \tilde{y}\right]$$

where $\sigma_j(\mathbf{p}, \tilde{y})$ is defined in (2.3).

Proof. See Appendix A.1. □

In practice, weakly positive serial correlation in the choice probabilities may be a tenuous assumption. In a model with no covariates, MTS only requires weakly positive serial dependence in the random utility shocks. However, in the presence of time-varying demand shifters, MTS may impose restrictions on their distribution. To illustrate, consider a binary-choice version of the myopic choice model ($\beta = 0$) with the following choice-specific values: $u_j(\mathbf{p}, l) = \gamma_j - \alpha p_j + \lambda \mathbb{I}_{\{l=j\}}$, where $(\gamma_1, ..., \gamma_J, \alpha, \lambda)$ are parameters and $\mathbb{I}_{\{l=j\}}$ indicates whether the consumer repeat-purchases the last brand chosen. If we allow for weakly positive serial-dependence in the random utility shocks, ϵ_j, it is straightforward to show then weakly positive serial-dependence in the price difference between the two products, $p_{1t} - p_{0t}$, is sufficient for the potential outcomes to satisfy MTS. While MTS may be violated by some covariate distributions, we find that it holds 100% of the time in each of our empirical applications in section 5.4 below using a parametric DDC and Markov-distributed prices. In empirical applications where MTS does not hold, one could condition the $\mathcal{HBL}$ parameter on price levels and modify the identifying restrictions accordingly.

Example: We again revisit the two-consumer example in Section 3.3.1, where we now observe outcomes $\boldsymbol{Y}_1 = (0,0,1)$ and $\boldsymbol{Y}_2 = (0,1,1)$. The observational equivalence condition alone gives us: $\hat{\mathcal{HBL}}_2 = P[\boldsymbol{\mu} \in (\cdot,(\cdot,\cdot),(0,1))] = P[\boldsymbol{\mu} \in (0,(1,\cdot),(0,1))] \leq \mathbb{P}[\boldsymbol{Y}_i = (0,1,1)] = 1/2$. Assumption MTS imposes the following additional restriction:

$$P[U_{i2}(0) = 1 | Y_{i1} = 1, Y_{i0} = 0] \geq P[U_{i2}(0) = 1 | Y_{i1} = 0, Y_{i0} = 0] \tag{3.14}$$

$$P[U_{i2}(1) = 1 | Y_{i1} = 1, Y_{i0} = 0] \geq P[U_{i2}(1) = 1 | Y_{i1} = 0, Y_{i0} = 0]. \tag{3.15}$$

The right-hand side of (3.14) is consistent with $\boldsymbol{Y}_1$ and occurs with probability 1: $P[U_{i2}(0) = 1 | Y_{i1} = 0, Y_{i0} = 0] = \mathbb{P}[Y_{i2} = 1 | Y_{i1} = 0, Y_{i0} = 0] = 1$. The inequality implies that the left-hand side must give $P[U_{i2}(0) = 1 | Y_{i1} = 1, Y_{i0} = 0] = 1$. Since the conditioning event on the left-hand-side can only be consistent with $\boldsymbol{Y}_2$, (3.14) rules out HBL for consumer 2 in $t = 2$: $\hat{\mathcal{HBL}}_2 = P[\boldsymbol{\mu} \in (0,(1,\cdot),(0,1))] = 0$. It does not impose additional restriction on $\hat{\mathcal{HBL}}_1$. In restriction (3.15), the left-hand side is consistent with $\boldsymbol{Y}_2$ and occurs with probability 1. But it has no information for either $\hat{\mathcal{HBL}}_1$ or $\hat{\mathcal{HBL}}_2$ because the inequality always holds. Therefore, Assumption MTS reduces $\hat{\mathcal{HBL}}_2$ from 0.5 to 0 and has no impact on $\hat{\mathcal{HBL}}_1$. Overall we have $\hat{\mathcal{HBL}} \in [0, \frac{1}{4}]$.[13]

Intuitively, Assumption MTS tightens the identified set by restricting the sources of serial-dependence in the data which can be confounded with HBL. The restriction bounds the degree of HBL from above.

3.4.4 Monotone Treatment Response (MTR)

A natural candidate restriction is directly on sign of the state dependence in brand choices. We could explicitly rule out negative state dependence (i.e., variety-seeking) through the

[13]In our two-consumer example, this suggests that MTS and ST are incompatible assumptions since ST led to point identification at $\hat{\mathcal{HBL}} = \frac{1}{2}$. If we assume both ST and MTS, the identified set $\mathcal{P}^{\star}_{ST+MTS}$ is empty. In our empirical applications, we will test for the non-emptiness of the sample identified set as a test of specific combinations of our identifying restrictions.

following restriction, as in Manski (1997):

Assumption MTR: For $\forall k \neq j$ and for all t, every $P \in \mathcal{P}^\dagger$ satisfies

$$\mathbb{P}_P[\mathbb{1}\{U_{it}(j) = j\} \geq \mathbb{1}\{U_{it}(k) = j\}] = 1. \tag{3.16}$$

This assumption captures the idea that being loyal to a brand weakly improves the utility of that brand more than the utility of any other brand.

The following proposition establishes a sufficient condition for the DDC model in section 2 to imply MTR. In the canonical linear utility version of the DDC, Assumption MTR holds if the *loyalty* coefficient is non-negative. More generally, this assumption is implied by a separable additive utility function where a change in state from k to j weakly increases the utility of choice j relative to all other options (Heckman and Vytlacil, 2007; Heckman et al., 2008; Lee and Salanié, 2018; Irace, 2018).

Proposition 4. *Suppose that every $P \in \mathcal{P}^\dagger$ is consistent with $U_{it}(y)$ being generated through (3.2) by the DDC model in section 2. If $u_j(\mathbf{p}, \ell = j) - u_j(\mathbf{p}, \ell = k) \geq 0$ and $u_l(\mathbf{p}, \ell = j) - u_l(\mathbf{p}, \ell = k) \leq 0$ at all $\mathbf{p}$ for all $k, l \neq j \in \mathcal{J}$, then Assumption MTR is satisfied.*

Proof. See Appendix A.1. □

While MTR may not always tighten the bounds of the $\mathcal{HBL}$ parameters, it reduces the computational burden by eliminating potential outcomes paths. In the two-consumer, three-period example of section 3.3.1, Assumption MTR reduces the number of admissible potential outcome paths from 8 to 6, so that $\boldsymbol{Y}_2$ is only consistent with $\boldsymbol{\mu} \in (0, (1, 1), (\cdot, 1))$.

3.4.5 Testing the Identifying Restrictions

There is some scope for testing the identifying restrictions. First, we can test the null hypothesis that the identified set is non-empty under a specific set of restrictions: $\mathbb{H}_0 : \mathcal{P}^\star \neq \emptyset$. If we reject the null hypothesis, we conclude there is no empirical evidence for the specific combination of restrictions. We use Chernozhukov et al. (2015) for this specification testing.

Second, from Manski and Pepper (2000), we know that the combination of MTS and MTR generates the following testable condition:

$$\mathbb{P}[Y_{it} = j | Y_{i(t-1)} = j, Y_{i(t-2)} = \tilde{y}] \geq \mathbb{P}[Y_{it} = j | Y_{i(t-1)} \neq j, Y_{i(t-2)} = \tilde{y}]. \tag{3.17}$$

Condition 3.17 implies that the empirical choice propensity for a brand j is higher in state j than under some other loyalty state $k \neq j$.

3.5 Computing the Identified Set of HBL Parameters

Every probability mass function P in the identified set $\mathcal{P}^\star$ produces an HBL instance that is consistent with data and prior assumptions ($\rho(P) \geq 0$, including Assumptions ST, MTS and MTR). To get the identified set of $\hat{\mathcal{HBL}}$, we use the linear programming approach from Torgovitsky (2019). After extending the objective functions and constraints to the multinomial choice context as we discussed above, we can compute the identified set of $\hat{\mathcal{HBL}}$ as a closed interval $\Theta^\star \equiv [\theta^\star_{lb}, \theta^\star_{ub}]$ by solving two optimization problems,

$$\begin{aligned} \theta^\star_{lb} \equiv & \min_{\{P(\boldsymbol{\mu}) \in [0,1] : \boldsymbol{\mu} \in \mathcal{U}\}} \theta(P) \\ & s.t. \ \rho(P) \geq 0, \ (3.7), \text{ and } (3.9) \end{aligned} \tag{3.18}$$

$$\begin{aligned} \theta^\star_{ub} \equiv & \max_{\{P(\boldsymbol{\mu}) \in [0,1] : \boldsymbol{\mu} \in \mathcal{U}\}} \theta(P) \\ & s.t. \ \rho(P) \geq 0, \ (3.7), \text{ and } (3.9) \end{aligned} \tag{3.19}$$

where θ could be any linear combination of $\{\mathcal{HBL}_{jt}(P)\}_{j \in \mathcal{J},\ t=1,\ldots,T}$. In Appendix B.2, we derive the analogous computation for a version of the DPO model that conditions on the price vector, where the prior assumptions $\rho(P) \geq 0$ can include Assumptions IV and MIV.

4 A Test of Forward-Looking Behavior in Brand Loyalty

The DPO nests a broad class of random utility models, including myopic consumers who fully discount the future ($\beta = 0$) and passively respond to their brand habits, and forward-looking consumers ($\beta > 0$) who base their choices in part on their future brand loyalty. The extant empirical literature on habitual brand loyalty typically assumes myopic consumer behavior.[14] We propose additional parametric restrictions under which the discount factor, β, can be jointly identified with the consumer's deterministic flow utilities, $u_j(\mathbf{s}), \forall j \in \mathcal{D}$. Our identification results are constructive and immediately suggest a semiparametric minimum-distance estimator of the discount factor and test for forward-looking behavior.

4.1 Identification of the Discount Factor

At least since Rust (1994)'s non-identification result, the conventional wisdom has been that the discount factor β in the DDC is not identified without additional restrictions (e.g., Magnac and Thesmar (2002)). We now show that the DDC in section 2 has built-in exclusion restrictions of the form discussed in Abbring and Daljord (2020) that resolve this non-identification problem.

Following Hotz and Miller (1993), the inversion of (2.3) identifies the choice-specific value contrasts:

$$\begin{aligned} \ln\left(\frac{\sigma_j(\mathbf{s})}{\sigma_0(\mathbf{s})}\right) &= v_j(\mathbf{s}) - v_0(\mathbf{s}) \\ &= u_j(\mathbf{s}) + \beta[\mathbf{F}_j(\mathbf{s}) - \mathbf{F}_0(\mathbf{s})][\mathbf{I} - \beta\mathbf{F}_0]^{-1}\mathbf{m}. \end{aligned} \tag{4.1}$$

Even if beliefs $\mathbf{F}$ in (4.1) are observed, an additional moment is needed to disentangle the utilities, $u_j(\mathbf{s})$, and the discount factor, β.

Appendix B.1 shows that the canonical dynamic discrete-choice model of demand for differentiated products with HBL or switching costs has a built-in set of exclusion restrictions that identify β. The current-period utility only depends on the price of the chosen alternative, p_j, and not on the prices of rival products, $\mathbf{p}_{-j}$:

$$u_j\left(p_j, \mathbf{p}_{-j}, \ell\right) = u_j\left(p_j, \cdot, \ell\right).$$

Let $\mathcal{S}^{id}$ denote the set of pairs of states that generate excluded rival price variation:

$$\mathcal{S}^{id} = \left\{\{\mathbf{s}, \mathbf{s}'\} \in \mathcal{S}^2 : p_j = p'_j,\ \mathbf{p}_{-j} \neq \mathbf{p}'_{-j},\ \text{and } \ell = \ell' = l, l \neq j, \text{for } j \in \mathcal{J}\right\}. \tag{4.2}$$

If we assume that prices are exogenous, $\mathbb{E}(\epsilon_{jt}|p_{jt}) = 0$, we can construct a new set of

[14]See for instance Guadagni and Little (1983), Erdem (1996), Keane (1997), Seetharaman et al. (1999), and Dubé et al. (2010). One exception is Chintagunta et al. (2001) who estimate the reduced-form of the DDC in 2.3 without separately identifying β and preferences.

moment conditions by differencing the Hotz-Miller conditions (4.1) for each pair of states $(\mathbf{s}, \mathbf{s}') \in \mathcal{S}^{id}$:

$$\ln\left(\frac{\sigma_j(\mathbf{s})}{\sigma_0(\mathbf{s})}\right) - \ln\left(\frac{\sigma_j(\mathbf{s}')}{\sigma_0(\mathbf{s}')}\right) = \beta[\mathbf{F}_j(\mathbf{s}) - \mathbf{F}_0(\mathbf{s}) - \mathbf{F}_j(\mathbf{s}') + \mathbf{F}_0(\mathbf{s}')][\mathbf{I} - \beta\mathbf{F}_0]^{-1}\mathbf{m}. \quad (4.3)$$

Intuitively, a change in a competing product's price can only affect the value of choosing a given brand through the change in the consumer's beliefs about future prices and, hence, her expected continuation value.

The following proposition, based on Theorem 1 in Abbring and Daljord (2020), characterizes the identified set for β based on the moments (4.3).

Proposition 5. *Suppose that we observe pairs of states as in (4.2) and that either the left-hand side of (4.3) is nonzero (that is,* $\sigma_j(\mathbf{s})/\sigma_0(\mathbf{s}) \neq \sigma_j(\mathbf{s}')/\sigma_0(\mathbf{s}')$*) or* $[\boldsymbol{F}_j(\mathbf{s}) - \boldsymbol{F}_0(\mathbf{s}) - \boldsymbol{F}_j(\mathbf{s}') + \boldsymbol{F}_0(\mathbf{s}')][\boldsymbol{I} - \beta\boldsymbol{F}_0]^{-1}\boldsymbol{m} \neq 0$. *Then the identified set of discount factors is a closed discrete subset of* $[0, 1)$.

Proof. See Abbring and Daljord (2020). □

The rank condition in proposition 5 does not guarantee that β is point identified.[15] However, conditional on β, the utilities are point identified from the Hotz-Miller conditions (4.1).

The following proposition[16] establishes point identification when the choice-specific utilities do not depend on the loyalty state except when the consumer repeat-purchases the same brand, as is typically assumed in parametric models of switching costs:

$$u_j(\mathbf{p}, \ell) = u_j(\mathbf{p}, k), \text{ for } \ell \neq k \text{ and } \ell, k \neq j.$$

Proposition 6. *Suppose that (i)* $J \geq 3$*; (ii) prices follow a first-order Markov process that evolves independently of the loyalty state; and (iii) the set of pairs of states* $\tilde{\mathcal{S}}^{id}$ *is non-empty,*

$$\tilde{\mathcal{S}}^{id} = \left\{\{\mathbf{s}, \mathbf{s}'\} \in \mathcal{S}^2 : \mathbf{p} = \mathbf{p}' \text{ and } \ell \neq \ell', \text{ for every } \mathbf{p} \in \mathcal{P}\right\}. \quad (4.4)$$

Then the discount factor is point identified.

[15] Abbring and Daljord (2020) conjecture that when multiple exclusion restrictions are available, as in our empirical applications below, β will be point identified.

[16] We are grateful to Aureo de Paula for inspiring this result.

Proof. See Appendix A.2. □

4.2 A Minimum Distance Estimator for the Discount Factor

We construct the empirical analog of the moments (4.3) for each observed pair $(\mathrm{s},\mathrm{s}') \in \mathcal{S}^{id}$, which satisfies $p_a = p'_a,\ \mathbf{p}_{-a} \neq \mathbf{p}'_{-a},$ and $\ell = \ell' = l, l \neq a$, respectively for $a \in \mathcal{J}$,

$$\mathbf{g}(\beta) = \ln\left(\frac{\sigma_a(\mathbf{s})}{\sigma_0(\mathbf{s})}\right) - \ln\left(\frac{\sigma_a(\mathbf{s}')}{\sigma_0(\mathbf{s}')}\right) - \beta\left[\boldsymbol{F}_a(\mathbf{s}) - \boldsymbol{F}_0(\mathbf{s}) - \boldsymbol{F}_a(\mathbf{s}') + \boldsymbol{F}_0(\mathbf{s}')\right]\left[\boldsymbol{I} - \beta\boldsymbol{F}_0\right]^{-1}\boldsymbol{m}. \tag{4.5}$$

Although we are unable to prove point identification in general, Proposition 6 provides a commonly-used empirical specification where it would hold. We therefore define the minimum distance estimator

$$\beta^{MD} = \underset{\beta}{\operatorname{argmin}}\, \mathbf{g}(\beta)'\, \mathbf{W}\, \mathbf{g}(\beta) \tag{4.6}$$

where $\mathbf{W}$ is an $(R \times R)$ weight matrix and $\mathbf{g}(\beta)$ is the $(R \times 1)$ vector of moments. Inference follows directly from Newey and McFadden (1994).

5 Empirical Application: CPG Demand

5.1 Data

We conduct a case study of demand for consumer brands in the CPG industry using the IRI Academic Dataset.[17] The data comprise a shopping panel for 15,079 households spanning 12 years between 2001 and 2012. In addition, the data include weekly price and sales data at the UPC level for chain grocery and drug stores in those markets. Our analysis focuses on the period from 2004 to 2007, since this data period provides the largest sample with qualifying purchases and that a longer period is more likely to contain non-stationary changes in the purchase environment that may invalidate the stationarity assumption in the DPO model.

For our analysis, we select mature CPG product categories that are unlikely to exhibit much short-term consumer learning, are likely to exhibit discrete choices, are sold in

[17] See Bronnenberg et al. (2008)

standardized package sizes, and have an oligopolistic market structure comprising established brands.[18] For each category, we focus on the top-selling product form and package size. We then retain the top two to six brands based on purchase incidence. Horizontally-differentiated brand variants are aggregated if their price correlation is high. Within each category, we retain those households that satisfy minimal requirements for reporting, make in-category purchases only at stores with complete price records for all brands in the choice set, purchase at most one brand from the choice set each week, and make at least three in-category purchases in the sample period. We define the outside option as a trip during which none of the brands in the choice set is purchased.

