

Randomized Markdowns and Online Monitoring*

Ken Moon

The Wharton School, University of Pennsylvania, Philadelphia, PA 19104
kenmoon@wharton.upenn.edu

Kostas Bimpikis, Haim Mendelson

Graduate School of Business, Stanford University, Stanford, CA 94305
{kostasb, haim}@stanford.edu

Online retail reduces the costs of obtaining information about a product's price and availability and of flexibly timing a purchase. Consequently, consumers can strategically time their purchases, weighing the costs of monitoring and the risk of inventory depletion against prospectively lower prices. At the same time, firms can observe and exploit their customers' monitoring behavior. Using a dataset tracking customers of a North American specialty retail brand, we present empirical evidence that monitoring products online is associated with successfully obtaining discounts. We develop a structural model of consumers' dynamic monitoring to find substantial heterogeneity, with consumers' opportunity costs for an online visit ranging from \$2 to \$25 in inverse relation to their price elasticities. Our estimation results have important implications for retail operations. The randomized markdown policy benefits retailers by combining price commitment with the exploitation of heterogeneity in consumers' monitoring costs. We estimate that the retailer's profit under randomized markdowns is 81% higher than from subgame-perfect, state-contingent pricing, because the retailer need not limit its inventory to credibly limit markdowns, which permits its jointly optimal inventory stock to expand by 133%. The welfare gain from these larger inventories splits nearly equally into retailer profit and consumer surplus. We also discuss targeting customers with price promotions using their online histories and the implications of reducing consumers' monitoring costs.

Key words: Continuous-time stochastic game; Dynamic consumer behavior; Price commitment; Price discrimination; Revenue management; Structural estimation.

* We gratefully acknowledge the Wharton Customer Analytics Initiative (WCAI) and our anonymous corporate sponsor for the dataset used in this paper. We thank our two editors, three anonymous referees, Rob Bray, and Harikesh Nair for their helpful comments. Seminar participants at the WCAI Symposium, MSOM Service Operations

1. Introduction

Customers increasingly complement their shopping experiences by using retailers' online channels to search for and monitor products, prices, and availability. This behavior is of growing economic and operational significance, as U.S. e-commerce retail sales have sustained annual growth rates in excess of 14% for each quarter of 2015, surpassing USD 340B annually (U.S. Census Bureau, May 17, 2016). Online channels, however, may be a double-edged sword: while they enable consumers to monitor their products and prices of interest, firms may observe such monitoring and adjust their operational strategies accordingly. Then, in whose favor does the blade ultimately cleave?

In this paper, we use a novel dataset comprehensively tracking the customers of an established, North American specialty retailer. The customer-level dataset contains detailed information regarding not only customers' purchases but also the timing and browsing of their visits to the retailer's online channel: we thus observe an individual consumer's monitoring process during the time leading up to her potential purchase. Panel data results support a relationship between consumers' monitoring intensity (i.e., beyond that induced by the availability of discounts) and the discounts they obtain, after controlling for customer characteristics, promotions, and marketing communications at the individual customer and transaction level.

Motivated by this relationship between monitoring and the propensity for obtaining discounts, we present a structural equilibrium model under which the firm's markdown decisions and consumers' visits and purchases are the endogenous outcomes of a continuous-time, stochastic game with embedded discrete choices. Using customer-level data, our model captures how consumers respond to changes in their payoffs from monitoring the retailer's online channel. We establish the existence of a Markov stationary equilibrium and its key properties. Methodologically, our compact framework allows us to estimate a dynamic game with unobserved heterogeneity that accommodates forward-looking players' sequential and separately endogenous decisions to monitor SIG (Kamalini Ramdas as discussant), INFORMS RMP, Stanford Econometrics Lunch, Stanford GSB, UNC Kenan-Flager, MIT Sloan, Yale SOM, Duke Fuqua, UCLA Anderson, Chicago Booth, Wharton, INSEAD, and UCSD Rady provided valuable feedback.

and purchase (or more generally, to interact and make a discrete choice upon interacting). We derive a lattice structure over the game’s equilibria, which proves crucial for computation and for ruling out multiplicity of equilibria in practice.

We estimate our model directly from customer-level data to find substantial heterogeneity in monitoring costs and price elasticities across two customer segments. Approximately 18% of the customers are price-elastic bargain hunters, with an individual opportunity cost of \$2 for an online visit per month. The remaining high valuation customers incur an opportunity cost of up to over \$25 to make an additional online visit per month.

Given our estimation results, we study the problem of optimizing the retailer’s pricing and inventory policies. As supported by the data, our retailer executes a state-independent pricing policy that randomizes the timing of markdowns over a small sequence of relatively predictable markdown levels. We find that committing to such a markdown policy is highly valuable in this setting, yielding profits 81% higher than state-contingent, dynamic pricing, and that randomizing rather than pre-announcing or committing to deterministic markdown times can better monetize consumers with heterogeneous monitoring costs. This may justify the widespread practice of imposing internal “business rule” constraints on prices, which can be seen as effectively carrying out retailers’ pricing commitments, and help explain well-documented rigidities in online pricing (Gorodnichenko et al. (2015)). Importantly, we find that accounting for the retailer’s joint pricing and inventory decision is critical to assess the gains collected from commitment — optimal inventories are 133% larger in this setting than under state-contingent, dynamic pricing. Moreover, the welfare gains from this inventory expansion accrue essentially equally to retailer profit and consumer surplus. While both customer segments benefit, high valuation customers capture 89% of the additional consumer surplus, reflecting the allocative gains from commitment.

Next, we exploit the heterogeneity in monitoring behavior to demonstrate that a simple metric, the customer’s purchase-to-visit ratio (PVR), is nearly as informative about the customer’s type (i.e., her price elasticity and her cost of monitoring), as tracking her entire online history. Exploring

the design of promotional strategies targeting consumer segments, we find that targeted pricing using the customer’s PVR leads to 6% in additional retailer profit. Finally, perhaps counterintuitively, our analysis illustrates that facilitating consumers’ monitoring may intensify the availability risk consumers face from the retailer’s limited inventory to significantly benefit the seller. Considering inventory alters the profit implications of strategic consumer behavior, where the predominant focus so far has been on the behavior’s detrimental effects.

To summarize, our main contribution is to study, both theoretically and empirically, the value of intertemporal pricing with commitment in conjunction with inventory management. Motivated by our retailer’s pricing practices, we introduce the randomized markdown policy as preserving key benefits from committing to prices while exploiting the heterogeneity in consumers’ monitoring costs. Our analysis uncovers a novel feature of strategic consumer behavior: by exploiting our dataset’s detailed tracking of customers’ dynamic behavior (in addition to their purchases), we establish that customers’ active monitoring underpins substantial value from randomized markdowns and the use of simple metrics such as PVR. In contrast, strategic consumer behavior with delay costs and/or willingness to wait alone cannot fully explain this value.

1.1. Related Literature

Broadly, two prior literatures relate to our work. First, an expansive empirical literature in economics and marketing has focused on consumers’ search costs, with the growth of online retail and e-commerce re-igniting interest in this area. Second, a sizeable literature on revenue management addresses strategic consumer behavior. We discuss each in turn.