After applying our screening criteria, we obtain panel data for the following 8 categories: coffee, deodorant, mayonnaise, margarine, peanut butter, spaghetti sauce, tooth brushes, and yogurt. Table 1 provides summary statistics of each of the 8 estimation samples.

category	brands	households	trips	purchases	brands
coffee	5	870	139.7	8.7	MAXWELL Sm, FOLGERS Sm, FOLGERS Lg, PL Sm, MAXWELL Lg
deodorant	6	1143	197.6	5.1	MENNEN, OLDSPICE, DOVE, DEGREE, RIGHT, SECRET
mayonnaise	3	2109	173	7	HELLMANNS, PL, KRAFT
margarine	4	1241	166.5	11	SMARTBALANCE, ICBINB, SHEDDS, PROMISE
peanut butter	4	4143	170.1	8.9	SKIPPY, JIF, PL, PETERPAN
spaghetti sauce	4	1899	185	11.5	RAGU, PREGO, FRANCESCORINALDI, HUNTS
toothbrushes	2	505	241.5	4.5	COLGATE, ORALB
yogurt	3	1457	168.4	13.3	YOPLAIT, DANNON, COLOMBO

This table summarizes the estimation samples drawn from the IRI Academic Dataset. The number of trips and number of purchases are computed as cross-household means.

Table 1: Descriptive Statistics

5.2 Empirical Evidence of Brand Loyalty

We now discuss our empirical estimates of HBL using the nonparametric DPO framework with the IRI data. The parameter of interest consists of the cross-time average share of observed consumer choices that were driven by HBL for a given brand j: $\mathcal{HBL}_j \equiv \frac{1}{T}\sum_{t=1}^{T}\mathcal{HBL}_{jt}(P)$, where $\mathcal{HBL}_{jt}$ is defined in (3.6). We compare the unrestricted bounds on $\mathcal{HBL}_j$ to the restricted cases using instruments, behavioral restrictions and combinations of the two. The estimated bounds and 95% confidence regions are constructed using the procedure of Torgovitsky (2019) and Chernozhukov et al. (2015), respectively. Since the MIV restriction using prices alone generates similar bounds to the fully unrestricted case,

[18] To ensure sufficient sample sizes, we also select categories in which the average household made 5 or more purchases in the sample period, except toothbrushes where the average was only 4 purchases.

we omit MIV from our results.

To address the curse of dimensionality from having $(J^{JT+1} \cdot |\mathcal{X}_{0:T}|$ potential outcomes paths when we condition on demand-shifters $\mathbf{X}_{0:T} \in \mathcal{X}_{0:T}$, we focus primarily on the binary choice between the top brand and a composite second choice comprising all other brands.[19] As a robustness check, we then report results for a trinomial DPO model ($J = 3$) that does not condition on demand-shifters and consists of the top two brands in each category and a third composite choice comprising all other brands.

Figure 1 and Table B1a report unconditional bounds for $\mathcal{HBL}_j$ under the binomial specification of the DPO. Figure 2 and Table B1b report the analogous $\mathcal{HBL}_j$ bounds under the trinomial formulation. These specifications already demonstrate the identifying role of behavioral restrictions. In four of our categories, the ST restriction generates lower bounds on $\mathcal{HBL}$ that are significantly larger than 0. When we additionally impose MTS, the upper bounds on $\mathcal{HBL}$ decline considerably, sometimes below 50%.

To impose the IV restriction, we need to condition on price and/or in-store marketing variables. When we combine IV with behavioral restrictions, we also need to adapt the ST restriction to ensure stationarity is not imposed on the data-generating process for prices or marketing variables. Conditioning on prices also reduces the concern that our HBL evidence is spuriously identified by non-stationarity in omitted demand shifting variables. We exclude the MTR restriction in the results below as it typically generated an empty identified set.

Figure 3 plots the bounds and confidence intervals for the binomial $\mathcal{HBL}_j$ with the IV restriction and the combination of the IV restriction and behavioral restrictions. See Table B2a in the Appendix for the complete set of estimated values. Using the lagged price vector as an IV restriction, the lower bounds on $\mathcal{HBL}$ are significantly larger than 0 in every category. In five of the categories, we find that HBL drives at least 15% of the choice behavior. Combining IV with the behavioral restrictions has a negligible impact on the lower bounds. But, the addition of behavioral restrictions, especially MTS, is informative, with HBL driving at most 62.7% of the choices in 6 of the categories. This finding is striking since many empirical applications assume a homogeneous loyalty (or "switching cost") parameter across consumers, which would likely overstate the extent of HBL in consumer behavior. Table B3a confirms that we obtain qualitatively similar results when

[19]To reduce computational dimensionality further, we follow Torgovitsky (2019)'s suggestion to construct shorter, overlapping DPO models. We also only use the final six observed purchases for each household ($T = 5$).

we instead use the vector of in-store display indicators as the excluded instrument.

Combining the IV, ST and MTS restrictions generates bounds indicating that at least 10.8% but no more than 72.2% of the consumers exhibit HBL across each of the product categories. These findings suggest that for many consumers, the psychological *switching costs* typically associated with HBL are the main driver of their choice behavior. These results have important economic implications for the market structure of CPG categories. A large body of work has studied how such *switching costs* can moderate the toughness of price competition (Farrell and Klemperer, 2007). In similar CPG categories as the ones used herein, such switching costs have been found to toughen price competition as firms' incentives to poach customers overwhelms their incentive to harvest current loyal consumers (Dubé et al., 2009b), leading to lower equilibrium prices.

In the data, the majority of observed purchases come from repeat buyers, ranging from 73.49% to 91.01% of observed purchases. Figure 4 decomposes the share of observed purchase persistence that is due to HBL by plotting bounds and confidence intervals conditional on the subpopulation of current repeat-buyers of the most popular brand, $\mathcal{HBL}_j|Y_t = j$ and $Y_{t-1} = j$. See Table B2b in the Appendix for the complete set of estimated values. For all categories, at least 10.5% of the choice-persistence for the top-selling brand (i.e., repeat purchases) is due to HBL; although we cannot rule out that the lower bounds are as low as 4.0%, 5.0% and 8.0% for peanut butter, spaghetti sauce, and yogurt categories, respectively, at the 5% significance level. Table B3b confirms that we obtain qualitatively similar results when we instead use the vector of in-store display indicators as the excluded instrument. From an economic perspective, a non-trivial proportion of a firm's retained customers appears to be due to HBL, as opposed to a persistent taste for the branded good. In the next section, we provide additional supporting evidence for the interpretation of this structural state-dependence as HBL by ruling out alternative mechanisms.

Table B4 in the Appendix summarizes the results of our tests of the IV and behavioral restrictions for the top brand. We fail to reject Manski and Pepper (2000)'s joint restriction of MTS and MTR in equation (3.17) across all 31 brands and 8 categories (although we only report the top brand in Table B4). Using Chernozhukov et al. (2015)'s test for a non-empty identified set, we mostly fail to reject each of the restrictions alone and in combination. An exception is the toothbrush category where we reject all but the IV restriction on its own. While failing to reject MTS is not conclusive evidence of this restriction, it is straightforward to show that MTS holds for the parametric DDC demand specification estimated for each of the IRI categories (results available upon request).

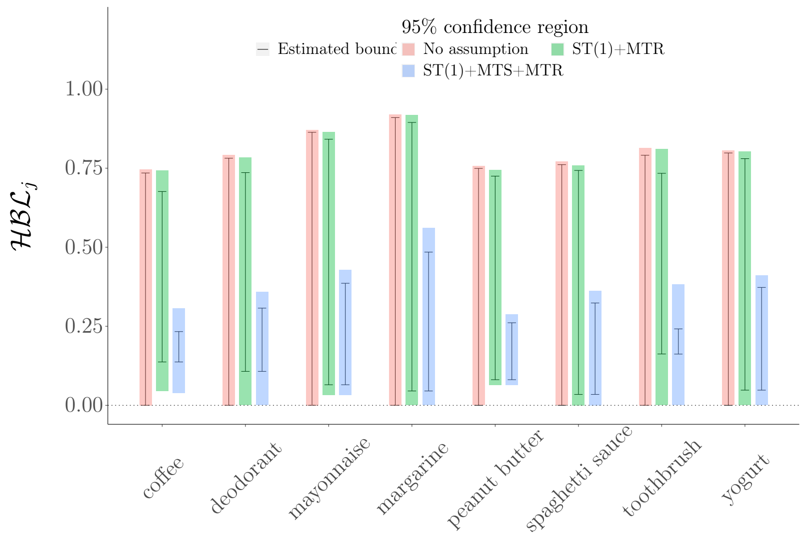

This figure presents $\mathcal{HBL}$ bounds and confidence intervals for the binary DPO under various sets of identifying restrictions using the IRI academic data. No marketing covariates are included in the binary DPO model. The bounds correspond to the share of the consumer population in each category who exhibit HBL for the top brand.

Figure 1: Binary DPO unconditional bounds on $\mathcal{HBL}$.

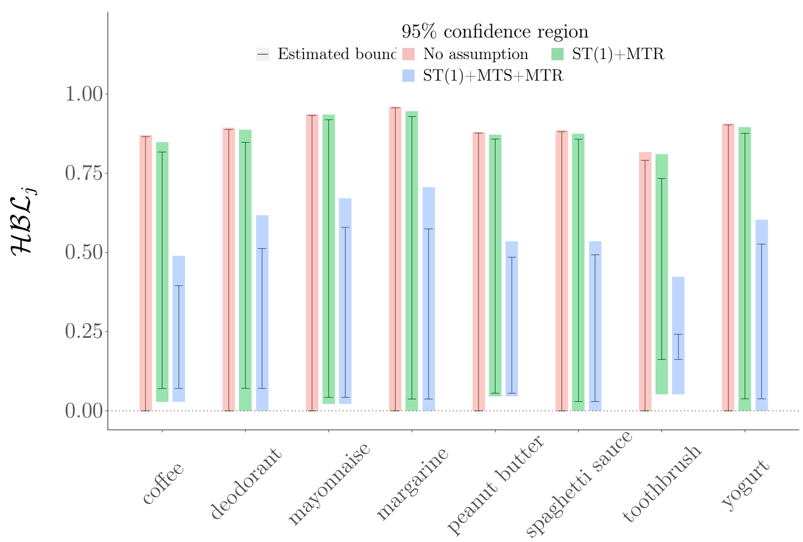

This figure presents $\mathcal{HBL}$ bounds and confidence intervals for the trinomial DPO under various sets of identifying restrictions using the IRI academic data. No marketing covariates are included in the trinomial DPO model. The bounds correspond to the share of the consumer population in each category who exhibit HBL for the focal brand.

Figure 2: Trinomial DPO unconditional bounds on $\mathcal{HBL}$.

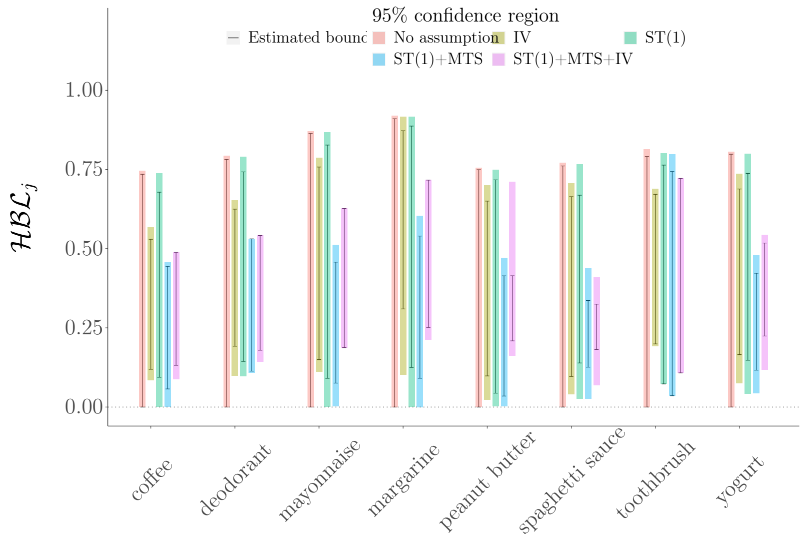

This figure presents $\mathcal{HBL}$ bounds and confidence intervals for the binary DPO under various sets of identifying restrictions using the IRI academic data. The restrictions are adapted to be conditional on realizations of the current prices of the top brand and that of its leading competitor. See appendix B.2.3 for details. The bounds correspond to the share of the consumer population in each category who exhibit HBL for the top brand.

Figure 3: Binary DPO bounds on $\mathcal{HBL}$ conditional on prices.

This figure presents $\mathcal{HBL}|Y_t = j, Y_{(t-1)} = j$ bounds and confidence intervals for the binary DPO under various sets of identifying restrictions using the IRI academic data. The restrictions are adapted to be conditional on realizations of the current prices of the top brand and that of its leading competitor. See appendix B.2.3 for details. The bounds correspond to the share of current repeat-buyers of the top brand in each category who exhibit HBL for the top brand.

Figure 4: Binary DPO bounds on $\mathcal{HBL}$ conditional on prices for repeat-purchasers of the top brand.

5.3 Testing HBL Mechanisms

We now provide more evidence for the structural interpretation of the state dependence as HBL – psychological switching costs – as opposed to learning (Osborne, 2009; Dubé et al., 2010) or search costs Dubé et al. (2010). We have already ruled out consumer learning since the evidence for HBL under the ST restriction explicitly rules out mechanisms which, like consumer learning, imply non-stationarity in the potential outcomes. While we cannot explicitly test ST, we typically fail to reject ST in our tests in section 5.2. Further, ST conditional on demand-shifting variables like prices is a commonly-used restriction in empirical papers using the parametric DDC or myopic choice model.

We also rule out search costs for the price of a previously-chosen item as the causal mechanism. In-store promotional displays for competing brands should eliminate their corresponding search costs and facilitate brand switching (e.g., Dubé et al., 2010). We therefore re-compute the $\mathcal{HBL}$ bounds for the subset of repeat-buyers of the most popular brand who are also exposed to a display for at least one other competing brand: $\mathcal{HBL}_j|Y_t = Y_{t-1} = j, \sum_{k \neq j} \mathbb{1}_{Display_{kt}} > 0$ where $\mathbb{1}_{Display_{kt}}$ indicates whether brand k is on display during trip t. Table 2 presents bounds and confidence intervals. In six categories, $\mathcal{HBL}$ for the most popular brand is bounded away from zero, suggesting that our evidence for HBL is robust to this control for search costs. For coffee, deodorant, peanut butter, and spaghetti sauce, we reject $\mathcal{HBL} = 0$ at the 5% significance level even after conditioning on competitor displays and at least 7.6% of the repeat purchases are still due to HBL.

5.4 Evidence of Forward-Looking Behavior

We now turn to the DDC specification and the empirical magnitude of consumers' discount factors, β. For our empirical application, we follow the convention in the literature and specify the transition process for the loyalty state ℓ as follows:

$$\ell_t = \begin{cases} j & \text{if } Y_{t-1} = j \in \mathcal{J}, \\ \ell_{t-1} & \text{if } Y_{t-1} = 0. \end{cases} \tag{5.1}$$

To accommodate the stochastic timing of trips, we use the cross-household mean probability of making a trip each week, π. We modify our moments (4.3) as follows:

$$\ln\left(\frac{\sigma_j(\mathbf{s})}{\sigma_0(\mathbf{s})}\right) - \ln\left(\frac{\sigma_j(\mathbf{s}')}{\sigma_0(\mathbf{s}')}\right) = \beta\pi\left(\mathbf{F}_j(\mathbf{s})\mathbf{B}_j - \mathbf{F}_0(\mathbf{s})\mathbf{B}_0 - \mathbf{F}_j(\mathbf{s}')\mathbf{B}_j + \mathbf{F}_0(\mathbf{s}')\mathbf{B}_0\right)\left[\mathbf{I} - \beta\pi\mathbf{F}_0\mathbf{B}_0\right]^{-1}\mathbf{m}$$

Category	Brand	(1) $\mathcal{HBL}_j \vert Y_t = Y_{t-1} = j$	(2) $\mathcal{HBL}_j \vert Y_t = Y_{t-1} = j, \sum_{k \neq j} \mathbb{1}_{Display_k t} > 0$
Coffee	Maxwell$_S$	[0.214,0.855]	[0.141,1.000]
		[0.136,0.908]	[0.125,1.000]
Deodorant	Mennen	[0.187,0.927]	[0.076,1.000]
		[0.053,1.000]	[0.011,1.000]
Mayonnaise	Hellmanns	[0.098,0.839]	[0.000,1.000]
		[0.097,0.839]	[0.000,1.000]
Margarine	Smart Balance	[0.129,0.652]	[0.054,1.000]
		[0.042,0.757]	[0.000,1.000]
Peanut Butter	Skippy	[0.144,0.734]	[0.196,0.917]
		[0.037,1.000]	[0.038,0.987]
Spaghetti Sauce	Ragu	[0.183,0.590]	[0.224,0.895]
		[0.079,0.930]	[0.080,1.000]
Toothbrush	Colgate	[0.055,1.000]	[0.000,1.000]
		[0.055,1.000]	[0.000,1.000]
Yogurt	Yoplait	[0.230,0.779]	[0.122,0.997]
		[0.070,0.878]	[0.000,1.000]

This table summarizes $\mathcal{HBL}_j | Y_t = j, Y_{t-1} = j, \sum_{k \neq j} \mathbb{1}_{Display_k t} > 1$ bounds and confidence intervals for the binary DPO under ST(1)+MTS+IV restriction using the IRI academic data. See appendix B.2.3 for details. The bounds correspond to the share of repeat-purchasing consumers of the top-selling brand j who exhibit HBL for j conditional on both j and at least one competitor brand being on display. We report $\mathcal{HBL}_j | Y_t = j, Y_{t-1} = j$ bounds as a benchmark.

Table 2: Binary DPO $\mathcal{HBL}$ for Current Repeat Buyers with Competitor Displays

where $\mathbf{B}_j = [\mathbf{I} - \beta(1-\pi)\mathbf{F}_j]^{-1}$. Finally, to control for time-invariant, unobserved heterogeneity, we use the group fixed-effects approach of Bonhomme et al. (2022) (see Appendix B.4). Although not reported herein, we find that the group fixed-effects and a Normal random coefficients approach generate comparable estimates of unobserved heterogeneity.

We start with our semiparametric estimates of β using the minimum distance estimator (4.6). Table 3 Columns (1) and (2) report the estimated population mean and standard error, respectively, for each category. Since most consumers do not shop in all the eight categories, it is difficult to interpret any heterogeneity in our estimates across categories due to selection. With the exception of mayonnaise, all the point estimates are significantly different from zero, implying that consumers are indeed forward-looking (i.e., we reject myopic choice behavior). However, our point estimates range from as low as 0.36 to very close to one. Given that we model a weekly choice, these results suggest that most categories have very impatient consumers.[20] Comparable degrees of consumer impatience have been documented in the extant literature (e.g., Frederick et al., 2002; Dubé et al., 2014).

[20]A weekly discount factor of $\beta = 0.9$, for instance, implies an annual interest rate of 238%.

category	(1) Min. Dist.	(2) Min. Dist. (s.e.)	(3) DDC	(4) DDC (s.e.)
coffee	0.691	0.183	0.799	0.167
deodorant	0.588	0.158	0.981	0.077
mayonnaise	0.357	0.201	0.867	0.057
margarine	0.979	0.127	0.874	0.062
peanut butter	0.999	0.115	0.845	0.113
spaghetti sauce	0.982	0.099	0.888	0.077
toothbrushes	0.542	0.234	0.671	0.238
yogurt	0.956	0.129	0.380	0.127

This table summarizes the discount factor estimates using the IRI academic data. "Min. Dist."indicates the semiparametric, minimum distance estimator. "DDC" indicates the dynamic discrete-choice model estimates. All specifications use the "group fixed-effects" approach to control for unobserved heterogeneity. Parameter values represent the population mean discount factor. Standard errors are computed via bootstrap.

Table 3: The Discount Factor β

We now turn to the parametric estimates from the DDC model. In practice, a parametric model would need to be estimated to conduct counterfactuals. Of interest is whether our parametric discount factor estimates agree with the semi-parametric estimates. For the empirical application, we assume:

$$u_j(\mathbf{s}_t) = \gamma_j - \alpha p_{jt} + \lambda \mathbb{I}_{\{\ell_t = j\}},\ \forall j \in \mathcal{J}$$

where $(\gamma_1 ...., \gamma_J, \alpha, \lambda)$ are parameters to be estimated. From Proposition 6, β is point-identified for this specification. We use an MPEC implementation of the maximum likelihood estimator. For details, see Appendix B.3. To control for unobserved heterogeneity, we once again use the group fixed-effects approach of Bonhomme et al. (2022) discussed in Appendix B.4.

Table 3 Columns (3) and (4) report the discount factor estimates from the DDC. All of our point estimates of β are statistically significant. Once again, we see a wide range of values, implying a fair amount of consumer impatience. For three of the eight categories, however, we reject the hypothesis that the minimum distance estimate and DDC estimate of β are equal (deodorant, mayonnaise, and yogurt).[21] It is possible that the parametric flow utilities are mis-specified in these categories.

We report the full set of taste coefficient estimates for the myopic random utility model and DDC specifications, respectively, in Tables B5 to B12 in Appendix B.7. We also

[21]We test this hypothesis using the z-statistic: $\frac{\beta^{MD} - \beta^{DDC}}{\sqrt{se^2_{MD} + se^2_{DDC}}}$ (Clogg et al., 1995)

compare specifications with versus without controls for unobserved heterogeneity (via group fixed-effects). As in the extant literature, the magnitude of the mean loyalty coefficient declines substantially in most categories once we control for unobserved heterogeneity. We nevertheless find economically significant coefficient estimates for brand loyalty, λ, even after controlling for unobserved heterogeneity. In six of the eight categories, the expected, dollar-denominated loyalty premium, $-\frac{\lambda}{\alpha}$, is over $\$0.6$. Surprisingly, we find almost no correlation between the estimated discount factor, β, and the other taste coefficients. Consequently, the taste coefficients for the myopic and dynamic models are similar in magnitude for most of the categories. Only the loyalty coefficient appears to change in some categories once we allow for forward-looking.

To compare our parametric results herein to the nonparametric results from section 5.2, we compute the parametric analog of the statistic $\mathcal{HBL}$, as in (3.6):

$$\mathcal{HBL}_j = \mathbb{P}_\epsilon[\exists k \neq j \in \mathcal{J} \text{ s.t. } j = \underset{l \in \mathcal{J}}{\text{argmax}}\{v_{il}(\mathbf{p}_t, j) + \epsilon_{ilt}\}, \text{ and } j \neq \underset{l \in \mathcal{J}}{\text{argmax}}\{v_{il}(\mathbf{p}_t, k) + \epsilon_{ilt}\}, \forall i, t].$$

See Appendix B.5 for details. Recall that HBL measures the propensity of consumers whose observed choice behavior is driven by loyalty, which is different from merely looking at the loyalty coefficient in a random utility model. For instance, someone with positive loyalty, $\lambda > 0$, may still not satisfy HBL if they also have very strong brand tastes, γ_j.

Figure 5 plots the estimated $\mathcal{HBL}$ bounds using the DPO and $\mathcal{HBL}$ point estimates using the DDC for the top-selling brand in each category.[22] See Table B13 in Appendix B.7 for the DDC point estimate values. When we assume homogeneous coefficients, both the DDC and myopic specifications generate $\mathcal{HBL}$ point estimates that lie well above the DPO's upper bounds in every category. In contrast, with the exception of toothbrushes, our DDC point estimates always lie within the DPO's empirical bounds, suggesting that the DDC's additional parametric assumptions are not inconsistent with the nonparametric evidence. We also look at a hybrid case of the myopic choice model that allows for heterogeneity on the price and brand coefficients but imposes homogeneity on the loyalty coefficient. For two categories, toothbrushes and margarine, the point estimates for $\mathcal{HBL}$ lie outside the DPO's bounds. Otherwise, the point estimates are quite close to those obtained when loyalty is also assumed to have heterogeneity. These results suggest that brand tastes and price sensitivity may in fact be driving most of the heterogeneity, leading to more reasonable $\mathcal{HBL}$ estimates (relative to the nonparametric DPO bounds) than specifications

[22] Our results are qualitatively similar if we use $\mathcal{HBL}$ bounds under the IV restrictions.

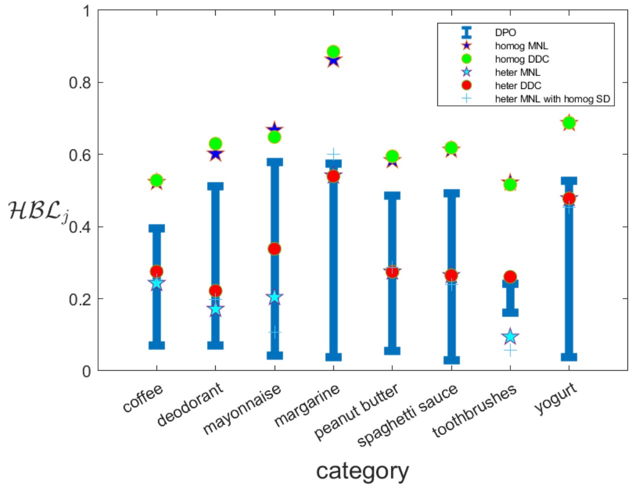

This table compares the $\mathcal{HBL}$ estimates using the nonparametric DPO and parametric DDC models estimated using the samples drawn from the IRI academic data. DPO imposes the restrictions ST(1), MTS and MTR.

Figure 5: $\mathcal{HBL}$ for most popular brand in each category using the DPO and DDC.

that do not allow for coefficient heterogeneity.

In sum, we find strong evidence for forward-looking consumer behavior across product categories, which is entirely consistent with the nonparametric HBL evidence using the DPO framework. However, consumers appear to be more myopic than would be implied by calibrating a discount factor using the real rate of interest. We explore the implications of this behavior in the next section.

6 Elasticities and Competitiveness

To assess the implications of consumer forward-looking behavior for the competitiveness of the market, we compare the short-run and long-run elasticity of demand. For each category, we simulate the long-term effect of a 10% increase in the price of the top-share brand in each state. We then compute the corresponding expected elasticities in equilibrium, using the steady-state probabilities of each state. For the short-run elasticities, we use the parameter estimates from the myopic discrete-choice model with $\beta = 0$. For the long-run elasticities, we use the DDC estimates, which allow β to vary freely in $[0, 1]$. See Appendix B.6 for the technical details regarding the elasticity simulation.

We report the results in Tables B14 to B21 in Appendix B.6. Even though the taste coefficients are fairly similar in the myopic discrete-choice model and DDC specifications, we nevertheless observe striking differences in the elasticities. With a few small exceptions, the long-run elasticities are all considerably larger in magnitude than the short-run elasticities, indicating that consumer demand is much more price-sensitive when consumers *plan* their future brand habits. The difference between the long-run and short-run own-price elasticities vary from close to 0, in mayonnaise, to 24.9%, in margarine. However, turning to the cross-price elasticities between brands, the long-run values are often hundreds or even thousands of percent higher than the short-run values. Intuitively, if consumers plan their future brand habits, then the "cost" of being loyal to the "wrong" brand is higher, making consumers more sensitive to permanent price changes. Given the similarities in the price coefficients and the fact that loyalty is unchanged by non-purchase, we see very small (typically less than 1%) differences in the effect on the outside good.

In sum, the myopic model appears to under-estimate the price-sensitivity of demand in response to permanent price changes. In particular, the myopic model drastically under-estimates the substitution between brands in the long term.

7 Conclusions

We have documented nonparametric evidence of *habitual brand loyalty* across several consumer goods product categories whereby a consumer's past brand choices have a causal effect on her current choice, similar to a "switching cost." We rule out alternative explanations for state dependence, such as learning and search costs. We have also documented semiparametric evidence that consumers make forward-looking purchase decisions that assign weight to their expected future loyalty habits. However, consumers appear to be more impatient than would be implied by the real rate of interest. Finally, using a parametric dynamic discrete-choice model of demand, we show that the long-run price elasticities of demand are considerably larger than would be implied by a myopic model with the deterministic restriction $\beta = 0$. In sum, the combination of habitual brand loyalty and forward-looking behavior likely have a material impact on demand, competition and market prices.

References

Jaap H Abbring and Øystein Daljord. Identifying the discount factor in dynamic discrete choice models. *Quantitative Economics*, 11(2):471–501, 2020.

T. W. Anderson and Leo A. Goodman. Statistical inference about markov chains. *The Annals of Mathematical Statistics*, 28(1):89–110, 1957.

Joe S. Bain. *Barriers to New Competition*. Cambridge: Harvard University Press, 1956.

Frank M. Bass, Moshe M. Givon, Manohar U. Kalwani, David Reibstein, and Gordon P. Wright. An investigation into the order of the brand choice process. *Marketing Science*, 3(4):267–287, 1984.

Gary S. Becker and Kevin M. Murphy. A theory of rational addiction. *Journal of Political Economy*, 1988.

Steven Berry, Amit Gandhi, and Philip Haile. Connected substitutes and invertibility of demand. *Econometrica*, 81(5):2087–2111, 2013.

Steven T. Berry and Giovanni Compiani. An instrumental variable approach to dynamic models. *Review of Economic Studies*, 90:1724–1758, 2023.

S. Bonhomme, T. Lamadon, and E. Manresa. Discretizing unobserved heterogeneity. *Econometrica*, 90(2):625–643, 2022.

J.W. Brehm. Post-decision changes in the desirability of alternatives. *Journal of Abnormal and Social Psychology*, 52:384–389, 1956.

Bart J Bronnenberg, Michael W Kruger, and Carl F Mela. Database paper—the iri marketing data set. *Marketing science*, 27(4):745–748, 2008.

Bart J. Bronnenberg, Jean-Pierre Dubé, and Matthew Gentzkow. The evolution of brand preferences: Evidence from consumer migration. *American Economic Review*, 102(6): 2472–2508, 2012.

G. H. Brown. Brand loyalty-fact or fiction. *The Trademark Reporter*, 43:251–258, 1953.

Gary Chamberlain. *Longitudinal Analysis of Labor Market Data*, chapter Heterogeneity, omitted variables, duration dependence. Cambridge University Press, 1985.

Victor Chernozhukov, Whitney K Newey, and Andres Santos. Constrained conditional moment restriction models. *arXiv preprint arXiv:1509.06311*, 2015.

Judith Chevalier and Austan Goolsbee. Are durable goods consumers forward-looking? evidence from college textbooks. *Quarterly Journal of Economics*, 124(4):1853–1854, 2009.

Pradeep Chintagunta, Ekaterini Kyriazidou, and Josef Perktold. Panel data analysis of household brand choices. *Journal of Econometric*, 103:111–153, 2001.

Clifford C Clogg, Eva Petkova, and Adamantios Haritou. Statistical methods for comparing regression coefficients between models. *American journal of sociology*, 100(5):1261–1293, 1995.

Carol Corrado, John Haltiwanger, and Daniel Sichel. Introduction to" measuring capital in the new economy". In *Measuring Capital in the New Economy*, pages 1–10. University of Chicago Press, 2005.

Ross M Cunningham. Brand loyalty-what, where, how much. *Harvard business review*, 34 (1):116–128, 1956.

Marnik G. Dekimpe, Dominique M. Hanssens, and Jorge M. Silva-Risso. Long-run eects of price promotions in scanner markets. *Journal of Econometrics*, 89:269–291, 1999.

Jean-Pierre Dubé, Günter J. Hitsch, and Peter E. Rossi. Do switching costs make markets less competitive? *Journal of Marketing Research*, 46:435–445, 2009a.

Jean-Pierre Dubé, Günter J. Hitsch, and Peter E. Rossi. Rejoinder to shin and sudhir and to cabral. *Journal of Marketing Research*, 46:451–452, 2009b.

Jean-Pierre Dubé, Günter J. Hitsch, and Peter E. Rossi. State dependence and alternative explanations for consumer inertia. *RAND Journal of Economics*, 41(3):417–445, 2010.

Jean-Pierre Dubé, Günter J. Hitsch, and Pranav Jindal. The joint identification of utility and discount functions from stated choice data: An application to durable goods adoption. *Quantitative Marketing and Economics*, 12:331–377, 2014.

James E. Egel, David T. Kollat, and Roger D. Blackwell. *Consumer Behavior*. Holt Rinehart and Winston, 1968.

Martin Eichenbaum, Nir Jaimovich, and Sergio Rebelo. Reference prices, costs, and nominal rigidities. *American Economic Review*, 101:4234–262, 2011.

Tulin Erdem. A dynamic analysis of market structure based on panel data. *Marketing Science*, 15(4):359–378, 1996.

Tülin Erdem and Michael P. Keane. Decision-making under uncertainty: Capturing dynamic brand choice processes in turbulent consumer goods markets. *Marketing Science*, 15(1):1–20, 1996.

Tülin Erdem and Baohong Sun. Testing for choice dynamics in panel data. *Journal of Business & Economic Statistics*, 19(2):142–152, 2001.

Tülin Erdem, Susumu Imai, and Michael P. Keane. Brand and quantity choice dynamics under price uncertainty. *Quantitative Marketing and Economics*, 1:5–64, 2003.

Joseph Farrell and Paul Klemperer. Coordination and lock-in: Competition with switching costs and network effects. In Mark Armstrong and R. Porter, editors, *Handbook of Industrial Organization*, volume 3, pages 1967–2072. Elsevier B.V., 2007.

Leon Festinger. *A Theory of Cognitive Dissonance*. Harper Row Publishers, Inc., 1957.

Ronald E. Frank. Brand choice as a probability process. *The Journal of Business*, 35(1): 43–56, 1962.

Shane Frederick, George Loewenstein, and Ted O'Donoghue. Time discounting and time preference: A critical review. *Journal of Economic Literature*, 40(2):351–401, 2002.

Peter M. Guadagni and John D.C. Little. A logit model of brand choice calibrated on scanner data. *Marketing Science*, 2:203–238, 1983.

L.P. Guest. Brand loyalty - twelve years later. *Journal of Applied Psychology*, 39:405–408, 1955.

Benjamin Handel. Adverse selection and switching costs in health insurance markets: When nudging hurts. *American Economic R*, 103:2643–2682, 2013.

Wesley R. Hartmann and V. Brian Viard. Do frequency reward programs create switching costs? a dynamic structural analysis of demand in a reward program. *Quantitative Marketing and Economics*, 6(2):109–137, 2008.

James J. Heckman. *Structural Analysis of Discrete Data and Econometric Applications*, chapter Statistical Models for Discrete Panel Data. The MIT Press, 1981.

James J Heckman and Edward J Vytlacil. Econometric evaluation of social programs, part i: Causal models, structural models and econometric policy evaluation. *Handbook of econometrics*, 6:4779–4874, 2007.

James J Heckman, Sergio Urzua, and Edward Vytlacil. Instrumental variables in models with multiple outcomes: The general unordered case. *Annales d'Economie et de Statistique*, pages 151–174, 2008.

Florian Heiss and Viktor Winschel. Likelihood approximation by numerical integration on sparse grids. *Journal of Econometrics*, 144:62–80, 2008.

Florian Heiss, Daniel McFadden, Joachim Winter, Amelie Wuppermann, and Bo Zhou. Inattention and switching costs as sources of inertia in medicare part d. *American Economic Review*, 111:2737–2781, 2021.

Igal Hendel and Aviv Nevo. Sales and consumer inventory. *RAND Journal of Economics*, 37(3):543–561, 2006.

Elisabeth Honka. Quantifying search and switching costs in the u.s. auto insurance industry. *Rand Journal of Economics*, 45:847–884, 2014.

Bo Honore and Ekaterini Kyriazidou. Panel data discrete choice models with lagged dependent variables. *Econometrica*, 68:839–874, 2000.

Bo E Honoré and Elie Tamer. Bounds on parameters in panel dynamic discrete choice models. *Econometrica*, 74(3):611–629, 2006.

V. Joseph Hotz and Robert A. Miller. Conditional choice probabilities and the estimation of dynamic models. *Review of Economic Studies*, 60(3):497–529, 1993.

Gaston Illanes. Switching costs in pension plan choice. *Northwestern University Working Paper*, 2017.

Mariano Irace. Patient loyalty in hospital choice: Evidence from new york. *University of Chicago, Becker Friedman Institute for Economics Working Paper*, (2018-52), 2018.

Michael P. Keane. Modeling heterogeneity and state dependence in consumer choice behavior. *Journal of Business & Economic Statistics*, 15(3):310–327, 1997.