Following [Stigler \(1961\)](#)’s classic paper on price dispersion, the search cost literature has focused on the economics of information. Search costs offer an explanation for price dispersion, e.g. in mutual-fund fees ([Hortaçsu and Syverson \(2004\)](#)) and auto insurance ([Honka \(2014\)](#)), as well as for information obfuscation in online retail ([Ellison and Ellison \(2009\)](#)). For a dynamic setting like ours, [Seiler \(2013\)](#)’s structural model of purchases allows a panel of consumers to decide whether to incur an incremental cost to observe detergent prices upon each grocery visit, where the

applicable search cost varies by customer type and purchase basket. A basic result is that search cost heterogeneity gives rise to relatively informed and uninformed customers in equilibrium — in fact, [Moraga-González and Wildenbeest \(2008\)](#) estimate that most consumers search very little (at most three prices) while a small segment of consumers searches intensively.

In contrast to this literature, we focus on the retailer’s pricing policies and operations. For this reason, we consider seasonal apparel products (as opposed to repeat purchase goods) for which managing limited inventories affects both pricing and consumer behavior. Rather than emphasizing informed versus uninformed consumers, we focus squarely on repeat, online customers who are most likely to engage dynamically in the monitoring that underpins strategic purchase timing.

To the literature on revenue management for dynamic, strategic demand, we contribute a theoretical and empirical treatment of intertemporal pricing with inventory management. With empirically grounded motivation from consumers’ heterogeneous monitoring costs, we introduce randomized markdown policies to the literature, where no prior work addresses the practice of generalizing price commitments with true randomization or quantifies its value.

Traditional models in revenue management do not emphasize strategic consumer behavior, focusing rather on dynamic demand stochasticity.¹ [Coase \(1972\)](#)’s durable-good monopolist operating in a dynamic setting serves as a precursor to subsequent work highlighting strategic consumer behavior. Absent credible commitment, the monopolist is tempted to lower its retail price to sell to residual demand once “high-value” customers purchase and exit the market, and sufficiently patient customers would wait. Modeling the subgame-perfect equilibrium, [Besanko and Winston \(1990\)](#) find that “prices are always lower with rational customers than with myopic customers.”

Subsequent research examines credible commitment devices to constrain sellers’ pricing. [Coase \(1972\)](#) and [Bulow \(1982\)](#) examine leasing and other contractual arrangements; [Liu and van Ryzin \(2008\)](#) consider the role of deciding production capacity. By restricting the retailer from selling additional volume in later periods (and subjecting consumers to availability risk), these devices

¹ See [Talluri and Van Ryzin \(2006\)](#) for an overview. See also, e.g., [Gallego and Van Ryzin \(1994\)](#) and [Belobaba \(1989\)](#).

allow firms to maintain some degree of market power.^{2,3} Even when the seller remains able to increase its (production or transaction) capacity while facing residual demand, McAfee and Wiseman (2008) show that capacity costs can credibly commit the seller to constrain its capacity.

Related research characterizes optimal pricing for a finite horizon or with finite inventory.⁴ A critical distinction is whether the seller is able to commit to prices or is constrained to pricing that satisfies subgame perfection looking forward at each point in time.⁵ For the no-commitment case, Su (2007) considers a flow of customers heterogeneous in delay costs and in product valuations, while Hörner and Samuelson (2011) consider the case where valuations are private, waiting is costless, and all consumption takes place at the end. For pricing with commitment, Besbes and Lobel (2015) solve the dynamic pricing problem for a continuous flow of customers that are heterogeneous in product valuations and delay costs, while recent contributions by Caldentey et al. (2015) and Chen and Farias (2015) consider robust formulations (the latter with finite inventory). Correa et al. (2016) (two discrete selling periods) and Cachon and Feldman (2015) (single discrete selling period) focus on price commitments that are state-contingent but otherwise deterministic.⁶ For two-period pricing with and without commitment, Aviv and Pazgal (2008) consider revenue maximization for a finite quantity of a seasonal good but do not focus on consumers' dynamic costs. Prescriptions

² In contrast, a different strand of the literature (e.g., Su and Zhang (2009)) contemplates consumers' stock-out costs, affecting whether they visit a retailer in the presence of availability risk. Typically, these models focus on consumers' relatively costly visits to physical stores with localized inventories and not on purchase timing. See also Swinney (2012) on the effects of inventory pooling, especially online, on availability and strategic behavior.

³ A retailer's option to place additional, expedited orders without undermining commitment (typically a matter of when the option may be exercised) is particularly valuable for strategic and uncertain demand (Cachon and Swinney (2009a,b, 2011)). See Taylor (2006) on supply chain considerations in timing sales to retailers.

⁴ Shen and Su (2007) and Netessine and Tang (2009) provide related overviews of strategic consumer behavior in revenue and operations management.

⁵ Stokey (1979) gives an early treatment of finite-horizon optimal pricing with commitment. For time-consistent subgame perfection in a Coasian setting, see the "dynamically consistent" policy defined in Gul et al. (1986).

⁶ In contrast, Jerath et al. (2010) explore a possibly mixed strategy whether to sell through an opaque intermediary. See also Etzion et al. (2006).

are very sensitive to the detailed settings of these stylized models. For instance, [Aviv and Pazgal \(2008\)](#) find that a fixed price performs approximately as well as two prices when committed prices are advantageous, while in [Hörner and Samuelson \(2011\)](#) the fixed price is always dominated by the subgame perfect policy.

In contrast, we explore monitoring as the consumer’s active decision to engage with the retailer. We show that markdowns would not be randomized in the absence of monitoring costs, motivating the assumptions in our structural empirical work; neither delay costs nor strategic consumer behavior alone explain the randomized markdown policy (or PVR) as beneficial for the retailer.

Unlike randomized markdowns, the theory literature has sought to explain fixed prices (e.g., [Hörner and Samuelson \(2011\)](#), p.385) and pre-announced (i.e., deterministic) pricing (e.g., [Aviv and Pazgal \(2008\)](#), p.353)⁷ as viable modes of price commitment. However, more flexible practices are commonplace for retailers committing to prices; in particular, retailers frequently impose “business rule” restrictions on permissible price levels when implementing revenue optimization ([Elmaghraby and Keskinocak \(2003\)](#) pp.1296-7):

Markdown price optimization tools take as input ...business rules specified by the retailers ...such as ...[t]he allowed number and frequency of markdowns [and] ...[t]he types of markdowns allowed (e.g., 10%, 25%, and so on); equivalently, the set of permissible prices.⁸

No prior literature addresses, motivates, or values the practice of generalizing price commitments with randomization.^{9,10}

⁷ See also [Elmaghraby and Keskinocak \(2003\)](#), [Aviv and Pazgal \(2008\)](#), [Aviv et al. \(2009\)](#), and [Elmaghraby et al. \(2008\)](#)’s description of the Filene’s Basement “Automatic Mark Down System” on p.145.

⁸ Business rules apply even in settings requiring large-scale, complex promotion planning. See [Cohen et al. \(2014\)](#).

⁹ Moreover, randomizing may make a retailer’s markdown policy more robust to competition, by reducing or eliminating the incentive (in equilibrium) to undercut with everyday low pricing. [Varian \(1980\)](#) and [Sobel \(1984\)](#) characterize such competitive-equilibrium outcomes in the context of temporary sales with search costs.

¹⁰ In the mechanism design literature, randomization over outcomes underpins the useful notion of ex ante incentive compatibility. See, e.g., [Dobzinski and Dughmi \(2013\)](#).

Several empirical studies have examined strategic consumer behavior and price discrimination in retail, largely with aggregate market data. Nair (2007) studies video-game pricing with forward-looking consumers. For seasonal (apparel) goods, Soysal and Krishnamurthi (2012) use store-level data to show that consumers consider availability risk when timing their purchases. Using market-level fare and booking data, Li et al. (2014) find that a subset of air-travel customers strategically delays its purchases. Chevalier and Goolsbee (2009) show that students anticipate near-future textbook revisions (affecting used-book saleability) when deciding whether to purchase new textbooks.

Finally, randomized markdowns by exploiting monitoring cost heterogeneity add to recent empirical work on using intertemporal pricing to exploit multi-faceted consumer heterogeneity. For instance, Hendel and Nevo (2013) use store-level data to explain a pattern of frequent yet unpredictably timed markdowns in the soda market as price discrimination exploiting the positive correlation of price sensitivity and storage capacity among heterogeneous customers.

In contrast to prior work, our customer-level dataset enables us to track individual consumers' decision processes, including whether to visit and whether to purchase upon visiting. This enables us to document strategic behavior at the customer level and to identify monitoring costs as a factor outside of valuations and price elasticities with novel implications. Our analysis delves beyond existing treatments in linking pricing, inventory, and availability risk in seasonal retail.

We contribute to the use of structural estimation in the growing literature employing empirical methods to explore questions of operational interest. In addition to Li et al. (2014), some leading papers in this area are Olivares et al. (2008), Allon et al. (2011), Bray and Mendelson (2012), Akşin et al. (2013), Kim et al. (2014), and Bray and Mendelson (2015). For the estimation of dynamic games, we contribute a model of discrete choices embedded in a continuous-time framework, substantially extending Doraszelski and Judd (2012)'s work on the latter. In particular, our equilibrium-lattice results and estimation show these games to be computationally tractable (where present-day workhorse techniques do not suffice) even for customer-level data while accommodating unobserved heterogeneity across players.

2. Data and Background

Our dataset was obtained by the Wharton Customer Analytics Initiative (WCAI) from a corporate sponsor whose identity is withheld by request. The sponsor is a North American specialty retailer selling goods, including apparel, shoes, and home furnishings, that are typically sold over their own (i.e., non-calendar) selling seasons. The two channels for purchases are online via the retailer’s website or in-person at the retailer’s store locations; the retailer neither makes its apparel products available elsewhere nor facilitates a second-hand market. The retailer’s annual revenues register in the billions of US dollars, and its online channel accounts for roughly 20% of its total revenue.

2.1. Customers

Our dataset tracks 25,965 customers of one of the retailer’s brands. These customers were randomly selected from those who interacted with the brand during the period from July 2010 through June 2012. All known transactions and interactions with these customers were recorded during this period. To this end, the retailer extensively links customer identities using all available data such as credit cards, addresses, phone numbers, email addresses, click-throughs, devices, and loyalty accounts; that said, attribution would not be possible in a relatively limited set of instances, such as a browsing session involving no identifying actions on a public device.¹¹ We focus on the 11,564 online active customers who interact with the brand’s online channel, i.e., visited the website at least once during this time.

Table 1 summarizes the characteristics of the customers in our sample, separating the subsamples of online active customers and the brand’s loyalty program members. As reflected in the gender breakdown, the primary target demographic consists of women between the ages of 28 to 45, reducing the incidence of multiple end customers in a single household. Membership in the loyalty program is substantial, at 57.7%. Members receive catalogs and notifications of certain events or promotions, which other than an annual birthday discount do not typically involve price discounts. As Table 1 shows, these subpopulations do not differ notably in terms of their observed attributes.

¹¹ Such unmatched instances would be sampled in the data as part of the one-time and newly arriving demand which we later discuss as part of our model.

As summarized in Table 2, we observe 35,958 orders and 765,070 customer sessions online; over both channels, the full sample of nearly 26,000 customers made over 195,000 orders. For each session, its date and time are recorded, and we observe the product(s) viewed in detail by the customer. On a product’s detailed page, a customer can view its price, availability (but not remaining inventory), and information about the product through reviews, descriptions (covering the product’s materials, design, dimensions, colors, and care instructions), and multiple images. A product’s detailed page is accessible from a common webpage displaying the relatively small number of items for its browsing category, which itself features the product’s large image and current price.

Additionally, we are able to trace orders made by the customer during her session. Importantly for our study, 18.5% of all orders are placed online, constituting 21.75% of the total net revenues, and 10,129 customers (88% of the online active customers) place at least one purchase order online during the two-year observation period.

Table 2 further details the application of promotions to online orders. “Shipping” promotions reduce the shipping charge for an order. A “System” promotion is a temporary offer available to all customers purchasing the relevant product(s) and distinct from a permanent price markdown. Promotions of this sort are not more prevalent in either the retail or online channel. Lastly, a “User-Applied” promotion is an offer applied by a user through a promotion code. As shown in Table 2, only 2.3% of the orders in our sample involve a “User-Applied” promotion.

The retailer has also included in the dataset its marketing-related communications with each customer. These include print and email catalogs, usually disseminated monthly, promotion-related emails, and even shipment and order cancellation notices.¹² The brand typically does not engage in promotion-centered marketing, in fact intending its markdowns and promotions to be difficult to predict, and introduces new products in an approximately constant stream over time, rather than with the calendar months or seasons. Communications for the loyalty program are also recorded.

¹² Catalogs may coordinate the timing of visits from high valuation customers in particular. Coordination would tend to accelerate availability risk among customers, causing us to possibly understate price elasticities and overstate valuations for this segment — crucially, our estimates would understate the relative importance of monitoring costs.

2.2. Products for Structural Estimation

For our structural estimation, we use cardigan sweaters. Cardigans are among our retailer’s flagship products, offered in a substantial variety supported by highly distinctive designs, including brand-specific collaborations with independent designers or boutiques. As we discuss later in Section 4.1, the differentiation across cardigans by style and features make it less likely for customers to substitute rather than wait for a better price.

Our dataset includes 365 cardigan styles with at least ten observed transactions. Of the women’s tops department, these represent 11.25% of the product count and 14.19% of revenue. We focus on the ninety-eight cardigan styles launched during 2011 and early January 2012, to ensure that we observe the product’s entire season and to identify and focus on the customers acquainted with the brand at the time of the product’s launch. Table 3 summarizes the list prices of cardigans in our sample; we discuss how they are priced over a selling season in the next subsection.

As shown in Table 4, online customers start monitoring products well in advance of buying, with 25.7% viewing a product’s detail page more than once. On average, a customer views a product’s detail page 1.6 times with 15 days between visits. Customers who ultimately purchase the product visit the product’s page more frequently, 2.3 times on average, and the intervals between their visits are shorter, 9.3 days on average. This suggests that customers monitor their cardigans of interest over an extended period, with 6% buying. Likewise, the data indicates monitoring at the browsing-category level, where customers can observe prices without accessing the detail page: among brand-acquainted customers (i.e., those who interacted with the brand within the six months prior to a cardigan’s introduction) who view a cardigan’s detail page at least once, over 92% of their visits to the cardigan’s browsing category during its selling season involve no purchase at all from the category. Lastly, in contrast to prevalent monitoring, we note that strategic use of the brand’s returns policy (i.e., returning a purchased item to repurchase it at a lower price) is relatively rare.¹³

¹³ As an upper bound on possibly strategic instances, only 2.1% of the overall single-unit purchases made in the women’s tops department involve a final purchase price lower than the customer’s first-in-time transaction.