Paul Klemperer. Competition when consumers have switching costs: An overview with applications to industrial organization, macroeconomics, and international trade. *Review of Economic Studies*, 62:515–539, 1995.

Alfred Arthur Kuehn. *An analysis of the dynamics of consumer behavior and its implications for marketing management*. PhD thesis, Carnegie Institute of Technology, 1958.

Sokbae Lee and Bernard Salanié. Identifying effects of multivalued treatments. *Econometrica*, 86(6):1939–1963, 2018.

Julia Levine and Stephan Seiler. Identifying state dependence in brand choice: Evidence from hurricanes. *Marketing Science*, 42(5):839–1028, 2023.

Thierry Magnac and David Thesmar. Identifying dynamic discrete decision processes. *Econometrica*, 70(2):801–816, 2002.

Charles F. Manski. *Identification Problems in the Social Sciences*. Harvard University Press, 1995.

Charles F. Manski. Monotone treatment response. *Econometrica*, 65:1311–1334, 1997.

Charles F. Manski and John V. Pepper. Monotone instrumental variables with an application to the returns to schooling. *Econometrica*, 68:997–1010, 2000.

Charles F Manski and John V Pepper. More on monotone instrumental variables. *The Econometrics Journal*, 12(suppl_1):S200–S216, 2009.

Klaus Miller, Navdeep Sahni, and Avner Schlain. Sophisticated consumers with inertia: Long-term implications from a large-scale field experiment. *HEC Paris Working Paper*, 2022.

Magne Mogstad, Andres Santos, and Alexander Torgovitsky. Using instrumental variables for inference about policy relevant treatment parameters. *Econometrica*, 86(5):1589–1619, 2018.

Francesca Molinari. Microeconometrics with partial identification. In Steven N. Durlauf, Lars Peter Hansen, James J. Heckman, and Rosa L. Matzkin, editors, *Handbook of Econometrics*, volume 7A, pages 355–486. Elsevier B.V., 2020.

Nickolay V. Moshkin and Ron Shachar. The asymmetric information model of state dependence. *Marketing Science*, 21(4):435–454, 2002.

Whitney K. Newey and Daniel McFadden. Large sample estimation and hypothesis testing. In Robert F. Engle and Daniel L. McFadden, editors, *Handbook of Econometrics*, volume IV, pages 2111–2245. Elsevier Science B.V., 1994.

Matthew Osborne. Consumer stockpiling behavior, price sensitivity and its implications for price elasticities and consumer welfare. *manuscript*, 2009.

Ariel Pakes, Jack R Porter, Mark Shepard, and Sophie Calder-Wang. Unobserved heterogeneity, state dependence, and health plan choices. Technical report, National Bureau of Economic Research, 2021.

Maria Polyakova. Regulation of insurance with adverse selection and switching costs: Evidence from medicare part d. *American Economic Journal: Applied Economics*, 8(3): 165–95, 2016.

John Rust. Optimal replacement of gmc bus engines: An empirical model of harold zurcher. *Econometrica*, 55(5):999–1033, 1987.

John Rust. Structural estimation of markov decision processes. In Robert F. Engle and Daniel L. McFadden, editors, *Handbook of Econometrics*, volume IV, pages 3081–3143. Amsterdam, London and New York: Elsevier, North-Holland, 1994.

Oleksandr Scherbakov. Measuring consumer switching costs in the television industry. *Rand Journal of Economics*, 47:366–393, 2016.

P.B. Seetharaman, Andrew Ainslie, and Pradeep Chintagunta. Investigating household state dependence effects across categories. *Marketing Science*, 36:488–500, 1999.

Matthew Shum. Does advertising overcome brand loyalty? evidence from the breakfast-cereals market. *Journal of Economics and Management Strategy*, 13(2):241–272, 2004.

Andrey Simonov, Jean-Pierre Dubé, Günter Hitsch, and Peter Rossi. State-dependent demand estimation with initial conditions correction. *Journal of Marketing Research*, 57(5):789–809, 2020.

Alexander Torgovitsky. Nonparametric inference on state dependence in unemployment. *Econometrica*, 87(5):1475–1505, 2019.

A Appendix

A.1 Proofs of the DPO Model

Proof of Proposition 1: Define $\boldsymbol{U}_{it}(-j) \equiv (U_{it}(1), ..., U_{it}(j-1), U_{it}(j+1), ..., U_{it}(J))$, $\boldsymbol{U}_{it}(-jk) \equiv (U_{it}(1), ..., U_{it}(j-1), U_{it}(j+1), ..., , U_{it}(k-1), U_{it}(k+1), ..., U_{it}(J))$ and $\mathcal{A}_j^k \equiv \mathcal{J}^k \setminus \{\underbrace{j, ..., j}_{k}\}$ for $k \in 1, 2, ..., J$. If $P \in \mathcal{P}^\star$, we can decompose $\mathcal{HLB}_{jt}$ from (3.6)

$$\mathcal{HLB}_{jt} = \mathbb{P}_P[Y_{i(t-1)} = j, Y_{it} = j, \boldsymbol{U}_{it}(-j) \in \mathcal{A}_j^{J-1}] + \mathbb{P}_P[Y_{i(t-1)} \neq j, U_{it}(j) = j, \boldsymbol{U}_{it}(-j) \in \mathcal{A}_j^{J-1}] \tag{A.1}$$

since the event $\{Y_{i(t-1)} = j, U_{it}(j) = j\}$ occurs if and only if $\{Y_{i(t-1)} = j, Y_{it} = j\}$.

When $|\mathcal{J}| \geq 3$, we can rewrite (A.1) by decomposing Y_{it} as follows

$$\begin{aligned}\mathcal{HBL}_{jt} =& \mathbb{P}_P[Y_{i(t-1)} = j, Y_{it} = j, \boldsymbol{U}_{it}(-j) \in \mathcal{A}_j^{J-1}] \\ &+ \sum_{l \neq j} \mathbb{P}_P[Y_{i(t-1)} = l, Y_{it} = j, U_{it}(j) = j, \boldsymbol{U}_{it}(-jl) \in \mathcal{A}_j^{J-2}] \\ &+ \sum_{l \neq j} \mathbb{P}_P[Y_{i(t-1)} = l, Y_{it} \neq j, U_{it}(j) = j, \boldsymbol{U}_{it}(-jl) \in \mathcal{J}^{J-2}].\end{aligned}$$

Since we only impose the following restrictions for all $l \neq j$,

$$\begin{aligned}&0 \leq \mathbb{P}_P[Y_{i(t-1)} = j, Y_{it} = j, U_{it}(k) = k] \leq \mathbb{P}[Y_{i(t-1)} = j, Y_{it} = j] \\ &0 \leq \mathbb{P}_P[Y_{i(t-1)} = l, Y_{it} = j, U_{it}(j) = j, \boldsymbol{U}_{it}(-jl) \in \mathcal{A}_j^{J-2}] \leq \mathbb{P}[Y_{i(t-1)} = l, Y_{it} = j] \\ &0 \leq \mathbb{P}_P[Y_{i(t-1)} = l, Y_{it} \neq j, U_{it}(j) = j, \boldsymbol{U}_{it}(-jl) \in \mathcal{J}^{J-2}] \leq \mathbb{P}[Y_{i(t-1)} = l, Y_{it} \neq j]\end{aligned}$$

we can find a $P \in \mathcal{P}^\star$ that attains all upper bounds and one that attains all lower bounds. Therefore, $\Theta^\star = [0, \mathbb{P}[Y_{i(t-1)} = j, Y_{it} \neq j] + \mathbb{P}[Y_{i(t-1)} \neq j]]$.

For $|\mathcal{J}| = 2$, we need a small modification. Let $k = \mathcal{J} \setminus j$ and simplify $\mathcal{HBL}_{jt}$ to be

$$\mathcal{HBL}_{jt} = \mathbb{P}_P[Y_{i(t-1)} = j, Y_{it} = j, U_{it}(k) = k] + \mathbb{P}_P[Y_{i(t-1)} = k, Y_{it} = k, U_{it}(j) = j].$$

Observational equivalence requires that

$$0 \leq \mathbb{P}_P[Y_{i(t-1)} = j, Y_{it} = j, U_{it}(k) = k] \leq \mathbb{P}[Y_{i(t-1)} = j, Y_{it} = j]$$

$$0 \leq \mathbb{P}_P[Y_{i(t-1)} = k, Y_{it} = k, U_{it}(j) = j] \leq \mathbb{P}[Y_{i(t-1)} = k, Y_{it} = k].$$

We can find a $P \in \mathcal{P}^\star$ that attains all upper bounds and one that attains all lower bounds. Therefore, we have $\Theta^\star = [0, \mathbb{P}[Y_{i(t-1)} = j, Y_{it} = j] + \mathbb{P}[Y_{i(t-1)} = k, Y_{it} = k]]$. $\square$

Proof of Proposition 2: For every $t, t' \geq r+1$, and candidate r-period potential outcomes path $\boldsymbol{\mu} \equiv (\mu_0(1), ..., \mu_0(J), ..., \mu_r(1), ..., \mu_r(J)) \in \{1, ..., J\}^{J(r+1)}$,

$$\begin{aligned}
&\mathbb{P}_P\left[U_{i(t-r)}(1) = \mu_0(1), ..., U_{it}(J) = \mu_r(J)\right] \\
&\quad = \int \prod_{\tau=t-r}^{t} \prod_{k=1}^{J} \mathbb{1}[\mu_{\tau-t+r}(k) = \operatorname*{argmax}_j \{v_j(\boldsymbol{x}_\tau, k) + \epsilon_{j\tau}\}] d\mathbf{F}_{\boldsymbol{X}}(\boldsymbol{x}_{t-r}, ..., \boldsymbol{x}_t) d\mathbf{F}_{\boldsymbol{\epsilon}}(\boldsymbol{\epsilon}_{t-r}, ..., \boldsymbol{\epsilon}_t) \\
&\quad = \int \prod_{\tau=t'-r}^{t'} \prod_{k=1}^{J} \mathbb{1}[\mu_{\tau-t'+r}(k) = \operatorname*{argmax}_j \{v_j(\tilde{\boldsymbol{x}}_\tau, k) + \epsilon_{j\tau}\}] d\mathbf{F}_{\boldsymbol{X}}(\boldsymbol{x}_{t'-r}, ..., \boldsymbol{x}_{t'}) d\mathbf{F}_{\boldsymbol{\epsilon}}(\boldsymbol{\epsilon}_{t'-r}, ..., \boldsymbol{\epsilon}_{t'}) \\
&\quad = \mathbb{P}_P[U_{i(t'-r)}(1) = \mu_0(1), ..., U_{it'}(J) = \mu_r(J)] \quad \square
\end{aligned}$$

Proof of Proposition 3: It's straightforward to show that

$$\mathbb{P}_P[U_{it}(y) = j, Y_{i(t-1)} = j | Y_{i(t-2)} = \tilde{y}] \geq \mathbb{P}_P[U_{it}(y) = j | Y_{i(t-2)} = \tilde{y}] \cdot \mathbb{P}_P[U_{i(t-1)}(\tilde{y}) = j | Y_{i(t-2)} = \tilde{y}].$$

is a sufficient condition for Assumption MTS. Therefore, Assumption MTS is satisfied if

$$\begin{aligned}
&\mathbb{P}[\operatorname*{argmax}_k \{v_k(\mathbf{p}_t, y) + \epsilon_{kt}\} = j, \operatorname*{argmax}_k \{v_k(\mathbf{p}_{t-1}, \tilde{y}) + \epsilon_{k(t-1)}\} = j | Y_{i(t-2)} = \tilde{y}] \\
&\quad \geq \mathbb{P}[\operatorname*{argmax}_k \{v_k(\mathbf{p}_t, y) + \epsilon_{kt}\} = j | Y_{i(t-2)} = \tilde{y}] \cdot \mathbb{P}[\operatorname*{argmax}_k \{v_k(\mathbf{p}_{t-1}, \tilde{y}) + \epsilon_{k(t-1)}\} = j | Y_{i(t-2)} = \tilde{y}]
\end{aligned} \tag{A.2}$$

Denoting $\Omega(\boldsymbol{\epsilon}, \mathbf{p}, y) \equiv \epsilon_{jt} - \max_k \{v_k(\mathbf{p}, y) - v_j(\mathbf{p}, y) + \epsilon_{kt}\}$, we can rewrite (A.2) as

$$\begin{aligned}
&\mathbb{P}[\Omega(\boldsymbol{\epsilon}_t, \mathbf{p}_t, y) \geq 0, \Omega(\boldsymbol{\epsilon}_{t-1}, \mathbf{p}_{t-1}, \tilde{y}) \geq 0 | Y_{i(t-2)} = \tilde{y}] \\
&\quad \geq \mathbb{P}[\Omega(\boldsymbol{\epsilon}_t, \mathbf{p}_t, y) \geq 0 | Y_{i(t-2)} = \tilde{y}] \cdot \mathbb{P}[\Omega(\boldsymbol{\epsilon}_{t-1}, \mathbf{p}_{t-1}, \tilde{y}) \geq 0 | Y_{i(t-2)} = \tilde{y}].
\end{aligned} \tag{A.3}$$

We can re-write the LHS of (A.3) as $\mathbb{E}\left[\sigma_j(\mathbf{p}_t, y) \cdot \sigma_j(\mathbf{p}_{t-1}, \tilde{y}) | Y_{i(t-2)} = \tilde{y}\right]$ since $Y_{i(t-1)}$ is determined by Y_{i0}, the history of $\boldsymbol{\epsilon}$, and $\mathbf{p}$, and $\boldsymbol{\epsilon}_t$ is $i.i.d.$.

If $\sigma_j(\mathbf{p}_t, y)$ and $\sigma_j(\mathbf{p}_{t-1}, \tilde{y})$ have weakly positive covariance, (A.2) holds:

$$
\begin{aligned}
&\mathbb{P}[\underset{k}{\text{argmax}}\{v_k(\mathbf{p}, y) + \epsilon_{kt}\} = j, \underset{k}{\text{argmax}}\{v_k(\mathbf{p}, \tilde{y}) + \epsilon_{k(t-1)}\} = j | Y_{i(t-2)} = \tilde{y}] \\
&\quad \geq \mathbb{E}_{\mathbf{p}_t}\left[\sigma_j(\mathbf{p}_t, y) | Y_{i(t-2)} = \tilde{y}\right] \cdot \mathbb{E}_{\mathbf{p}_{t-1}}\left[\sigma_j(\mathbf{p}_{t-1}, \tilde{y}) | Y_{i(t-2)} = \tilde{y}\right] \\
&\quad = \mathbb{P}[\underset{k}{\text{argmax}}\{v_k(\mathbf{p}, y) + \epsilon_{kt}\} = j | Y_{i(t-2)}] \cdot \mathbb{P}[\underset{k}{\text{argmax}}\{v_k(\mathbf{p}, \tilde{y}) + \epsilon_{k(t-1)}\} = j | Y_{i(t-2)} = \tilde{y}]. \quad \square
\end{aligned}
$$

Proof of Proposition 4: A sufficient condition for Assumption MTR is that

$$
\mathbb{P}\left[\mathbb{1}[j = \underset{l}{\text{argmax}}\{v_l(\mathbf{p}_t, j) + \epsilon_{lt}\}] \geq \mathbb{1}[j = \underset{l}{\text{argmax}}\{v_l(\mathbf{p}_t, k) + \epsilon_{lt}\}]\right] = 1. \tag{A.4}
$$

In the DDC model,

$$
v_l(\mathbf{p}, j) - v_l(\mathbf{p}, k) = u_l(\mathbf{p}, j) - u_l(\mathbf{p}, k), \text{ for } j, k, l \in \mathcal{J} \tag{A.5}
$$

since prices are exogenous and loyalty evolves deterministically after a purchase.

If $u_j(\mathbf{p}, j) - u_j(\mathbf{p}, k) \geq 0$ and $u_l(\mathbf{p}, j) - u_l(\mathbf{p}, k) \leq 0, \forall p \in \mathcal{P}, k, l \in \mathcal{D} \setminus \{0, j\}$, then exogenously switching loyalty state from k to j weakly increases the choice-specific value function of j and weakly decreases choice-specific value functions of all other alternatives. Therefore, $\mathbb{1}[j = \text{argmax}_l\{v_l(\mathbf{p}_t, j) + \epsilon_{lt}\}] \geq \mathbb{1}[j = \text{argmax}_l\{v_l(\mathbf{p}_t, k) + \epsilon_{lt}\}]$ holds. $\square$

A.2 Proof of Proposition 6

Proof. For $\mathbf{p} \in \mathcal{P}$, the moment corresponding to $\{(\mathbf{p}, k), (\mathbf{p}, \ell) | k, \ell \neq j\} \in \tilde{\mathcal{S}}^{id}$ is

$$
\begin{aligned}
\ln \frac{\sigma_j(\mathbf{p}, k)}{\sigma_0(\mathbf{p}, k)} - \ln \frac{\sigma_j(\mathbf{p}, \ell)}{\sigma_0(\mathbf{p}, \ell)} &= -v_0(\mathbf{p}, k) + v_0(\mathbf{p}, \ell) \\
&= \beta \sum_{\mathbf{p}_r \in \mathcal{P}} [-v(\mathbf{p}_r, k) + v(\mathbf{p}_r, \ell)] G(\mathbf{p}_r | \mathbf{p}) \\
&= \beta \sum_{\mathbf{p}_r \in \mathcal{P}} [\ln \sigma_j(\mathbf{p}_r, k) - \ln \sigma_j(\mathbf{p}_r, \ell)] G(\mathbf{p}_r | \mathbf{p}) \\
&= \beta \mathbf{G}(\mathbf{p}) \left[\ln \frac{\sigma_j(\mathbf{p}_1, k)}{\sigma_j(\mathbf{p}_1, \ell)}, ..., \ln \frac{\sigma_j(\boldsymbol{p}_{L_p}, k)}{\sigma_j(\mathbf{p}_{L_p}, \ell)}\right]'
\end{aligned} \tag{A.6}
$$

where the first line follows from $u_j(\mathbf{p}, k) = u_j(\mathbf{p}, \ell)$ and assumption (ii), which implies $\mathbf{F}_j(\mathbf{p}, k) = \mathbf{F}_j(\mathbf{p}, \ell)$. Inverting (A.6) point identifies β. $\square$

B Online Appendix

B.1 DDC and the exclusion restriction

Consider a consumer with budget y and preferences over its consumption $q_j \geq 0$ of a set of $j \in \{1, ..., J\}$ market goods and an essential numeraire, z with price $p_z = 1$:

$$U(\mathbf{q}, z) = \tilde{U}(\sum_{j \in \mathcal{D}} \gamma_j q_j, z)$$

where $\gamma_j \geq 0$ is the constant marginal utility per unit of product j. Assume $\gamma_0 = 0$ and $p_0 = 0$. The perfect substitutes preferences over the market goods ensure that at most one product will have an interior optimal quantity (discrete-choice). If the market goods also have unit-inelastic demand, $q_j \in 0, 1, \ \forall j \in \mathcal{D}$, then the consumer purchases one unit of the product with the highest choice-specific values

$$v_j = \tilde{U}(\gamma_j, y - p_j), \ \forall j \in \mathcal{D}.$$

The choice-specific value of a product j does not depend on any of the other prices, $\mathbf{p_{-j}}$.

B.2 Estimation and Statistical Inference

Our estimation and statistical inference closely follow the estimation and inference appendix of Torgovitsky (2019). We explicitly write down the complete optimization problem, and explain several important modifications based on Torgovitsky (2019). The modifications include concatenating shorter DPO models to describe longer observed choice sequences, and constructing identifying assumptions conditional on time-varying covariates. We explain how to construct a consistent estimator of the $\mathcal{HBL}$ parameters, the confidence regions, as well as a mis-specification for falsifying the model.

B.2.1 The complete optimization problem

The primitive object is a collection of probability mass functions $P \equiv \{P_{t_0}\}_{t_0=0}^{T-ML}$, each of which describes the joint distribution of

$$(Y_{it_0}, U_{i(t_0+1)}(0), ..., U_{i(t_0+ML)}(J))$$

The set of all potential outcome paths $\mathcal{U}$ is $\{1, ..., J\}^{J*ML+1}$, and the set of all observed choice sequences $\mathcal{Y}$ is $\{1, ..., J\}^{ML+1}$

$$
\begin{aligned}
\theta^{\star}_{lb} \equiv & \min_{P} \theta(P) \\
s.t. \quad & P_{t_0}(\boldsymbol{\mu}) \in [0,1] \text{ for all } \boldsymbol{\mu} \in \mathcal{U}, t_0 \in \{0, ..., T-ML\} \\
& \sum_{\boldsymbol{\mu} \in \mathcal{U}} P_{t_0}(\boldsymbol{\mu}) = 1 \text{ for all } t_0 \in \{0, ..., T-ML\} \\
Observational\ equivalence \quad & \sum_{\boldsymbol{\mu} \in \mathcal{U}_{oeq}} P_{t_0}(\boldsymbol{\mu}) = \mathbb{P}[\boldsymbol{Y}_{i(t_0:t_0+ML)} = \boldsymbol{y}] \\
& \quad \text{for all } \boldsymbol{y} \in \mathcal{Y} \text{ and all } t_0 \in \{0, ..., T-ML\} \\
ST(1) \quad & \mathbb{P}_{P_{t_0}}[\boldsymbol{U}_{i(t-1:t)}(1) = \boldsymbol{\mu}_1(1), ..., \boldsymbol{U}_{i(t-1:t)}(J) = \boldsymbol{\mu}_1(J)] \\
& \quad = \mathbb{P}_{P_{t_0}}[\boldsymbol{U}_{i(t'-1:t')}(1) = \boldsymbol{\mu}_1(1), ..., \boldsymbol{U}_{i(t'-1:t')}(J) = \boldsymbol{\mu}_1(J)] \\
& \quad \text{for every } \boldsymbol{\mu}_1 \equiv (\boldsymbol{\mu}_1(1), ..., \boldsymbol{\mu}_1(J)) \in \mathcal{J}^{J*ML+1}, \\
& \quad \text{every } t_0 + 2 \leq t, t' \leq t_0 + ML, \\
& \quad \text{and all } t_0 \in \{0, ..., T-ML\} \\
MTS \quad & \mathbb{P}_{P_{t_0}}[U_{it}(y) = j | Y_{i(t-1)} = j, Y_{i(t-2)} = \tilde{y}] \\
& \quad \geq \mathbb{P}_{P_{t_0}}[U_{it}(y) = j | Y_{i(t-1)} \neq j, Y_{i(t-2)} = \tilde{y}] \\
& \quad \text{for every } t_0 \in \{0, ..., T-ML\}, \\
& \quad \text{for every } y, \tilde{y} \in \mathcal{J}, \\
& \quad \text{and all } t_0 + 2 \leq t \leq t_0 + ML \text{ such that } \mathbb{P}[Y_{i(t-1)} = j | Y_{i(t-2)} = \tilde{y}] \in (0,1) \\
MTR \quad & \mathbb{P}_{P_{t_0}}[\exists k \neq j : U_{it}(j) \neq j,\ U_{it}(k) = j] = 0 \\
& \quad \text{for every } t_0 \in \{0, .., T-ML\}, \\
& \quad \text{and every } \ \leq t \leq t_0 + ML \\
Coherency \quad & \mathbb{P}_{P_{t_0}}[Y_{i(t_0+1)} = y_0, \boldsymbol{U}_{i(t_0+2:t_0+ML)} = \boldsymbol{\mu}] \\
& \quad = \mathbb{P}_{P_{t_0+1}}[Y_{i(t_0+1)} = y_0, \boldsymbol{U}_{i(t_0+2:t_0+ML)} = \boldsymbol{\mu}] \\
& \quad \text{for all } (y_0, \boldsymbol{\mu}) \in \mathcal{J}^{1+J(ML-1)}
\end{aligned}
\tag{B.1}
$$

B.2.2 Combining Shorter Models

We face the curse of dimensionality problem since there are $J^{(JT+1)}$ probabilities, each associated with a potential outcomes path. Even after reducing the choice set to a trinomial problem ($J = 3$), we still face $3^{(3*5+1)}$, i.e., over 43 million, potential outcome paths for choice sequences of $T = 5$. We further address the dimensionality problem by constructing shorter, overlapping models to describe longer observed choice sequences.

The solution starts from fixing a model length $ML \in \{2, ..., T\}$. We then construct a DPO model for the observed sequence $(Y_{it_0}, Y_{i(t_0+1)}, ..., Y_{i(t_0+ML)})$ for every initial period $t_0 \in \{0, 1, ..., T - ML\}$. The shorter, overlapping models P_{t_0} relate potential outcomes to an observed choice sequence starting from initial period t_0 to the end period $t_0 + ML$ via the recursive model (3.1). The identified set is a collection of shorter models $P \equiv \{P_{t_0} : P_{t_0} \in \mathcal{P}^{\dagger}_{t_0}\}_{t_0=0}^{T-ML}$, such that P_{t_0} is a probability mass function, satisfies the identifying assumptions, and rationalizes the (shorter) observed choice sequences, for every $t_0 \in \{0, ..., T - ML\}$. The $\mathcal{HBL}$ parameter θ is a function of the collection of shorter models P.

We impose an additional *Coherency* restriction to ensure that two overlapping models, for example, P_{t_0} and P_{t_0+1}, are logically consistent. More specifically, the coherency condition states that when two overlapping models can both assign a probability to an event, this probability must be the same regardless of which model it is generated from. That is,

$$\mathbb{P}_{P_{t_0}}[Y_{i(t_0+1)} = y_0, \boldsymbol{U}_{i(t_0+2:t_0+ML)} = \boldsymbol{\mu}] = \mathbb{P}_{P_{t_0+1}}[Y_{i(t_0+1)} = y_0, \boldsymbol{U}_{i(t_0+2:t_0+ML)} = \boldsymbol{\mu}]$$
$$\text{for all } (y_0, \boldsymbol{\mu}) \in \mathcal{J}^{J*(ML-1)+1}$$

In Section 5.3, we implement the estimation with $T = 5$ and $ML = 3$ for all versions of the DPO model except the binomial DPO, where no marketing covariates are included.

B.2.3 Incorporating Marketing Variables

We can re-define the $\mathcal{HBL}$ parameter and the identifying assumptions conditional on marketing variables. In addition to observing choice sequences $\boldsymbol{Y}_i = (Y_{i0}, Y_{i1}, ..., Y_{iT})$, the researcher also observes a vector $\boldsymbol{X}_i = (\boldsymbol{X}_{i0}, \boldsymbol{X}_{i1}, ..., \boldsymbol{X}_{iT})$ with support $\mathcal{X}$. The primitive of the DPO model becomes a probability mass function P defined over the support of $\mathcal{U} \times \mathcal{X}$, that is,

$$\mathcal{P} = \left\{ P : \mathcal{U} \times \mathcal{X} \to [0, 1] \quad s.t. \sum_{\boldsymbol{\mu} \in \mathcal{U}, \boldsymbol{x} \in \mathcal{X}} P(\boldsymbol{\mu}, \boldsymbol{x}) = 1 \right\}$$

The observational equivalence condition is re-defined as

$$\mathbb{P}[\boldsymbol{Y}_i = \boldsymbol{y}, \boldsymbol{X}_i = \boldsymbol{x}] = \sum_{\boldsymbol{\mu} \in \mathcal{U}_{oeq}(\boldsymbol{y})} P(\boldsymbol{\mu}, \boldsymbol{x})$$

The ST, MTS, and MTR assumptions are adapted to be

$$\begin{aligned} &\mathbb{P}_P[\boldsymbol{U}_{i(t-m:t)}(1) = \boldsymbol{\mu}_m(1), ..., \boldsymbol{U}_{i(t-m:t)}(J) = \boldsymbol{\mu}_m(J) | \boldsymbol{X}_{i(t-m:t)} = \boldsymbol{x}_m] \\ &\quad = \mathbb{P}_P[\boldsymbol{U}_{i(t'-m:t')}(1) = \boldsymbol{\mu}_m(1), ..., \boldsymbol{U}_{i(t'-m:t')}(J) = \boldsymbol{\mu}_m(J) | \boldsymbol{X}_{i(t'-m:t')} = \boldsymbol{x}_m] \quad \text{(ST)} \end{aligned}$$

for every $\boldsymbol{\mu}_m \in \mathcal{J}^{J(m+1)}$, every $\boldsymbol{x}_m \in \mathcal{X}_m$ where $\mathcal{X}_m$ is the support of $\boldsymbol{X}_{i(s:s+m)}$, and every $t, t' \geq m+1$.

$$\begin{aligned} &\mathbb{P}_P[U_{it}(y) = j | Y_{i(t-1)} = j, Y_{i(t-2)} = \tilde{y}, \boldsymbol{X}_{it} = \boldsymbol{x}, \boldsymbol{X}_{i(t-1)} = \tilde{\boldsymbol{x}}] \\ &\quad \geq [U_{it}(y) = j | Y_{i(t-1)} \neq j, Y_{i(t-2)} = \tilde{y}, \boldsymbol{X}_{it} = \boldsymbol{x}, \boldsymbol{X}_{i(t-1)} = \tilde{\boldsymbol{x}}] \quad \text{(MTS)} \end{aligned}$$

for any $y, \tilde{y} \in \mathcal{J}$, any $\boldsymbol{x}, \tilde{\boldsymbol{x}} \in \mathcal{X}_0$ where $\mathcal{X}_0$ is the support of $\boldsymbol{X}_{is}$, and all $t \geq 2$.

$$\mathbb{P}_P[\mathbb{1}\{U_{it}(j) = j\} \geq \mathbb{1}\{U_{ik} = j\} | \boldsymbol{X}_i = \boldsymbol{x}] = 1 \quad \text{(MTR)}$$

for any $k \neq j$, any $\boldsymbol{x} \in \mathcal{X}$, and all t.

The $Coherency$ restrictions in Appendix B.2.2 are similarly adapted to be

$$\begin{aligned} &\mathbb{P}_{P_{t_0}}[Y_{i(t_0+1)} = y_0, \boldsymbol{U}_{i(t_0+2:t_0+ML)} = \boldsymbol{\mu}, \boldsymbol{X}_{i(t_0+1:t_0+ML)} = \boldsymbol{x}] \\ &\quad = \mathbb{P}_{P_{t_0+1}}[Y_{i(t_0+1)} = y_0, \boldsymbol{U}_{i(t_0+2:t_0+ML)} = \boldsymbol{\mu}, \boldsymbol{X}_{i(t_0+1:t_0+ML)} = \boldsymbol{x}] \quad \text{(Coherency)} \\ &\quad \text{for every } y_0 \in \mathcal{J}, \boldsymbol{\mu} \in \mathcal{J}^{ML-1}, \boldsymbol{x} \in \mathcal{X}_{ML-1}, \text{ and } t_0 \in \{0, ..., T-ML\} \end{aligned}$$

The incorporation of marketing variables also facilitates the MIV restriction:

$$\mathbb{P}_P[U_{it}(y) = j | X_{ijt} = x_1, \boldsymbol{X}_{i,-j,t} = \boldsymbol{x}_0] \geq \mathbb{P}_P[U_{it}(y) = j | X_{ijt} = x'_1, \boldsymbol{X}_{i,-j,t} = \boldsymbol{x}_0] \quad \text{(MIV)}$$

for each $y \in \mathcal{J}$, each $\boldsymbol{x}_0$, every $x_1 \leq x'_1$, every $j \in \mathcal{J}$ and every $t \geq 1$.

In Section 5.3, we implement three versions of the model: one that focuses on pricing by conditioning on a price vector including "own price" and "price of largest-share competitor;" another that focuses on promotions by conditioning on an in-store display vector including either "own display" and "display of largest-share competitor", or "own display" and "whether at least one competitor is on display." In particular, the version where we control for "own display“ and "whether at least one competitor is on display" is used to test search costs as an alternative mechanism for the HBL evidence. To keep the dimensionality manageable (4 support points for the marketing covariate in a single period), we discretize own and competitor prices by using a median-split to define them as "high" versus "low."

B.2.4 The Criterion Function

Let $\boldsymbol{W}_i \equiv (\boldsymbol{Y}_i, \boldsymbol{X}_i)$ denote the observable data, the joint support of which is $\mathcal{W} \equiv supp(\boldsymbol{Y}_i, \boldsymbol{X}_i)$. For each observable data $\boldsymbol{w} \equiv (\boldsymbol{w}_y, \boldsymbol{w}_x) \in \mathcal{W}$, define the *observational equivalence* moment of $\boldsymbol{w}$ as

$$m_{oeq,\boldsymbol{w}} \equiv \mathbb{1}[\boldsymbol{Y}_i = \boldsymbol{w}_y, \boldsymbol{X}_i = \boldsymbol{w}_x] - \sum_{\boldsymbol{\mu} \in \mathcal{U}_{oeq}(\boldsymbol{w}_y)} P(\boldsymbol{\mu}, \boldsymbol{w}_x)$$

Next, we divide the restriction function ρ into a deterministic component $\rho_d : \mathcal{P} \to \mathbb{R}^{d_\rho - d_s}$ and a stochastic component $\rho_s : \mathcal{P} \to \mathbb{R}^{d_s}$. The division is determined by whether the identifying restriction depend on the distribution of data $\boldsymbol{W}_i$. Since the stochastic component ρ_s depends on data, it is assumed to be representable as a moment condition, i.e., assuming there exists a function $m_\rho : \mathcal{W} \times \mathcal{P} \to \mathbb{R}^{d_s}$ for which $\rho_s(P) = \mathbb{E} m_\rho(\boldsymbol{W}_i, P)$.

For example, our Assumption MTR and the Coherency restriction are deterministic restrictions, i.e., they are a part of ρ_d, because they rules out certain potential outcomes paths without depending on the observed distribution of the data. Assumption MTS is a stochastic restriction, because the conditioning depends on the distribution of $Y_{i(t-1)}, Y_{i(t-2)}$. For assumption ST, it is a part of ρ_d when there is no marketing covariates, and it is a part of ρ_s after we include marketing covariates in (B.2.3) as it now depends on the distribution of $\boldsymbol{X}_{i(t-m:t)}$. Similarly, assumptions IV and MIV are part of ρ_s because it depends on the observed distribution of marketing covariates.

With such division, we can use the deterministic component ρ_d to first restrict the set of probability mass functions $\mathcal{P}$, and then rewrite the identified set $\mathcal{P}^\star$ as

$$\begin{aligned} \mathcal{P}^\star = &\{P \in \mathcal{P}_d^\dagger : \mathbb{E} m_{oeq,w}(\boldsymbol{W}_i, P) = 0 \quad \forall \boldsymbol{w} \in \mathcal{W} \\ &\text{and } \mathbb{E} m_{\rho,s}(\boldsymbol{W}_i, P) \geq 0 \quad \forall s = 1, ..., d_s\} \end{aligned}$$

where $\mathcal{P}_d^\dagger \equiv \{P \in \mathcal{P} : \rho_d(P) \geq 0\}$, and $m_{\rho,s}(\boldsymbol{W}_i, P)$ is the sth component of $m_\rho(\boldsymbol{W}_i, P)$.

In other words, we restate the DPO model as a moment inequality model with parameter space $\mathcal{P}_d^\dagger$ (under deterministic restrictions), moment equalities $\{\mathbb{E}m_{oeq,\boldsymbol{w}}(\boldsymbol{W}_i, P) = 0\}_{\boldsymbol{w}\in\mathcal{W}}$ (under observational equivalence restrictions), and moment inequalities $\{\mathbb{E}m_{\rho,s}(\boldsymbol{W}_i, P) \geq 0\}_{s=1}^{d_s}$ (under stochastic restrictions).

We construct a *population* criterion function as

$$Q(P,\eta) = \sum_{s=1}^{d_s} |\mathbb{E}m_{\rho,s}(\boldsymbol{W}_i, P) - \eta_s| + \sum_{\boldsymbol{w}\in\mathcal{W}} |\mathbb{E}m_{oeq,\boldsymbol{w}}(\boldsymbol{W}_i, P)|$$

where the first term corresponds to the moment inequalities under stochastic restrictions, and the second term corresponds to moment equalities under observational equivalence. We use a vector of non-negative slackness variables $\eta \in \mathbb{R}_+^{d_s}$ to transform the moment inequalities to moment equalities.