2.3. Retailer’s Pricing

Table 3 summarizes the distribution of discounts obtained by cardigan buyers. About a third of purchases occur close to the list price, followed by substantial discounts. Overall, 71.4% of transactions take place within 0.5% (i.e., rounding error) of three chronologically decreasing price levels per product. Accordingly, the data suggest that the price path for each of the retailer’s products consistently involves up to three predominant price levels: a list price, a sale price (typically about 50% of the list price), and a clearance price (typically around 28% of the list price). As discussed later, the customers in our model expect to see an offered price with idiosyncratic variation around one of these three price levels upon a visit. Confirmed in our discussions with the retailer and also tested within the context of our model, the retailer makes the precise markdown timing unpredictable to avoid “conditioning” or “training” its customers to wait for markdowns.¹⁴ The price levels and markdown times are common across the online and offline channels.

3. Panel Data Analysis

A panel data analysis of customers’ online monitoring motivates our subsequent structural model and estimation. We hypothesize that differences in customers’ monitoring explain whether they consistently purchase at discounted prices, instead of both their monitoring and discounted purchasing being driven by common factors such as promotions or purchasing volume. We aggregate the data to the customer-month level, so as to appropriately control for individual- and time-specific effects in addition to promotions. Further discussion and descriptive background on customers’ monitoring and purchasing behaviors can be found in Appendix A.

To measure a customer’s intrinsic monitoring intensity (or alternatively, proxy for her cost to monitor), we use the ratio of her in-month purchase count over her number of in-month online visits, which we call her monthly purchase-to-visit ratio (PVR). We then employ panel-data methodology to test the relationship between the customer’s PVR and the discounts she obtains.

¹⁴ Reflecting substantial variation, the 2011 sample’s 25%, 50%, and 75% quantiles for a cardigan’s markdown times (in days since its introduction) to the sale and clearance prices are (48, 69, 90) and (81, 112, 161), respectively.

We must resolve that PVR is endogenous for dependent variables measuring concurrent discount-share outcomes, since the customer’s visits, purchases, and discounts are the inter-related outcomes of her decision process. As such, they are likely to be commonly influenced by factors characterizing the current selling season, including, *inter alia*, the available products and discounts.¹⁵ We address these issues with instrumental variables, including in particular the customer’s PVR lagged by five and six months in the most conservative case to avoid confounding factors specific to the current selling season. This approach is intuitively appealing, as we are primarily interested in the component of the customer’s monitoring that is persistent, hence predictable via such instruments.¹⁶

Table 5 presents our unbalanced panel (random effects) results. For each customer-month pair $\{i, t\}$ in which a purchase is made, the dependent variable is the customer’s total amount paid as a fraction of the sum of her purchases’ list prices (PDY_{it} , her “post-discount yield”), $\frac{\sum \text{Transaction Price}}{\sum \text{List Price}}$, which can be viewed as a list-price-weighted measure of her overall discount. We find similar results when we use instead as the dependent variable the fraction, by number, of purchases made at a substantial discount (at various percentage thresholds).

We control for observed, customer-specific covariates such as her age, gender, loyalty enrollment, and distance to the retail store nearest to her residence, which we collectively denote as X_i . At the customer-month level, we control for X_{it} : the number of promotions applied (with and without instruments), and the list-price totals of purchases made in each channel (with instruments). The regression without instruments represents a linear approximation to an informative conditional expectation since the promotions applied by a customer directly affect her discounts. Therefore, the linear model specifies:

$$PDY_{it} = \alpha_i + \delta_t + X_i^T \beta + X_{it}^T \beta' + \gamma \cdot \text{PVR}_{it} + \epsilon_{it},$$

with individual (random) effect α_i , time effect δ_t , coefficients β , β' , and γ , and the error term ϵ .

¹⁵ We also expect measurement error, as monthly realizations of purchases and visits are random variables reflecting her monitoring intensity or cost with stochastic noise. Nonlinearity complicates any direct decomposition of this error.

¹⁶ This virtually precludes a fixed-effects approach to the panel.

As detailed in the notes accompanying Table 5, we use lags of five and six months of multiple variables, including PVR, as instruments. As the table elaborates, slightly over half of the sample's month-customer pairs are used, due to the availability of such lags as instruments, with these pairs including 42% of the customers in the sample. We find a statistically and economically significant relationship between a customer's PVR and her post-discount yield: the sample average impact associated with increasing a customer's monitoring by just a single visit per month is a greater than 1% decrease in the revenue drawn from the customer, holding her purchases fixed. Panel results supplied in Appendix A largely rule out the possibility that our results are driven by a selection bias from using only the relatively experienced customers with available instruments. Likewise, pooled OLS results are consistent with our random effects assumption, for which formal Hausman test results are provided in Table 5.

Now focusing on cardigans, our structural model builds on the foregoing evidence that relates customers' online monitoring to their discounted purchasing outcomes, after controlling for other factors.

4. Model

We devise a structural, dynamic-equilibrium model endogenizing customer and firm decisions. In our model, a customer decides not only whether to purchase a product (i.e., the cardigan) but also how often to monitor the product by visiting the retailer's website. A customer is motivated to monitor a cardigan to obtain and process information that determines whether she purchases, including subjectively assessing how her own tastes and time-varying state might fit with the cardigan. Since her monitoring incurs an opportunity cost, she will typically monitor the cardigan more intensively when she is more likely to make a purchase. On the other hand, even after making a visit she may forego a purchase to wait for a prospective markdown in price or to simply continue her monitoring process.

Methodologically, we embed the customer's discrete-choice purchase decisions within a continuous-time stochastic game in which she also selects her hazard rate of monitoring based on

the current state of the product’s market, which includes the product’s price level and availability as well as the forward-looking customers remaining in the market for its limited inventory. Within this framework, we will discuss and model the retailer’s endogenous pricing decisions based on the corporate sponsor’s policy in practice. Finally, an equilibrium in our game specifies a continuous-time Markov jump process over market states, i.e., a well-defined data-generating process with a likelihood over the observed data. In the remainder of this section, we describe the model’s state space and players’ decisions before defining and characterizing our equilibrium concept.

4.1. Product Market

Following [Caro and Gallien \(2010\)](#) and other works in retail operations, we treat each highly differentiated cardigan product (see [Section 2.2](#)) as its own separate market. To justify this approach, we examine whether substitution effects are sufficiently weak among our cardigans, which we preface with the following observations. First, earlier models addressing substitution are constructed with aggregated, discrete-time data in mind; in our continuous-time setting using detailed customer-level data, a customer decides whether to purchase or wait within the context of a single visit, i.e., where deciding to purchase another cardigan or sweater fails to reliably signal substitution, as she could very well return to buy within even a day or less. Second, since visits without any purchases are prevalent in the data (see [Section 2.2](#)), cross-product browsing is similarly rationalizable with customers monitoring multiple products. Both of these observations arise because a customer may purchase (single) units of multiple products in our setting. Lastly, because the retailer introduces products continually and does not engage in seasonal or promotional markdowns that take effect simultaneously across its offerings, we instead focus on whether customers tend to find certain products to be closer substitutes to the cardigan than others, suggesting that they would (and we should) devote special attention to these items’ prices and markdowns in deciding whether to buy.

Accordingly, customer-level data permits us to test for a customer’s market exit upon making other purchases. Specifically, since the strongest candidates as substitutes are other cardigans, we examine whether a customer more likely exits a cardigan’s market following her purchase of another

cardigan as opposed to following her purchase of a non-cardigan sweater. In Appendix A, we show that we do not find such evidence in our data but still discuss how substitution may affect our findings, in particular making them conservative in several important respects.

4.2. Market States

The Markov stationary equilibrium we use for our stochastic game will follow the paradigm of Markov perfect equilibria (MPE) in defining a state space that satisfies consistency and payoff-relevance (see Maskin and Tirole (2001)). Informally, while MPE impose the strong assumption that players know the current state, this state (and therefore the amount of knowledge we attribute to the players) is minimal in the sense that it is the coarsest partition of histories that naturally distinguishes the players' expected time-discounted payoffs over actions going forward.

Our model's state includes the cardigan's available inventory, I , and its price level, p . (To reiterate, the offered price in each customer's visit is the sum of this price level and any idiosyncratic price variation which we capture separately.) Moreover, players' payoffs from their strategies are affected by the level of demand for the limited inventory, captured by the number of consumers of each type in the market, where the types' respective product valuations and price elasticities are common knowledge. In the case of two types, which we will later label as high valuation customers and bargain hunters, respectively, we denote these as N^1 and N^2 . Lastly, based on empirical evidence of this pattern, we allow for the cardigan to depreciate in its utility value to consumers as it falls out of fashion during its selling season.¹⁷

Accordingly, we now define the state of the cardigan's market at time t as:

$$X_t := \{I_t, N_t^1, N_t^2, p_t, D_t\}, \quad \forall t \in \mathbb{T}, \quad (1)$$

¹⁷ As evidence supporting such depreciation, running a naïve logit static-demand specification only through each cardigan's mid-season yields higher price elasticity estimates than for the entire season: the cardigan's first markdown appears to draw a stronger demand response per dollar than later markdowns. Generally, we note that depreciation may affect availability in settings where retailers terminate unpopular styles, especially offline where retail space may be at a premium.

where D_t is an indicator variable for whether it has depreciated as of time t . In equilibrium, X_t is a stochastic process, namely a continuous-time jump process, with the state transitioning in jumps upon the occurrence of well-defined random events. For instance, a markdown would shift the price level p_t down, whereas a cardigan purchased by a bargain hunter would decrement by one unit each of the inventory I_t and the market's count of remaining bargain hunters N_t^1 . More generally, we use X to denote a market state and \mathbb{X} for the state space of market states.

To round out our introduction of a cardigan's state space, we define the start and end of its selling season. The cardigan's season starts upon its introduction by the retailer, when a fixed inventory is made available without replenishment, consistent with the retailer's operations. At this time, the cardigan's customers are randomly drawn from a finite-mixture population of heterogeneous types, have single-unit demand, and exit the market (only) upon purchase, consistent with the data in which multiple-unit purchases are rare.¹⁸ These are the customers engaged with the brand over time; we separately model one-time and newly arriving demand, including foot traffic in stores, as a price-elastic but not forward-looking stream (see Appendix B). The cardigan's season ends when its inventory depletes, a random event occurring in finite time with probability one (but no deterministic bound).

To reflect the information-rich online environment, the current state of the cardigan's market is modeled as common knowledge for all players. This may be a strong assumption, as consumers' website visits could serve to inform them (possibly incompletely) about the state. While tractably modeling consumers' private updates within a dynamic equilibrium is extremely challenging,¹⁹ Appendix D presents a tractable, alternative equilibrium framework which substantially limits the customers' knowledge of a cardigan's price level, inventory, and market demand. Our empirical

¹⁸ In the women's tops department (including cardigans), less than 3.1% of all customer-product pairs with the first transaction occurring in 2011 involved multiple units purchased over 2011 and 2012.

¹⁹ Notably, a perfect Bayesian equilibrium would track not only consumers' beliefs about the state, which entails solving partially observed Markov decision processes for the best responses, but also their beliefs about other consumers' histories of beliefs (which include beliefs about others' beliefs about beliefs, and so on).

estimates, obtained by a different estimation approach, prove robust to these relaxations. A natural interpretation is that customers who are unable to perfectly observe all payoff-relevant states can nonetheless essentially replicate the Markov stationary equilibrium strategies by tracking (as individual, private information) a few informationally relevant states, such as the time elapsed since their last-in-time visits and whether the price observed in those visits were observed before.

4.3. Consumer’s Decisions

In our model, a utility-maximizing customer decides not only when to purchase a cardigan but also how often to monitor the cardigan by visiting the retailer’s website. We capture this sequence of decisions by embedding the customer’s discrete-choice decisions to purchase or wait within a continuous-time framework where she decides her monitoring hazard rate. As we explain further, these endogenous decisions resolve separate but related trade-offs, which we discuss in turn.

First, within her visits to the retailer’s website, she decides whether to purchase or wait after observing an offered price possibly including idiosyncratic variation around the cardigan’s current price level. Because observed price variation alone is not sufficient to rationalize decisions in the data, incorporating customer-specific sources of uncertainty (including time-varying tastes and browsing) within visits is necessary.

Consider a customer, indexed by i , visiting the retailer’s website at time t . She may choose to purchase the cardigan, receiving the conditional, instantaneous utility payoff:

$$v_i - \alpha_i \cdot p_t - \omega \cdot D_t + \epsilon_{i,t}^{\text{Buy}}, \quad (2)$$

where v_i is her mean present-valued utility from a purchase including the cardigan, α_i is her price elasticity, ω is the cardigan’s utility value lost when it falls out of fashion, and $\epsilon_{i,t}^{\text{Buy}}$ the idiosyncratic utility shock affecting her current valuation of the cardigan and its subjective fit. In addition to idiosyncratic price variation, unobserved states captured by this visit-specific shock may include idiosyncratic shocks to her fashion taste or needs and chance variation in her browsing and shopping behavior.²⁰ We take these to be i.i.d. over customers and visits, private, and unknown ex ante.²¹

²⁰ In our model and setting, such sources of uncertainty persist even after fully learning a product’s attributes.

²¹ This satisfies the conditional independence property stated by Rust (1994) for the discrete-time setting.

Alternatively, the customer could elect to wait, either by making other purchases from the retailer or by not buying at all. When making other purchases, she derives the instantaneous net payoff:

$$w_i + \epsilon_{i,t}^{\text{Other}}, \quad (3)$$

where w_i is her mean net payoff (after accounting for expenditures) while the unobserved state $\epsilon_{i,t}^{\text{Other}}$ captures variation in the particular purchases she makes as well as the subjective value she assigns to them in that visit. To define other purchases, we include the cardigans and sweaters available within the cardigan's common browsing category; since they can be monitored together with the cardigan within a single viewing, these products involve a complementarity in monitoring that we account for in this way. Finally, we normalize her instantaneous net payoff from a visit without a purchase as the mean-zero idiosyncratic shock, $\epsilon_{i,t}^{\text{No Purchase}}$. Note that with these choices, she retains the option to buy the cardigan later. To simplify the dynamic optimization problem to follow, we assume that consumers' idiosyncratic utility shocks, $\epsilon_{i,t}$, follow the mean-zero, type 1 extreme value distribution.