We can then construct a *sample analogue* of the criterion function Q,

$$Q_n(P,\eta) = \sum_{s=1}^{d_s} \sqrt{n} \left| \frac{1}{n}\sum_{i=1}^{n} m_{\rho,s}(\boldsymbol{W}_i, P) - \eta_s \right| + \sum_{\boldsymbol{w}\in\mathcal{W}} \sqrt{n} \left| \frac{1}{n}\sum_{i=1}^{n} m_{oeq,\boldsymbol{w}}(\boldsymbol{W}_i, P) \right|$$

B.2.5 Estimation

The estimation of the $\mathcal{HBL}$ parameter is done in two steps. In the first step, we find the set of *probability mass functions* P, a high dimensional primitive, that are close to minimizing the sample criterion function. We define the minimal sample criterion

$$\bar{Q}_n \equiv \min_{(P,\eta)\in\mathcal{P}_d^\dagger\times\mathbb{R}_+^{d_s}} Q_n(P,\eta)$$

and the set of P that are close enough, by a factor of τ_n, to minimizing the sample criterion function

$$\mathcal{P}_n \equiv \left\{ P \in \mathcal{P}_d^\dagger : Q_n(P,\eta) \leq \bar{Q}_n(1+\tau_n) \quad \text{for some } \eta \in \mathbb{R}_+^{d_s} \right\}$$

We choose $\tau_n = 0.8$ because it performed well in recovering the true $\mathcal{HBL}$ in synthetic choice panels generated by a DDC model estimated from the observed IRI data.

In the second step, we search among the probability mass functions in $\mathcal{P}_n$ to find the upper and lower bounds of the $\mathcal{HBL}$ parameter, a low dimensional target parameter. We define the estimator of upper and lower bounds of the $\mathcal{HBL}$ parameters as

$$\hat{\theta}^{\star}_{lb} \equiv \min \theta(\mathcal{P}_n) = \min_{(P,\eta)\in\mathcal{P}^{\dagger}_d\times\mathbb{R}^{d_s}_+} \theta(P) \quad s.t.\ Q_n(P,\eta) \le \bar{Q}_n(1+\tau_n) \quad \text{and}$$

$$\hat{\theta}^{\star}_{ub} \equiv \max \theta(\mathcal{P}_n) = \max_{(P,\eta)\in\mathcal{P}^{\dagger}_d\times\mathbb{R}^{d_s}_+} \theta(P) \quad s.t.\ Q_n(P,\eta) \le \bar{Q}_n(1+\tau_n)$$

We refer readers to Mogstad et al. (2018) for the consistency of the estimator.

B.2.6 Confidence Region

We construct confidence regions by test inversion, i.e., collecting all scalar values t, for which we fail to reject the null hypothesis $\mathbb{H}_0 : t \in \Theta^{\star}$ at some significance level α. We borrow the procedure for testing shape constraints proposed by Chernozhukov et al. (2015), which is also implemented in Torgovitsky (2019), to operationalize the test.

We first define a test statistic for the null hypothesis $\mathbb{H}_0 : t \in \Theta^{\star}$,

$$\bar{Q}_n(t) \equiv \inf_{(P,\eta)\in\mathcal{P}^{\dagger}_d(t)\times\mathbb{R}^{d_s}_+} Q_n(P,\eta), \tag{B.2}$$

where $\mathcal{P}^{\dagger}_d(t) \equiv \{P \in \mathcal{P}^{\dagger}_d : \theta(P) = t\}$. The distribution of $\bar{Q}_n(t)$ is nonstandard, and is approximated by the distribution of $\tilde{Q}_n(t)$,

$$\begin{aligned}
\tilde{Q}_n(t) \equiv \min_{(P,\eta,g,h)} \sum_{s=1}^{d_s} & \left| \xi^{\star}_{n,s}(P) + \frac{1}{n}\sum_{i=1}^{n} \nabla m_{\rho,s}(\boldsymbol{W}_i, P)[g] - h_s \right| \\
& + \sum_{\boldsymbol{w}\in\mathcal{W}} \left| \xi^{\star}_{n,\boldsymbol{w}}(P) + \frac{1}{n}\sum_{i=1}^{n} \nabla m_{oeq,\boldsymbol{w}}(\boldsymbol{W}_i, P)[g] \right| \\
s.t. \qquad & (P,\eta) \in \mathcal{P}^{\dagger}_d(t) \times \mathbb{R}^{d_s}_+, \\
& (P,\eta) + \frac{1}{\sqrt{n}}(g,h) \in \mathcal{P}^{\dagger}_d(t) \times \mathbb{R}^{d_s}_+, \text{ and} \\
& Q_n(P,\eta) \le \bar{Q}_n(t)(1+\tau_n),
\end{aligned}$$

where

$$\xi_{n,s} \equiv \frac{1}{\sqrt{n}}\sum_{i=1}^{n} \left[m_{\rho,s}(\boldsymbol{W}^{\star}_i, P) - \bar{m}_{\rho,s}(P) \right], \quad \xi_{n,\boldsymbol{w}} \equiv \frac{1}{\sqrt{n}}\sum_{i=1}^{n} \left[m_{oeq,\boldsymbol{w}}(\boldsymbol{W}^{\star}_i, P) - \bar{m}_{oeq,\boldsymbol{w}}(P) \right]$$

is derived from a bootstrapped sample $\{\boldsymbol{W}_i^\star\}_{i=1}^n$ drawn i.i.d. with replacement from the observed sample $\{\boldsymbol{W}_i\}_{i=1}^n$. $\nabla m_{\rho,s}(\boldsymbol{W}_i, P)[g]$ and $\nabla m_{oeq,\boldsymbol{w}}(\boldsymbol{W}_i, P)[g]$ denote the directional derivative of $m_{\rho,s}$ and $m_{oeq,\boldsymbol{w}}$ with respect to P, evaluated at P, in the direction g. The first constraint requires that (P, η) is in the parameter space under the null hypothesis. The second constraint requires that (g, h) are small deviations of (P, η) that remain in the parameter space under the null hypothesis. The third constraint requires that (P, η) approximately solve the test statistic (B.2).

We approximate the distribution of $\bar{Q}_n(t)$ by the distribution of $\tilde{Q}_n(t)$, which is further approximated by redrawing bootstrap samples $\{\boldsymbol{W}_i^\star\}_{i=1}^n$ a large number of times and computing $\tilde{Q}_n(t)$ for each draw. If $\bar{Q}_n(t)$ is greater than the $1-\alpha$ quantile of the bootstrapped distribution of $\tilde{Q}_n(t)$, we reject the null that t is in the identified set. A $1-\alpha$ confidence region of $\Theta^\star$ is the set of all t at which we do not reject the null test.

We implemented the test procedure with 30 bootstrapped samples and $\tau_n = 0.8$. We pick values of t by triangulation. For example, at the lower bound, we start from testing at $t_1 = (0 + \hat{\theta}_{lb}^\star)/2$. If we reject the null hypothesis at t_1, i.e., t_1 is not in the identified set, we proceed to the next testing point $t_2 = (t_1 + \hat{\theta}_{lb}^\star)/2$. If we then fail to reject t_2, we move to the third testing point $t_3 = (t_1 + t_2)/2$. The triangulation stops when the distance between left and right end of the triangulated interval, e.g., $t_2 - t_1$, is less than some tolerance threshold. The upper bound of the confidence region is constructed in a similar manner.

B.2.7 Mis-specification Tests

We use the same test statistic $\bar{Q}_n$ as in equation (B.2), except that we replace $\mathcal{P}_d^\dagger(t)$ with $\mathcal{P}_d^\dagger$. We reject the null hypothesis that $\mathbb{H}_0 : \mathcal{P}^\star = \emptyset$ if $\bar{Q}_n$ is greater than the $1-\alpha$ quantile of the simulated distribution of $\tilde{Q}_n$.

B.3 MPEC Estimator for DDC

Let $i = 1, ..., N$ index the observed trips. From the empirical DDC model in section 5.4, a consumer's choice-specific values are:

$$v_j\left(\mathbf{s}; \theta, \beta\right) = \gamma_j - \alpha p_j + \lambda \mathbb{I}_{\{\ell = j\}} + \beta \pi \mathbf{F}_j(\mathbf{s}) \mathbf{B}_j\left(\beta\right) \ln \left\{ \sum_{k \in \mathcal{D}} \exp\left(\mathbf{v}_k\right) \right\}$$

where $\theta = (\gamma_1, ..., \gamma_J, \alpha, \lambda)$ is the vector of taste parameters, $\beta \in [0, 1)$ is the discount factor, $\mathbf{F}_j$ is the $L \times L$ state transition matrix conditional on current choice j and $\mathbf{B}_j(\beta) = [\mathbf{I} - \beta(1-\pi)\mathbf{F}_j]^{-1}$ is an $(L \times L)$ matrix. Let $\Gamma = (\beta, \theta)$ denote all the structural parameters.

Let $\mathbf{v} = (v_{j1}, v_{j2}, ..., v_{J+1,L})$ denote the choice-specific values in each of the L states. We can now define an MPEC estimator:

$$\Theta^* = \underset{\Theta=(\Gamma,\mathbf{v})}{\arg\max} \sum_{i=1}^{N} \sum_{j=0}^{J} \left(Y_{ij} \ln\left(\sigma_j\left(\mathbf{p}_i, \ell_i; \theta, \mathbf{v}\right)\right)\right)$$

subject to the constraints

$$G_{js}(\Theta) = v_{js} - u_j(\mathbf{s}; \Gamma) - \beta\pi\mathbf{F}_j(\mathbf{s})\mathbf{B}_j(\beta) \ln\left\{\sum_{k\in\mathcal{D}} \exp\left(\mathbf{v}_k\right)\right\} \text{ , } \forall j \in \mathcal{D}, \forall s \in \mathcal{S}$$

where

$$\sigma_j(\mathbf{p}_i, \ell_i; \theta, \mathbf{v}) = \frac{\exp\left(u_j\left(\mathbf{p}_i, \ell_i; \theta\right) + v_{js_i} - u_j\left(\mathbf{s}_i; \theta\right)\right)}{\sum_{k\in\mathcal{D}} \exp\left(u_k\left(\mathbf{p}_i, \ell_i; \theta\right) + v_{ks_i} - u_k\left(\mathbf{s}_i; \theta\right)\right)}$$

and

$$u_j(\mathbf{p}, \ell; \theta) = \gamma_j - \alpha p_j + \lambda \mathbb{I}_{\{\ell=j\}}.$$

We use automatic differentiation to compute the gradient of the objective function.[23] To improve speed, we derive the constraint Jacobian in closed form:

$$\begin{aligned}
\frac{\partial G_{js}(\Theta)}{\partial \beta} &= -\pi\mathbf{F}_j(\mathbf{s})\mathbf{B}_j(\beta) \ln\left\{\sum_{k\in\mathcal{D}} \exp\left(\mathbf{v}_k\right)\right\} - \beta\pi\mathbf{F}_j(\mathbf{s})\mathbf{B}_j(\beta)(1-\pi)\mathbf{F}_j(\mathbf{s})\mathbf{B}_j(\beta) \ln\left\{\sum_{k\in\mathcal{D}} \exp\left(\mathbf{v}_k\right)\right\} \\
\frac{\partial G_{js}(\Theta)}{\partial \theta_k} &= -x_{sjk} \\
\frac{\partial G_{js}(\Theta)}{\partial v_{js}} &= 1 - \beta\pi\mathbf{F}_{js}(\mathbf{s})\mathbf{B}_j(\beta)\sigma_{js}, \forall j \in \mathcal{D}, \forall \mathbf{s} \in \mathcal{S} \\
\frac{\partial G_{js}(\Theta)}{\partial v_{ns'}} &= -\beta\pi\mathbf{F}_{js'}(\mathbf{s})\mathbf{B}_j(\beta)\sigma_{ns'}, \forall j \in \mathcal{D}, \forall n \in \mathcal{D}_{-j}, \forall \mathbf{s}, \mathbf{s}' \in \mathcal{S}
\end{aligned}$$

where $\sigma_{js} \equiv \sigma_j(\mathbf{s}; \theta, \mathbf{v})$ and $\nabla_\beta \mathbf{B}_j(\beta) = \mathbf{B}_j(\beta)(1-\pi)\boldsymbol{F}_j(\mathbf{s})\mathbf{B}_j(\beta)$.

B.4 Group Fixed-Effects and Unobserved Heterogeneity

We use Bonhomme et al. (2022)'s *grouped fixed effects* (GFE) approach to control for persistent, unobserved heterogeneity in consumer tastes. We assume (1) unobserved heterogeneity has a low underlying dimension, and (2) researchers have individual-specific moments that

[23] We use the Deep Learning Toolbox in Matlab to obtain exact derivatives.

are injective functions of the heterogeneity.

Let consumers be indexed by $h = 1, ..., H$. We use a vector of individual-specific moments w_h for the classification step. The moments consist of a household's income, propensity to purchase each of the available brands conditional on purchase, total number of observed purchases and total number of runs (i.e., repeat-purchase strings for the same brand). The choice of moments is important, as the injectivity assumption requires that one can separate the types of individuals just by using their moments w_h and $w_{h'}$: individuals with the same population moments have the same "type." From Berry et al. (2013), we know that our discrete-choice model with absolutely-continuous random utility shocks is invertible. So our propensity moments will be injective. However, it is less clear whether our runs moments will also satisfy injectivity.

We use a variable-neighborhood search for the k-means clustering algorithm to classify each consumer into groups. We use the variable-neighborhood search heuristic to overcome the local optimum from the use of Lloyd's algorithm.[24] Given our moderate-sized cross-section, we use $K = 6$ groups. Standard errors are calculated by block-bootstrapping consumers and re-running the variable-neighborhood-kmeans clustering. Comparing the GFE approach to standard random coefficients models using the same IRI data generate strikingly similar results, which are available upon request.

B.5 Computing the $\mathcal{HBL}$ statistic under the DDC

To simulate choices in each observed state, we first use the DDC estimates to compute the predicted choice-specific values at each observed state, $\mathbf{s}_{it}$ where $i = 1, ..., N$ and $t = 1, ..., T$: $\{\hat{v}_{ilt}(\mathbf{s}_{it})\}_{l \in \mathcal{D}}$. We then simulate random utilities ϵ_{ilt} for each choice alternative $l \in \mathcal{J}$, consumer i and time period t. We then approximate $\mathcal{HBL}_j$ as follows:

$$\begin{aligned} \mathcal{HBL}_j &= \mathbb{P}_{\boldsymbol{\epsilon}}[\exists k \neq j \in \mathcal{J} \text{ s.t. } j = \underset{l \in \mathcal{J}}{\text{argmax}}\{v_{il}(\mathbf{p}_t, j) + \epsilon_{ilt}\}, \text{ and } j \neq \underset{l \in \mathcal{J}}{\text{argmax}}\{v_{il}(\mathbf{p}_t, k) + \epsilon_{ilt}\}, \forall i, t] \\ &\approx \frac{1}{NTR} \sum_{i=1}^{N} \sum_{t=1}^{T} \sum_{r=1}^{R} \mathbb{1}\left[j = \arg\max_{l \in \mathcal{J}}\{v_{il}(\mathbf{p}_t, j) + \epsilon^r_{ilt}\}, \text{ and } j \neq \arg\max_{l \in \mathcal{J}}\{v_{il}(\mathbf{p}_t, k) + \epsilon^r_{ilt}\}, \forall i, t \right] \end{aligned} \tag{B.3}$$

We set $R = 10,000$.

[24] We use Matlab's *kmeans* function to initialize the search.

B.6 Computing the Long-Run Elasticity of Demand

The factual expected choice probabilities in equilibrium are $\mathbf{\Pi}\sigma\left(\mathbf{p},\ell\right)$, where $\sigma\left(\mathbf{p},\ell\right)$ is the $(L\times J+1)$ matrix of choice probabilities in each of the L states and $\mathbf{\Pi}$ is the $(1\times L)$ row-vector of equilibrium probabilities of being in each state. We then create a counterfactual set of L states which are identical to $(\mathbf{p},\ell)$ except for the modification for product k: $\tilde{p}_{kr}=1.1p_{kr}$. The corresponding counter-factual expected choice probabilities in equilibrium are $\tilde{\mathbf{\Pi}}\sigma\left(\tilde{\mathbf{p}},\ell\right)$ where $\tilde{\mathbf{\Pi}}$ is the vector of counter-factual equilibrium probabilities of being in each state. The long-run elasticity of demand is:

$$\varepsilon_j^{LR}=\frac{\frac{\sigma(\tilde{\mathbf{p}},\ell)\tilde{\mathbf{\Pi}}-\sigma(\mathbf{p},\ell)\mathbf{\Pi}}{\sigma(\mathbf{p},\ell)\mathbf{\Pi}}}{0.1}. \tag{B.4}$$

We compute the steady-state probabilities as follows. Let $\mathbf{Q}$ be the $(R\times J)$ transition matrix for prices. We compute the $(1\times J)$ row-vector of steady-state price probabilities, $\mathbf{q}=\mathbf{q}\mathbf{Q}$, by taking the left eigenvector of $\mathbf{Q}$ corresponding to the eigenvalue 1. We then use forward-simulation to find the steady-state probabilities Π corresponding to the combined state, $(\mathbf{p},\ell)$,with transition matrix, $\mathbf{F}$. We repeat the following steps $r=1,...,R$ times:

1. Initialize prices by assuming they are in equilibrium and drawing $\mathbf{p}_0^r\in\mathcal{S}$ using the probabilities $\mathbf{q}$. Then simulate a sequence of prices of length T, $\mathbf{p}^r\equiv(\mathbf{p}_1^r,...,\mathbf{p}_T^r)$, using the transition process $\mathbf{Q}$.

2. Initialize the loyalty state by drawing $\ell_0\in\mathcal{J}$ from uniform discrete distribution. Then use the sequence of prices, $\mathbf{p}^r$, and the DDC to simulate a choice sequence of length T and the corresponding sequence of loyalty states, $\boldsymbol{\ell}^r=(\ell_1^r,...,\ell_T^r)$.

We can then approximate the steady-state probabilities, $\mathbf{\Pi}$, as follows:

$$\mathbf{\Pi}_l=\frac{1}{R}\sum_{r=1}^{R}\mathbb{1}\left[\mathbf{s}_T^r=\mathbf{s}_l\right],\ \forall l\in\mathcal{S}. \tag{B.5}$$

We use $R=3000$ and $T=200$. Results for the CPG categories are displayed in tables B14 to B21.