Second, the customer weighs her monitoring against her incurred opportunity costs, which are increasing in the frequency with which she monitors. Her cost function should be flexible enough to capture both her baseline monitoring as well as her responsiveness to changes in her expected payoff from a visit. To this end, we use the cost function δ in the form of a second-order polynomial in her monitoring hazard rate λ :²²

$$\delta(\lambda) = r_{i,1} \cdot \lambda + \frac{r_{i,2}}{2} \cdot \lambda^2, \text{ with parameters } r_{i,1}, r_{i,2} > 0, r_{i,1} \leq \Gamma, \quad (4)$$

where Γ is Euler's constant.

With these primitives, the customer solves a dynamic optimization problem to derive her best response (maximizing her expected ρ -discounted payoffs²³) to the strategies of others. As we discuss

As standard for customer-level or scanner data in the literature, our econometric error structure does not account for potential endogeneity of the unobserved states and price. For potential issues in this widespread approach, see [Villas-Boas and Winer \(1999\)](#).

²² The upper bound on $r_{i,1}$ will guarantee a positive monitoring hazard rate and is not binding in estimation.

²³ The customer's time discounting of her payoffs by the discount factor ρ is distinct from the cardigan possibly depreciating in its present-valued utility to her during the season (see Expression (2) above).

and prove in Section 4.5, she has a best-response strategy that is Markov stationary given that her opponents are playing Markov stationary strategies. In this case, her best-response monitoring hazard rate $\lambda_X^* \in [0, M]$ solves the following continuous-time Bellman equation for all states $X \in \mathbb{X}$:

$$\rho \cdot V_i(X) = \max_{\lambda \in [0, M]} \left\{ -r_1 \cdot \lambda - \frac{r_2}{2} \cdot \lambda^2 + \sum_{y \in S(X)} q(X, y) \cdot (V_i(y) - V_i(X)) \right. \\ \left. + \lambda \cdot [\Gamma + \log(1 + \exp\{w_i\} + \exp\{\gamma_i + v - \omega \cdot D_t - \alpha_i \cdot p_t + V_i^0 - V_i(X)\})] \right\}, \quad (5)$$

where $S(X) \subset \mathbb{X}$ denotes the set of successor states of X into which a direct state transition is feasible, $q(X, y)$ is the exogenous transition hazard rate from X into a successor state $y \in S(X)$ taking as fixed the strategy profile (i.e., other players' strategies), and V_i^0 is shorthand for her continuation value following market exit. Her closed-form, Markov stationary monitoring hazard rate is then:

$$\lambda_X^* = \frac{\Gamma + \log(1 + \exp\{w_i\} + \exp\{\gamma_i + v - \omega \cdot D_t - \alpha_i \cdot p_t + V_i^0 - V_i(X)\}) - r_1}{r_2}, \quad (6)$$

and, for any visit time $t \in \mathbb{T}$ for customer i , the conditional probability of her purchasing the product at time t is given by:

$$\frac{\exp\{\gamma_i + v - \omega \cdot D_t - \alpha_i \cdot p_t + V_i^0 - V_i(X_t)\}}{1 + \exp\{w_i\} + \exp\{\gamma_i + v - \omega \cdot D_t - \alpha_i \cdot p_t + V_i^0 - V_i(X_t)\}}. \quad (7)$$

For further discussion of the customer's dynamic optimization problem, see Appendix C.

4.4. Retailer's Decisions

We model the retailer's pricing policy based on what we observe in the data as well as our conversations with the corporate sponsor. First, the retailer pre-commits to three price levels, which we refer to as the list, sale, and clearance prices. As discussed earlier, an overwhelming percentage of observed cardigan purchases occur at these prices, which are roughly customary percentages of their list prices. This leads us to believe that brand-acquainted customers should form reasonably accurate expectations of these price levels, which we assume to be common knowledge as of the start of the cardigan's season.

Second, the markdowns that transition the price level to the sale and clearance prices follow from the retailer’s constant, state-independent markdown hazard rate, known to all players. Confirmed in our discussions with the retailer and estimates, the retailer relies on making the precise markdown timing unpredictable to avoid “conditioning” or “training” its customers to wait for markdowns. To test whether the retailer in fact follows a constant markdown hazard rate, we more generally allow the retailer to endogenously select at each point in time a markdown hazard rate from among a pre-determined pair of high and low hazard rates, λ^{R-High} and λ^{R-Low} , to maximize its expected profit going forward. (We test for state independence as the high and low hazard rates being equal.) In this case, the price levels and the high and low markdown hazard rates are common knowledge.

4.5. Markov Stationary Equilibrium

We describe our stochastic game and its Markov stationary equilibrium. In the Electronic Companion, we re-state and elaborate on Propositions 2-4; likewise, all formal definitions and proofs are found in the Electronic Companion’s Appendices B and C.

Importantly, we restrict our attention to equilibria in which all participants play Markov stationary strategies: however, for each player, her Markov stationary equilibrium strategy is a bona fide best response to the equilibrium play of others without artificially imposing that she must play only stationary Markovian strategies.²⁴ We first guarantee that such an equilibrium exists.

PROPOSITION 1. An equilibrium in stationary Markovian strategies (a “Markov stationary equilibrium”) exists for the stochastic game defined in Appendix B.

In this equilibrium, each player at each point in time plays her Markov stationary strategy that unilaterally maximizes her expected discounted payoff going forward under her time discount factor ρ . Informally, we can think of a customer’s Markov stationary strategy as deciding her monitoring hazard rate and her purchase probability (i.i.d. across her visits) for each state that she could

²⁴ For stochastic games, proving the existence of a Markov stationary best response can involve technical issues: see Appendix C for references on the “self-referential” problem whereby properties required of the player’s environment must be shown for her best-response policy itself, since it is part of the environment for other players.

possibly face. Each strategy profile of Markov stationary strategies across all players specifies a continuous-time Markov jump process for states, $\{X_t, t \geq 0\}$, under which each player’s expected ρ -discounted payoff is well-defined. Since the cardigan’s current state history constitutes common knowledge, a player’s rational expectations essentially amount to knowledge of the state transition rates (including others’ purchase rates) she faces under equilibrium play.

For estimation, Proposition 1 allows us to search only the subspace of Markov stationary strategy profiles, while immediately implying that any equilibrium found in this subspace remains a valid equilibrium when we remove the restriction to Markov stationary strategies. We highlight the key properties that we exploit to make this search computationally tractable while also providing useful intuition about our model’s identification.

First, a contraction mapping (provided as the operator T in our proof of Proposition 1) guarantees that we can compute a player’s unique Markov stationary best response for any Markov stationary strategy profile. A key advantage is that our continuous-time Bellman equation simultaneously captures both the customer’s monitoring and purchasing decisions, obviating the typical need to solve multiple, nested Bellman equations (lacking guaranteed properties) for her sequential, forward-looking decision-making.

Second, we prove a special lattice structure over the model’s equilibria that derives from a property of strategic complementarity across players. This property can be explained intuitively for our setting: for a given customer, her expected discounted payoffs are affected by other customers only through their Markov stationary purchase rates (i.e., the state-wise product of a customer’s monitoring rate and purchase probability). We define a Markov stationary strategy to be more aggressive than another if it entails a weakly higher purchase rate in every state, defining a partial order \geq over Markov stationary strategies. The following Proposition 2 formalizes the idea that when a customer faces a more aggressive strategy from a competitor, her own best response is likewise more aggressive. This dynamic of meeting aggression with aggression is the outcome of consumers’ behavior affecting downstream availability risk due to a cardigan’s finite inventory.

Concepts and definitions are further clarified in Appendix B, where we also discuss why existing notions of strategic complementarity are inadequate.

PROPOSITION 2. *Fix the retailer’s markdown hazard rate policy, and consider the resulting dynamic game among the customers. Furthermore, consider only the class, \mathbb{C} , of strategy profiles, $s \in [0, M]^{|X| \times N}$, such that each player is weakly more aggressive (higher in effective purchase rate) in the successor states that immediately follow the exit (with purchase) of another customer. Then for any customer $i \in N$, her best response s_i^* is increasing in s_{-i} , i.e., given $s_{-i}, \hat{s}_{-i} \in \mathbb{C}_{-i}$:*

$$\{s_i^*(s_{-i}), s_i^*(\hat{s}_{-i})\} \subset \mathbb{C}_i$$

$$s_{-i} \geq \hat{s}_{-i} \implies s_i^*(s_{-i}) \geq s_i^*(\hat{s}_{-i}).$$

Proposition 3 highlights two important implications from the lattice structure over the model’s equilibria. First, we exploit for computation that best-response dynamics converge to equilibria. Appendix B explains how our model is tractable even for very large numbers of customers, whereas the size of the associated state space falls well beyond the viable scope of present-day workhorse techniques. Second, notwithstanding that the Markov stationary best response is unique, there may exist multiple equilibria. The lattice structure imposes “maximal” and “minimal” equilibria, which are one and the same in the case of uniqueness, with stable convergence dynamics that can be used to rule out multiple equilibria in practice, which we have done throughout.

PROPOSITION 3. *Fix the retailer’s markdown hazard rate policy, and consider the resulting dynamic game among the customers. Assume that all equilibria are in the complete lattice, \mathbb{C} , defined in Proposition 2. Let τ be a parameter such that the (Markov stationary) best response functions defined in the proof of Proposition 1 are weakly increasing in τ for all players. Then:*

- (i) *There exist extremal (smallest and largest) equilibria that are each increasing in τ .*
- (ii) *Starting from any pre-change equilibrium, best-response dynamics lead to a weakly larger equilibrium in response to an increase in τ .*

Lastly, the lattice structure elucidates our model’s monotone comparative statics in Proposition 4 of Appendix B — these results are helpful for identification. In Appendix B, we discuss in greater detail our model and Propositions 2-4. In the next section, we design an estimation procedure to computationally handle customers’ unobserved heterogeneity with customer-level data.

5. Structural Estimation and Results

In this section, we outline our maximum likelihood estimation procedure and present our results.

For given parameter values and an assignment of customers into types, the Markov stationary equilibrium prescribes a conditional likelihood over observed outcomes. Since the true customer segmentation is unobserved, we treat customers’ type assignments as latent random variables to be integrated out for the unconditional likelihood. Given the high-dimensional support, we use Markov Chain Monte Carlo (MCMC) for our numerical evaluations. Rather than directly simulating and maximizing the model’s log likelihood function, we use the EM algorithm to avoid nonlinearity bias and machine precision issues while benefiting from more selective sampling from the support.^{25,26} We run the E-Step of our Monte Carlo EM (MCEM) independently in parallel across cardigans. Statistical inference is performed by estimating the observed Fisher information, using a simulation approach (Louis (1982)) that relies on the missing information principle.

Our maximum likelihood estimates are consistent in the number of products. Each observed cardigan market, $j \in \mathbb{J}$, is taken to be an independent instance of an equilibrium outcome. However, we expect consumers’ demand for a given cardigan to be correlated with certain observables. First, a rational, profit-maximizing retailer would tend to assign higher list prices and to procure larger inventories for cardigans expected to draw higher consumer valuations and demand. Furthermore, demand exhibits some predictable seasonality. For these reasons, we adopt the correlated random

²⁵ See Dempster et al. (1977) on the EM algorithm and Chib (2001) for an overview of MCMC and Monte Carlo EM.

²⁶ Addressing nonlinearity bias for Simulated Maximum Likelihood Estimation (SMLE) is computationally expensive in our case, whether by sampling or by Gouriéroux and Monfort (1997)’s first-order bias correction. Typical alternatives, such as the Simulated Method of Moments (McFadden (1989)), forego the advantages of maximum likelihood estimation including the efficiency and quasi-MLE rationales. See, e.g., White (1982).

effects specification described in Appendix B to permit customers' valuations of cardigans to be correlated with these factors.²⁷ For each cardigan, the correlated random effects' residual and the depreciation date are also latent random variables sampled in the estimation's E-Step.

For each cardigan, we restrict the sample of customers to those: (i) who are brand-acquainted (i.e., interacted with the brand within the six months prior to the cardigan's launch); and (ii) who viewed the cardigan's detail page at least once. Cardigans' detail pages are accessible from a relatively small, common webpage for the browsing category, which itself features pertinent information about each product including its image and currently offered price. For this reason, we count as an instance of monitoring each online session where a customer views a product detail page accessible (only) from the cardigans' browsing category. Because online customers also buy in the retail channel, we allow for purchases in-store, as a de facto visit with purchase, following online monitoring, and we assume that all inventory, which we do not observe directly, is reflected in our dataset's transactions as sales. For tractability, we also treat a purchase-and-return instance as a de facto visit by the customer without modeling possibly richer information that may flow from such an interaction; as discussed before, strategic use of the returns policy is rare. For each player, we assume an instantaneous time discount rate equivalent to a monthly 0.975 discount factor as in Nair (2007). For the retailer, we discount at the same rate, but also find that our counterfactual results hold when the seller maximizes its expected undiscounted revenues.

Our structural parameter estimates are presented in Table 6 and highlight the significant differences across two customer segments. About 18% are "bargain hunters", who are highly price elastic and incur low costs to visit the retailer's website. The remaining "high valuation" customers exhibit lower price elasticity and incur substantially higher costs in visiting the retailer's website.²⁸

As shown in Table 7, our estimates imply that inducing an additional visit per month from a

²⁷ Estimating product-specific valuations instead would introduce the incidental parameters problem. See Section 23.2.3 of Cameron and Trivedi (2005) for discussion of correlated random effects in the panel regression setting.

²⁸ For price elasticities, the segments we find are consistent with earlier studies. For instance, Soysal and Krishnamurthi (2012) assigns 80% of customers to a "fashion-sensitive" segment and 20% to a "price-sensitive" segment.

bargain hunter would require a per-visit, expected compensation of \$2, while drawing the same additional interaction from a high valuation customer would exact over \$25 per visit. As shown in Appendix E, estimating three customer segments yields qualitatively consistent findings.

The segments' contrasting behavior is evident in the empirical distributions shown in Table 8. Bargain hunters visit the website much more intensively, accounting for 62% of all visits made while constituting about 18% of the customer population. On the other hand, their purchases compose slightly less than 18% of all customer purchases by number. In terms of timing, bargain hunters are least active in visits and purchases by share at the list price and most active at clearance — they register 61% of visits and 16% of purchases made at the list price compared to 65% of visits and 21% of purchases made at clearance.²⁹

We can also turn to the estimated model's predictions for an indication of how the segments behave across market states. Looking at the behavior predicted for the initial state in each market, bargain hunters are expected to visit the retailer's website every 2.5 days on average, compared with every 20 days for others. While the in-visit probability of purchase is always lower for a bargain hunter, her higher visit frequency implies that her effective purchasing rates in the first post-markdown and first post-clearance states are respectively 38.4% and 94.1% higher than for high valuation customers. On the other hand, her effective purchasing rate is 33.5% lower on average at the list price. In summary, bargain hunters who constitute roughly one-fifth of the retailer's (repeat) customer population delay their purchases, monitor much more frequently, and anticipate and exploit price markdowns more effectively.