B.7 Additional Tables

Category	Brand	(1) No assumptions	(2) ST(1)+MTR	(3) ST(1)+MTS+MTR
Coffee	Maxwell_S	[0.000,0.735]	[0.137,0.676]	[0.137,0.233]
		[0.000,0.747]	[0.045,0.743]	[0.039,0.308]
Deodorant	Mennen	[0.000,0.782]	[0.107,0.736]	[0.107,0.307]
		[0.000,0.792]	[0.000,0.784]	[0.000,0.358]
Mayonnaise	Hellmanns	[0.000,0.864]	[0.065,0.842]	[0.065,0.386]
		[0.000,0.871]	[0.033,0.865]	[0.033,0.429]
Margarine	Smart Balance	[0.000,0.910]	[0.045,0.895]	[0.045,0.485]
		[0.000,0.920]	[0.000,0.917]	[0.000,0.560]
Peanut Butter	Skippy	[0.000,0.749]	[0.081,0.725]	[0.081,0.261]
		[0.000,0.757]	[0.064,0.745]	[0.064,0.288]
Spaghetti Sauce	Ragu	[0.000,0.761]	[0.034,0.743]	[0.034,0.324]
		[0.000,0.771]	[0.000,0.759]	[0.000,0.362]
Toothbrush	Colgate	[0.000,0.791]	[0.162,0.734]	[0.162,0.242]
		[0.000,0.814]	[0.000,0.810]	[0.000,0.382]
Yogurt	Yoplait	[0.000,0.798]	[0.048,0.780]	[0.048,0.373]
		[0.000,0.806]	[0.000,0.803]	[0.000,0.411]

(a) $\mathcal{HBL}$ - Binomial DPO

Category	Brand	(1) No assumptions	(2) ST(1)+MTR	(3) ST(1)+MTS+MTR
Coffee	Maxwell_S	[0.000,0.866]	[0.070,0.817]	[0.070,0.395]
		[0.000,0.870]	[0.029,0.849]	[0.029,0.489]
Deodorant	Mennen	[0.000,0.888]	[0.071,0.848]	[0.071,0.513]
		[0.000,0.893]	[0.000,0.888]	[0.000,0.617]
Mayonnaise	Hellmanns	[0.000,0.933]	[0.043,0.919]	[0.042,0.579]
		[0.000,0.936]	[0.022,0.935]	[0.023,0.671]
Margarine	Smart Balance	[0.000,0.957]	[0.037,0.929]	[0.037,0.574]
		[0.000,0.961]	[0.000,0.947]	[0.000,0.706]
Peanut Butter	Skippy	[0.000,0.877]	[0.056,0.858]	[0.056,0.485]
		[0.000,0.879]	[0.046,0.872]	[0.046,0.535]
Spaghetti Sauce	Ragu	[0.000,0.881]	[0.029,0.858]	[0.029,0.493]
		[0.000,0.886]	[0.000,0.876]	[0.000,0.536]
Toothbrush	Colgate	[0.000,0.791]	[0.162,0.734]	[0.162,0.242]
		[0.000,0.816]	[0.052,0.809]	[0.053,0.423]
Yogurt	Yoplait	[0.000,0.903]	[0.038,0.876]	[0.038,0.526]
		[0.000,0.906]	[0.000,0.895]	[0.000,0.603]

(b) $\mathcal{HBL}$ - Multinomial DPO

This table summarizes the bounds on $\mathcal{HBL}$ using the estimation samples drawn from the IRI academic data. We report bounds for the cross-time average value of $\mathcal{HBL}_j$ where j is the top-selling brand in each category. Estimated bounds and 95% confidence regions are reported in large and small fonts, respectively. Estimation of panel (a) is based on the unconditional binomial DPO model with $J = 2$ and $T = 5$. Estimation on panel (b) is based on the unconditional trinomial DPO model with $J = 3$ (except for the toothbrush category where $J = 2$) and $T = 5$ using the dimension-reduction technique in Appendix B.2.2. In the binomial DPO model, the choice sets consist of the top-selling brand and a composite brand representing all other choice alternatives, while in the trinomial DPO model, the choice sets consist of each of the top-two selling brands and a composite brand representing all other choice alternatives. The 95% confidence regions are constructed using the method described in Appendix B.2.6, implemented using 30 bootstrap draws and a tuning parameter of $\tau_n = 0.8$.

Table B1: Habitual Brand Loyalty in IRI Data

Category	Brand	(1) No assumptions	(2) IV	(3) ST(1)	(4) ST(1)+MTS	(5) ST(1)+MTS+IV
Coffee	Maxwell_S	[0.000,0.735]	[0.119,0.530]	[0.094,0.678]	[0.057,0.444]	[0.132,0.489]
		[0.000,0.745]	[0.084,0.567]	[0.000,0.737]	[0.000,0.456]	[0.088,0.489]
Deodorant	Mennen	[0.000,0.782]	[0.192,0.625]	[0.144,0.742]	[0.113,0.531]	[0.179,0.542]
		[0.000,0.794]	[0.099,0.653]	[0.098,0.790]	[0.109,0.531]	[0.143,0.542]
Mayonnaise	Hellmanns	[0.000,0.864]	[0.150,0.758]	[0.091,0.827]	[0.076,0.458]	[0.188,0.627]
		[0.000,0.871]	[0.112,0.788]	[0.002,0.868]	[0.002,0.511]	[0.187,0.627]
Margarine	Smart Balance	[0.000,0.910]	[0.310,0.872]	[0.125,0.887]	[0.091,0.540]	[0.251,0.716]
		[0.000,0.920]	[0.103,0.917]	[0.000,0.916]	[0.000,0.604]	[0.213,0.717]
Peanut Butter	Skippy	[0.000,0.749]	[0.098,0.650]	[0.044,0.717]	[0.035,0.414]	[0.209,0.414]
		[0.000,0.755]	[0.023,0.701]	[0.002,0.749]	[0.002,0.472]	[0.161,0.711]
Spaghetti Sauce	Ragu	[0.000,0.761]	[0.097,0.664]	[0.139,0.669]	[0.126,0.336]	[0.182,0.325]
		[0.000,0.771]	[0.040,0.707]	[0.026,0.766]	[0.025,0.439]	[0.068,0.409]
Toothbrush	Colgate	[0.000,0.791]	[0.199,0.672]	[0.073,0.764]	[0.036,0.743]	[0.108,0.722]
		[0.000,0.814]	[0.193,0.690]	[0.073,0.801]	[0.036,0.797]	[0.108,0.722]
Yogurt	Yoplait	[0.000,0.798]	[0.165,0.688]	[0.148,0.738]	[0.116,0.423]	[0.224,0.518]
		[0.000,0.806]	[0.076,0.737]	[0.042,0.800]	[0.043,0.478]	[0.118,0.544]

(a) $\mathcal{HBL}$

Category	Brand	(1) No assumptions	(2) IV	(3) ST(1)	(4) ST(1)+MTS	(5) ST(1)+MTS+IV
Coffee	Maxwell_S	[0.000,1.000]	[0.155,0.877]	[0.179,0.925]	[0.103,0.838]	[0.122,0.883]
		[0.000,1.000]	[0.077,1.000]	[0.000,1.000]	[0.000,1.000]	[0.118,0.884]
Deodorant	Mennen	[0.000,1.000]	[0.175,0.927]	[0.215,0.951]	[0.157,0.942]	[0.176,0.934]
		[0.000,1.000]	[0.046,1.000]	[0.055,1.000]	[0.054,1.000]	[0.101,0.995]
Mayonnaise	Hellmanns	[0.000,1.000]	[0.098,0.885]	[0.055,0.983]	[0.046,0.514]	[0.105,0.729]
		[0.000,1.000]	[0.070,0.935]	[0.001,1.000]	[0.002,0.548]	[0.105,0.730]
Margarine	Smart Balance	[0.000,1.000]	[0.310,0.968]	[0.196,0.980]	[0.124,0.656]	[0.259,0.907]
		[0.000,1.000]	[0.092,1.000]	[0.000,1.000]	[0.000,0.719]	[0.214,1.000]
Peanut Butter	Skippy	[0.000,1.000]	[0.134,0.892]	[0.063,0.991]	[0.050,0.770]	[0.183,0.704]
		[0.000,1.000]	[0.018,1.000]	[0.003,1.000]	[0.004,0.838]	[0.040,0.918]
Spaghetti Sauce	Ragu	[0.000,1.000]	[0.098,0.937]	[0.171,0.917]	[0.162,0.576]	[0.188,0.538]
		[0.000,1.000]	[0.025,1.000]	[0.036,1.000]	[0.036,0.796]	[0.050,0.992]
Toothbrush	Colgate	[0.000,1.000]	[0.222,0.883]	[0.059,0.989]	[0.030,0.974]	[0.105,0.958]
		[0.000,1.000]	[0.198,0.980]	[0.058,1.000]	[0.029,1.000]	[0.105,0.958]
Yogurt	Yoplait	[0.000,1.000]	[0.112,0.930]	[0.185,0.956]	[0.150,0.661]	[0.168,0.779]
		[0.000,1.000]	[0.044,1.000]	[0.049,1.000]	[0.051,0.763]	[0.080,0.990]

(b) $\mathcal{HBL}$ Among Current Repeat Buyers

This table summarizes the bounds for the cross-time average value of $\mathcal{HBL}_j$ in panel (a) and $\mathcal{HBL}_j|Y_t = j$ and $Y_{t-1} = j$ in panel (b), where j is the top-selling brand in each category, using the estimation samples drawn from the IRI academic data. Estimated bounds and 95% confidence regions are reported in large and small fonts, respectively. Estimation is based on a binomial DPO model with marketing covariates, where $J = 2$, $T = 5$ using the dimension-reduction technique in Appendix B.2.2. The choice sets consist of the top-selling brand and a composite brand representing all other choice alternatives, and marketing covariates include the prices of the focal brand and its largest competing brand. The IV assumption uses lagged price vector as the excluded instrument. The 95% confidence regions are constructed using the method described in Appendix B.2.6, implemented using 30 bootstrap draws and a tuning parameter of $\tau_n = 0.8$.

Table B2: Habitual Brand Loyalty in IRI Data - Binomial DPO

Category	Brand	(1) No assumptions	(2) IV	(3) ST(1)	(4) ST(1)+MTS	(5) ST(1)+MTS+IV
Coffee	Maxwell_S	[0.000,0.735]	[0.133,0.507]	[0.166,0.669]	[0.093,0.454]	[0.138,0.486]
		[0.000,0.744]	[0.098,0.509]	[0.033,0.736]	[0.033,0.484]	[0.137,0.486]
Deodorant	Mennen	[0.000,0.782]	[0.122,0.682]	[0.094,0.748]	[0.075,0.496]	[0.156,0.532]
		[0.000,0.794]	[0.026,0.743]	[0.000,0.789]	[0.000,0.561]	[0.135,0.533]
Mayonnaise	Hellmanns	[0.000,0.864]	[0.154,0.799]	[0.080,0.839]	[0.074,0.523]	[0.191,0.683]
		[0.000,0.871]	[0.154,0.800]	[0.028,0.867]	[0.028,0.524]	[0.191,0.683]
Margarine	Smart Balance	[0.000,0.910]	[0.306,0.847]	[0.042,0.887]	[0.037,0.550]	[0.348,0.560]
		[0.000,0.918]	[0.236,0.888]	[0.000,0.916]	[0.000,0.680]	[0.299,0.594]
Peanut Butter	Skippy	[0.000,0.749]	[0.075,0.614]	[0.036,0.711]	[0.028,0.404]	[0.157,0.373]
		[0.000,0.755]	[0.010,0.642]	[0.000,0.748]	[0.000,0.458]	[0.029,0.496]
Spaghetti Sauce	Ragu	[0.000,0.761]	[0.134,0.633]	[0.095,0.712]	[0.080,0.441]	[0.202,0.420]
		[0.000,0.770]	[0.064,0.664]	[0.013,0.762]	[0.013,0.499]	[0.086,0.477]
Toothbrush	Colgate	[0.000,0.791]	[0.188,0.739]	[0.107,0.741]	[0.072,0.492]	[0.101,0.770]
		[0.000,0.816]	[0.187,0.739]	[0.000,0.797]	[0.000,0.554]	[0.101,0.771]
Yogurt	Yoplait	[0.000,0.798]	[0.180,0.713]	[0.058,0.770]	[0.047,0.497]	[0.176,0.543]
		[0.000,0.806]	[0.028,0.756]	[0.000,0.805]	[0.000,0.630]	[0.108,0.552]

(a) $\mathcal{HBL}$

Category	Brand	(1) No assumptions	(2) IV	(3) ST(1)	(4) ST(1)+MTS	(5) ST(1)+MTS+IV
Coffee	Maxwell_S	[0.000,1.000]	[0.241,0.783]	[0.321,0.879]	[0.168,0.844]	[0.214,0.865]
		[0.000,1.000]	[0.132,0.899]	[0.069,1.000]	[0.085,1.000]	[0.214,0.865]
Deodorant	Mennen	[0.000,1.000]	[0.110,0.991]	[0.146,0.969]	[0.118,0.892]	[0.121,0.961]
		[0.000,1.000]	[0.009,1.000]	[0.000,1.000]	[0.000,1.000]	[0.043,1.000]
Mayonnaise	Hellmanns	[0.000,1.000]	[0.083,0.953]	[0.056,0.990]	[0.051,0.612]	[0.116,0.799]
		[0.000,1.000]	[0.083,0.996]	[0.020,1.000]	[0.021,0.612]	[0.115,0.800]
Margarine	Smart Balance	[0.000,1.000]	[0.087,0.866]	[0.043,0.980]	[0.040,0.634]	[0.200,0.638]
		[0.000,1.000]	[0.015,1.000]	[0.000,1.000]	[0.000,0.765]	[0.044,0.759]
Peanut Butter	Skippy	[0.000,1.000]	[0.107,0.940]	[0.057,0.987]	[0.046,0.768]	[0.200,0.773]
		[0.000,1.000]	[0.016,1.000]	[0.000,1.000]	[0.000,0.877]	[0.016,0.995]
Spaghetti Sauce	Ragu	[0.000,1.000]	[0.211,0.854]	[0.090,0.978]	[0.076,0.689]	[0.277,0.702]
		[0.000,1.000]	[0.110,0.933]	[0.014,1.000]	[0.014,0.733]	[0.090,0.873]
Toothbrush	Colgate	[0.000,1.000]	[0.216,0.993]	[0.079,0.986]	[0.055,0.590]	[0.055,1.000]
		[0.000,1.000]	[0.216,1.000]	[0.000,1.000]	[0.000,0.665]	[0.055,1.000]
Yogurt	Yoplait	[0.000,1.000]	[0.154,0.937]	[0.071,0.978]	[0.058,0.719]	[0.167,0.819]
		[0.000,1.000]	[0.006,1.000]	[0.000,1.000]	[0.000,0.868]	[0.093,0.849]

(b) $\mathcal{HBL}$ Among Current Repeat Buyers

This table summarizes the bounds for the cross-time average value of $\mathcal{HBL}_j$ in panel (a) and $\mathcal{HBL}_j|Y_t = j$ and $Y_{t-1} = j$ in panel (b), where j is the top-selling brand in each category, using the estimation samples drawn from the IRI academic data. Estimated bounds and 95% confidence regions are reported in large and small fonts, respectively. Estimation is based on a binomial DPO model with marketing covariates, where $J = 2$, $T = 5$ and we use the dimension-reduction technique in Appendix B.2.2. The choice sets consist of the top-selling brand and a composite brand representing all other choice alternatives, and marketing covariates include the in-store display of the focal brand and its largest competing brand. The IV assumption uses lagged display vector as the excluded instrument. The 95% confidence regions are constructed using the method described in Appendix B.2.6, implemented using 30 bootstrap draws and a tuning parameter of $\tau_n = 0.8$.

Table B3: Habitual Brand Loyalty - Binomial DPO controlling for in-store display

Category	Brand	IV		ST(1)		ST(1)+MTS		ST(1)+MTS+IV	
		$\Theta^\star = \emptyset$ in sample	p-value for $\mathbb{H}_0: \Theta^\star \neq \emptyset$	$\Theta^\star = \emptyset$ in sample	p-value for $\mathbb{H}_0: \Theta^\star \neq \emptyset$	$\Theta^\star = \emptyset$ in sample	p-value for $\mathbb{H}_0: \Theta^\star \neq \emptyset$	$\Theta^\star = \emptyset$ in sample	p-value for $\mathbb{H}_0: \Theta^\star \neq \emptyset$
Coffee	Maxwell$_S$	Yes	0.267	Yes	0.800	Yes	0.500	Yes	0.233
Deodorant	Mennen	Yes	0.500	Yes	0.500	Yes	0.333	Yes	0.467
Mayonnaise	Hellmanns	Yes	0.367	Yes	0.833	Yes	0.633	Yes	0.167
Margarine	Smart Balance	Yes	0.633	Yes	0.967	Yes	0.767	Yes	0.000
Peanut Butter	Skippy	Yes	0.767	Yes	0.700	Yes	0.567	Yes	0.833
Spaghetti Sauce	Ragu	No	1.000	No	1.000	Yes	1.000	Yes	1.000
Toothbrush	Colgate	Yes	0.467	Yes	0.033	Yes	0.000	Yes	0.000
Yogurt	Yoplait	Yes	1.000	Yes	1.000	Yes	0.567	Yes	0.500

(a) Binary DPO Controlling for Prices

Category	Brand	IV		ST(1)		ST(1)+MTS		ST(1)+MTS+IV	
		$\Theta^\star = \emptyset$ in sample	p-value for $\mathbb{H}_0: \Theta^\star \neq \emptyset$	$\Theta^\star = \emptyset$ in sample	p-value for $\mathbb{H}_0: \Theta^\star \neq \emptyset$	$\Theta^\star = \emptyset$ in sample	p-value for $\mathbb{H}_0: \Theta^\star \neq \emptyset$	$\Theta^\star = \emptyset$ in sample	p-value for $\mathbb{H}_0: \Theta^\star \neq \emptyset$
Coffee	Maxwell$_S$	Yes	0.100	Yes	1.000	Yes	0.233	Yes	0.000
Deodorant	Mennen	Yes	0.233	Yes	0.800	Yes	0.767	Yes	0.267
Mayonnaise	Hellmanns	Yes	0.033	Yes	0.100	Yes	0.167	Yes	0.000
Margarine	Smart Balance	Yes	0.100	Yes	0.667	Yes	0.600	Yes	0.700
Peanut Butter	Skippy	Yes	0.967	Yes	1.000	Yes	0.800	Yes	0.867
Spaghetti Sauce	Ragu	Yes	0.933	Yes	0.200	Yes	0.467	Yes	0.667
Toothbrush	Colgate	Yes	0.300	Yes	0.267	Yes	0.367	Yes	0.000
Yogurt	Yoplait	Yes	0.400	Yes	0.600	Yes	0.400	Yes	0.167

(b) Binary DPO Controlling for In-store Display

This table summarizes the misspecification test results for DPO models. For each brand-assumption pair, the first column reports whether the identified set is empty in the sample and the second column reports the p-value for rejecting the null hypothesis that the identified set is non-empty under a set of assumptions. Each household's last six purchases are reconstructed to be a focal brand and a composite brand for the binary DPO model, while observed marketing covariates include prices or in-store display of the focal brand as well as of its leading competitor. The DPO model combines shorter models to model longer sequences of choices, as described in Appendix B.2.2. The tuning parameter is set to be $\tau_n = 0.8$ and the number of bootstrap samples 30.

Table B4: Misspecification Tests

coefficient	(1) MNL	(2) s.e. (MNL)	(3) SDMNL	(4) s.e. (SDMNL)	(5) MNLGFE	(6) s.e. (MNLGFE)	(7) SDMNLGFE	(8) s.e. (SDMNLGFE)	(9) DDCGFE	(10) s.e. (DDCGFE)
beta	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.799	0.167
price	-1.011	0.041	-1.019	0.041	-1.079	0.035	-1.089	0.035	-1.084	0.035
loyalty	0.000	0.000	1.447	0.071	0.000	0.000	0.621	0.056	0.720	0.073
FOLGERS Sm	0.553	0.163	-0.141	0.179	0.403	0.156	0.174	0.159	0.191	0.161
FOLGERS Lg	-1.315	0.167	-1.692	0.169	-1.768	0.202	-1.840	0.200	-1.742	0.197
MAXWELL Sm	0.760	0.129	0.033	0.147	0.539	0.150	0.308	0.158	0.212	0.157
MAXWELL Lg	-1.855	0.168	-2.045	0.172	-2.177	0.182	-2.198	0.180	-2.099	0.178
PL Sm	-1.406	0.192	-1.650	0.190	-1.997	0.189	-2.040	0.184	-1.951	0.192

This table summarizes the myopic random utility model (denoted "MNL") and DDC coefficient estimates from the Coffee category in the IRI academic data. Specifications indicating "GFE" use the "group fixed-effects" approach to control for unobserved heterogeneity. Reported parameter values represent the population mean value of a parameter. Standard errors are computed via bootstrap.

Table B5: Maximum Likelihood Estimates for Coffee

coefficient	(1) MNL	(2) s.e. (MNL)	(3) SDMNL	(4) s.e. (SDMNL)	(5) MNLGFE	(6) s.e. (MNLGFE)	(7) SDMNLGFE	(8) s.e. (SDMNLGFE)	(9) DDCGFE	(10) s.e. (DDCGFE)
beta	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.981	0.077
price	-0.729	0.046	-0.716	0.046	-0.715	0.047	-0.718	0.047	-0.718	0.046
loyalty	0.000	0.000	1.780	0.061	0.000	0.000	0.541	0.053	0.730	0.087
DEGREE	-3.450	0.153	-3.924	0.162	-4.305	0.226	-4.384	0.218	-4.201	0.228
DOVE	-3.386	0.153	-4.013	0.155	-4.107	0.166	-4.204	0.167	-4.051	0.178
MENNEN	-2.579	0.150	-3.614	0.157	-3.112	0.157	-3.303	0.144	-3.480	0.146
OLDSPICE	-2.897	0.158	-3.591	0.175	-3.650	0.149	-3.754	0.151	-3.631	0.164
RIGHT	-3.557	0.164	-4.055	0.160	-4.358	0.296	-4.489	0.293	-4.329	0.282
SECRET	-3.212	0.184	-3.580	0.175	-4.139	0.183	-4.195	0.181	-3.977	0.191

This table summarizes the myopic random utility model (denoted "MNL") and DDC coefficient estimates from the Deodorants category in the IRI academic data. Specifications indicating "GFE" use the "group fixed-effects" approach to control for unobserved heterogeneity. Reported parameter values represent the population mean value of a parameter. Standard errors are computed via bootstrap.

Table B6: Maximum Likelihood Estimates for Deodorants

coefficient	(1) MNL	(2) s.e. (MNL)	(3) SDMNL	(4) s.e. (SDMNL)	(5) MNLGFE	(6) s.e. (MNLGFE)	(7) SDMNLGFE	(8) s.e. (SDMNLGFE)	(9) DDCGFE	(10) s.e. (DDCGFE)
beta	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.874	0.062
price	-0.753	0.056	-0.719	0.057	-0.850	0.055	-0.839	0.056	-0.840	0.055
loyalty	0.000	0.000	2.998	0.067	0.000	0.000	1.459	0.093	1.459	0.105
ICBINB	-2.205	0.148	-4.247	0.134	-3.186	0.151	-3.889	0.123	-3.742	0.114
SMARTBALANCE	-1.971	0.105	-4.163	0.117	-2.684	0.152	-3.464	0.124	-3.396	0.121
SHEDDS	-3.240	0.125	-4.813	0.119	-4.462	0.157	-4.878	0.126	-4.685	0.132
PROMISE	-3.104	0.162	-4.273	0.152	-4.551	0.167	-4.892	0.146	-4.684	0.161

This table summarizes the myopic random utility model (denoted "MNL") and DDC coefficient estimates from the Margarine category in the IRI academic data. Specifications indicating "GFE" use the "group fixed-effects" approach to control for unobserved heterogeneity. Reported parameter values represent the population mean value of a parameter. Standard errors are computed via bootstrap.

Table B7: Maximum Likelihood Estimates for Margarine

coefficient	(1) MNL	(2) s.e. (MNL)	(3) SDMNL	(4) s.e. (SDMNL)	(5) MNLGFE	(6) s.e. (MNLGFE)	(7) SDMNLGFE	(8) s.e. (SDMNLGFE)	(9) DDCGFE	(10) s.e. (DDCGFE)
beta	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.867	0.057
price	-2.382	0.054	-2.430	0.053	-2.408	0.046	-2.417	0.045	-2.414	0.046
loyalty	0.000	0.000	1.882	0.047	0.000	0.000	0.498	0.074	0.941	0.186
HELLMANNS	0.455	0.078	-1.110	0.086	0.258	0.074	-0.104	0.112	-0.526	0.188
PL	-2.445	0.088	-2.960	0.078	-3.730	0.104	-3.808	0.104	-3.312	0.152
KRAFT	-1.897	0.110	-2.269	0.107	-2.863	0.130	-2.915	0.128	-2.429	0.171

This table summarizes the myopic random utility model (denoted "MNL") and DDC coefficient estimates from the Mayonnaise category in the IRI academic data. Specifications indicating "GFE" use the "group fixed-effects" approach to control for unobserved heterogeneity. Reported parameter values represent the population mean value of a parameter. Standard errors are computed via bootstrap.

Table B8: Maximum Likelihood Estimates for Mayonnaise

coefficient	(1) MNL	(2) s.e. (MNL)	(3) SDMNL	(4) s.e. (SDMNL)	(5) MNLGFE	(6) s.e. (MNLGFE)	(7) SDMNLGFE	(8) s.e. (SDMNLGFE)	(9) DDCGFE	(10) s.e. (DDCGFE)
beta	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.845	0.113
price	-1.367	0.067	-1.248	0.064	-1.147	0.063	-1.168	0.062	-1.156	0.063
loyalty	0.000	0.000	1.541	0.024	0.000	0.000	0.691	0.020	0.692	0.038
JIF	-1.451	0.129	-2.425	0.128	-2.410	0.120	-2.633	0.116	-2.616	0.122
PETERPAN	-3.057	0.124	-3.538	0.119	-4.225	0.138	-4.276	0.134	-4.182	0.136
PL	-2.148	0.100	-2.974	0.101	-3.140	0.108	-3.304	0.108	-3.255	0.106
SKIPPY	-1.347	0.115	-2.426	0.120	-2.164	0.107	-2.429	0.108	-2.438	0.109

This table summarizes the myopic random utility model (denoted "MNL") and DDC coefficient estimates from the Peanut Butter category in the IRI academic data. Specifications indicating "GFE" use the "group fixed-effects" approach to control for unobserved heterogeneity. Reported parameter values represent the population mean value of a parameter. Standard errors are computed via bootstrap.

Table B9: Maximum Likelihood Estimates for Peanut Butter

coefficient	(1) MNL	(2) s.e. (MNL)	(3) SDMNL	(4) s.e. (SDMNL)	(5) MNLGFE	(6) s.e. (MNLGFE)	(7) SDMNLGFE	(8) s.e. (SDMNLGFE)	(9) DDCGFE	(10) s.e. (DDCGFE)
beta	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.888	0.077
price	-3.326	0.074	-3.386	0.072	-3.333	0.073	-3.357	0.073	-3.352	0.073
loyalty	0.000	0.000	1.644	0.045	0.000	0.000	0.722	0.041	0.722	0.064
FRANCESCORINALDI	-1.311	0.079	-1.915	0.090	-1.825	0.085	-1.972	0.089	-1.911	0.092
HUNTS	-2.473	0.087	-2.811	0.093	-3.217	0.093	-3.298	0.092	-3.197	0.097
PREGO	-0.026	0.103	-0.591	0.105	-0.701	0.109	-0.865	0.110	-0.791	0.108
RAGU	0.995	0.089	0.024	0.110	0.496	0.104	0.133	0.112	0.133	0.124

This table summarizes the myopic random utility model (denoted "MNL") and DDC coefficient estimates from the Spaghetti Sauce category in the IRI academic data. Specifications indicating "GFE" use the "group fixed-effects" approach to control for unobserved heterogeneity. Reported parameter values represent the population mean value of a parameter. Standard errors are computed via bootstrap.

Table B10: Maximum Likelihood Estimates for Spaghetti Sauce

coefficient	(1) MNL	(2) s.e. (MNL)	(3) SDMNL	(4) s.e. (SDMNL)	(5) MNLGFE	(6) s.e. (MNLGFE)	(7) SDMNLGFE	(8) s.e. (SDMNLGFE)	(9) DDCGFE	(10) s.e. (DDCGFE)
beta	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.671	0.238
price	-0.510	0.068	-0.470	0.072	-0.553	0.066	-0.552	0.065	-0.551	0.065
loyalty	0.000	0.000	1.187	0.089	0.000	0.000	-0.043	0.156	0.486	0.288
COLGATE	-3.093	0.188	-4.086	0.214	-3.721	0.212	-3.688	0.221	-4.212	0.284
ORALB	-3.060	0.222	-3.829	0.236	-4.142	0.292	-4.154	0.291	-3.950	0.341

This table summarizes the myopic random utility model (denoted "MNL") and DDC coefficient estimates from the Toothbrushes category in the IRI academic data. Specifications indicating "GFE" use the "group fixed-effects" approach to control for unobserved heterogeneity. Reported parameter values represent the population mean value of a parameter. Standard errors are computed via bootstrap.

Table B11: Maximum Likelihood Estimates for Toothbrushes

coefficient	(1) MNL	(2) s.e. (MNL)	(3) SDMNL	(4) s.e. (SDMNL)	(5) MNLGFE	(6) s.e. (MNLGFE)	(7) SDMNLGFE	(8) s.e. (SDMNLGFE)	(9) DDCGFE	(10) s.e. (DDCGFE)
beta	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.380	0.127
price	-1.134	0.049	-1.149	0.050	-1.213	0.050	-1.221	0.050	-1.221	0.050
loyalty	0.000	0.000	1.792	0.048	0.000	0.000	1.111	0.039	1.111	0.041
YOPLAIT	-0.806	0.096	-1.906	0.117	-1.326	0.104	-1.955	0.109	-1.963	0.114
DANNON	-1.558	0.087	-2.510	0.103	-1.891	0.107	-2.369	0.111	-2.356	0.113
COLOMBO	-2.136	0.099	-2.918	0.100	-2.638	0.107	-3.025	0.103	-3.009	0.104

This table summarizes the myopic random utility model (denoted "MNL") and DDC coefficient estimates from the Yogurt category in the IRI academic data. Specifications indicating "GFE" use the "group fixed-effects" approach to control for unobserved heterogeneity. Reported parameter values represent the population mean value of a parameter. Standard errors are computed via bootstrap.

Table B12: Maximum Likelihood Estimates for Yogurt

category	brand	(1) $\hat{\mathcal{HBL}}_{MNL}$	(2) $\hat{\mathcal{HBL}}_{MNL}$ (s.e.)	(3) $\hat{\mathcal{HBL}}_{DDC}$	(4) $\hat{\mathcal{HBL}}_{DDC}$ (s.e.)	(5) $\hat{\mathcal{HBL}}_{MNLheter}$	(6) $\hat{\mathcal{HBL}}_{MNLheter}$ (s.e.)	(7) $\hat{\mathcal{HBL}}_{DDCheter}$	(8) $\hat{\mathcal{HBL}}_{DDCheter}$ (s.e.)	(9) $\hat{\mathcal{HBL}}_{MNLhomogSD}$	(10) $\hat{\mathcal{HBL}}_{MNLhomogSD}$ (s.e.)
coffee	MAXWELL Sm	0.523	0.027	0.528	0.027	0.243	0.026	0.276	0.031	0.253	0.027
deodorant	MENNEN	0.602	0.016	0.630	0.016	0.171	0.022	0.222	0.029	0.198	0.040
mayonnaise	HELLMANNS	0.668	0.012	0.648	0.012	0.204	0.037	0.338	0.067	0.107	0.011
margarine	SMARTBALANCE	0.863	0.008	0.885	0.012	0.541	0.030	0.540	0.035	0.601	0.070
peanut butter	SKIPPY	0.584	0.007	0.595	0.007	0.275	0.016	0.275	0.018	0.286	0.139
spaghetti sauce	RAGU	0.613	0.014	0.619	0.014	0.265	0.019	0.265	0.026	0.240	0.028
toothbrushes	COLGATE	0.522	0.034	0.516	0.034	0.094	0.052	0.261	0.068	0.057	0.032
yogurt	YOPLAIT	0.687	0.013	0.688	0.013	0.478	0.022	0.478	0.046	0.454	0.024

This table summarizes the $\mathcal{HBL}$ statistics corresponding to the DDC estimates from the IRI academic data. $\mathcal{HBL}$ computed for the top brand in each of the categories. All specifications use the "group fixed-effects" approach to control for unobserved heterogeneity. Reported parameter values represent the population mean value of a parameter. Standard errors are computed via bootstrap.

Table B13: The Extent of Habitual Brand Loyalty: DDC evidence

brand	(1) SR Elasticity (RUM)	(2) LR Elasticity (DDC)	(3) % difference
$FOLGERS_S$	-3.298	-3.777	14.5
$FOLGERS_L$	0.068	0.076	11.3
$MAXWELL_S$	0.044	0.029	-32.5
$MAXWELL_L$	0.074	0.114	54.7
$PRIVATELABEL_S$	0.066	0.095	44.8
no purch	0.067	0.077	14.1

Table B14: LR elasticities: Coffee

brand	(1) SR Elasticity (RUM)	(2) LR Elasticity (DDC)	(3) % difference
DEGREE	-2.009	-2.467	22.8
DOVE	0.008	0.039	377.6
MENNEN	0.009	0.033	280.3
OLDSPICE	0.01	0.017	74
RIGHT	0.015	0.133	796
SECRET	0.008	-0.003	-134.1
no purch	0.016	0.02	26

Table B15: LR elasticities: Deodorant

brand	(1) SR Elasticity (RUM)	(2) LR Elasticity (DDC)	(3) % difference
HELLMANNS	-3.778	-3.774	-0.1
PRIVATELABEL	0.041	0.624	1440.1
KRAFTMAYO	0.035	0.659	1759.5
no purch	0.054	0.046	-13.5

Table B16: LR elasticities: Mayonnaise

brand	(1) SR Elasticity (RUM)	(2) LR Elasticity (DDC)	(3) % difference
ICANTBELIEVEITSNOTBUTTER	-1.418	-1.772	24.9
SMARTBALANCE	0.006	0.011	69.1
SHEDDSCOUNTRYCROCK	0.012	0.016	29.5
PROMISE	0.015	0.019	21.9
no purch	0.012	0.02	75.5

Table B17: LR elasticities: Margarine

brand	(1) SR Elasticity (RUM)	(2) LR Elasticity (DDC)	(3) % difference
JIF	-2.062	-2.406	16.7
PETERPAN	0.023	0.032	38.2
PRIVATELABEL	0.023	0.035	51.9
SKIPPY	0.028	-0.014	-148.5
no purch	0.025	0.032	28

Table B18: LR elasticities: Peanut Butter

brand	(1) SR Elasticity (RUM)	(2) LR Elasticity (DDC)	(3) % difference
FRANCESCORINALDI	-3.382	-3.733	10.4
HUNTS	0.061	0.065	6.6
PREGO	0.082	0.099	21.2
RAGU	0.062	0.05	-18.5
no purch	0.067	0.084	24.9

Table B19: LR elasticities: Spaghetti Sauce

brand	(1) SR Elasticity (RUM)	(2) LR Elasticity (DDC)	(3) % difference
COLGATE	-1.747	-1.74	-0.4
ORALB	0.009	0.131	1308.6
no purch	0.026	0.026	-2.2

Table B20: LR elasticities: Toothbrushes

brand	(1) SR Elasticity (RUM)	(2) LR Elasticity (DDC)	(3) % difference
YOPLAIT	-2	-2.006	0.3
DANNON	0.091	0.578	537.9
COLOMBO	0.081	0.422	420.1
no purch	0.086	0.07	-19.1

Table B21: LR elasticities: Yogurt