Returning to Table 6, our model's nuisance parameters highlight certain aspects of the retail environment. They indicate that, on average, a cardigan depreciates in its value to customers roughly 100 days into its season, which is consistent with typical seasonal apparel life cycles at many specialty retailers.³⁰ The retailer does not engage in an aggressive discounting policy, allowing on

²⁹ Because monitoring is cheap, bargain hunters can choose to monitor a product even if their valuation is quite low.

Accordingly, while quite responsive to markdowns, not all bargain hunters purchase immediately or aggressively.

³⁰ See, e.g., Mattioli (2012) and p.85 of Kumar (2005).

average 56 days to discount, implying that clearance is typically reached shortly after the cardigan tends to depreciate. We do not find a statistically significant difference in its discounting hazard for states where a discount would improve its expected, discounted revenue. Our correlated random effects estimates suggest that a cardigan’s initial inventory and its list price are both positively correlated with its valuation. Moreover, demand is stronger in November and December, consistent with the retailer’s overall revenue and traffic.

6. Implications for Retail Operations

Consumers’ online monitoring behavior translates into several implications for the retailer. First, we examine the seller’s pricing policy, focusing on how to price to exploit heterogeneity in monitoring costs specifically via our retailer’s randomized markdown policy. Aligning the retailer’s pricing and inventory policies captures an important synergy — compared to state-contingent, dynamic pricing, our retailer sets its inventories substantially larger, with considerable welfare implications.

Second, we seek to understand what information is most valuable for retailers. While our retailer records rich online histories in addition to purchases, we show that tracking a simple metric, the customer’s purchase-to-visit ratio (PVR), is as informative as her entire online history. Our finding is important for implementing targeted price discrimination: using PVR to design a well-targeted promotion campaign does not require the retailer to store and analyze troves of customer data.

Finally, we consider the strategic implications of monitoring costs for the retailer. At the end of the day, a seller will need to understand whether facilitating customers’ monitoring behavior would increase or decrease its profits.

6.1. Commitment and Randomized Markdowns

Prior literature has suggested that in the presence of strategic consumer behavior, pricing commitments can be rewarding for the seller, and that commitments may be more credible for seasonal products. [Aviv and Pazgal \(2008\)](#) use numerical studies to suggest that pricing commitments can increase expected revenues by up to 8.32% over subgame-perfect, contingent pricing for seasonal products. More recently, [Dasu and Tong \(2010\)](#)’s numerical study suggests that neither may dominate, with a maximum performance disparity of 1.6% between the alternatives. To the best of

our knowledge, there is no prior empirical treatment of this question using actual data which is precisely the focus of this subsection.

In addition to differences in our setting vis-à-vis prior literature, numerical studies typically do not fully reflect that a seller’s inventory levels are chosen to optimally complement its chosen pricing strategy.³¹ In this subsection, we present counterfactual comparisons accounting for the seller’s joint decision over pricing and inventories.

We estimate profits using our retailer’s reported gross margin. For each product, the seller chooses its pricing policy and inventory level to jointly maximize its expected profit across customer types and one-time visitors, including foot traffic — for this purpose, all our counterfactuals assume complete information about repeat customers’ types, which we justify in further detail in the next subsection 6.2.

We compare the randomized markdown policy to a state-contingent, dynamic pricing policy (“state-contingent pricing” hereafter). Under state-contingent pricing, the seller may offer a new price upon each state transition (e.g., sale transaction). State-contingent pricing policies are subgame-perfect, solving a continuous-time Bellman equation to set a price to maximize the retailer’s revenue-to-go (a common objective in practice) in each new state as it is realized.³²

For the randomized markdown policy, the corporate sponsor’s clearance price typically falls near or below marginal cost, which we interpret to indicate that at clearance, the seller attempts to dispose of its remaining inventory. Including clearance, the average and median length for a product season is consistently very close to 150 days. To preserve the retailer’s inventory turnover cycle, we maintain the retailer’s markdown hazard rate and its clearance price for each cardigan but optimize over its inventories and the list and sale prices offered. For a fair comparison, state-contingent pricing includes a season-ending hazard rate with an expected 150-day duration (with our comparisons

³¹ Theory (Liu and van Ryzin (2008)) and empirics (Soysal and Krishnamurthi (2012)) show that inventory constraints impact strategic consumer behavior but do not compare committed and contingent pricing in their presence.

³² Customers’ equilibrium behaviors are likewise Markovian. Retailers’ detailed practices for state-contingent pricing and retailing may affect and alter what different customers can track and know. However, any additional returns to creating and using uncertainty to price discriminate are likely bounded by our results on targeting in Section 6.3.

qualitatively unchanged when allotting a more favorable 180 days to state-contingent pricing). No clearance pricing is required, and the seller may price for profit through the season’s last day. Neither class of policies enforces a monotonic price path, although we find that the optimal randomized markdowns for our sample are in fact monotonically decreasing in price level over time.

We present our counterfactual comparison in Table 9. Expected profits register 81% higher with randomized markdowns over state-contingent pricing. This improvement is more than a five-fold multiple of the figures reached by applying our retailer’s reported gross margins to prior work. Our findings suggest the value of committed prices in seasonal retail to be much higher than previously documented or predicted.

Importantly, the two pricing policies induce substantially different inventory levels, with broad welfare implications. In particular, the retailer in the absence of commitment chooses a lower initial inventory level as a way to increase the availability risk among its consumers. Because the randomized markdown policy contrastingly does not rely on availability risk for its pricing credibility, its jointly optimal inventories are a striking 133% greater than those optimal under state-contingent dynamic pricing. These results are illustrated in Figure 10, which plots for one cardigan from our sample the cost of inventory and expected revenues over a range of inventory order quantities. As the quantity increases, the expected revenue under the state-contingent policy lags substantially, because the seller cannot credibly sustain high prices. Furthermore, rather than the seller capturing the entire welfare gain from this inventory expansion under committed pricing, we find that customers share almost equally in the added surplus. While both customer segments benefit, 89% of customers’ gains accrue to high valuation customers, as markdowns exploiting the retailer’s commitment power better allocate retail inventory across the two segments.

These concerns are not lost upon practitioners. Aggressively applying state-contingent revenue management tools may seriously erode a retailer’s bottom line, and the business rules commonly imposed on pricing tools in practice (see, e.g., [Elmaghraby and Keskinocak \(2003\)](#)), especially those pre-determining permissible prices, may reflect valuable commitment power. (In Appendix

E, we discuss why using only a few price levels in the retailer’s markdown policy can be important for making its commitment credible.) As retailers’ transaction costs of adjusting posted prices have fallen dramatically, especially online, short-run concerns about seasonal lost sales and excess inventory have been used to explain the recent popularity of dynamic revenue management tools in the retail sector. However, while highly contingent dynamic pricing promises short-run gains from better matching supply and demand, long-run inventory effects bolster the relative profitability of committing to a pricing policy.³³ Committed dynamic policies are an area of active research today.

Returning to our data, we find that our corporate sponsor deployed its randomized markdown policy very effectively in practice, achieving over 82% of our counterfactual optimal markdown profits. It is likely that some of the remaining difference is attributable to actual sale prices being set at more credibly customary percentages of the list price. The sponsor’s observed markdown policies clock in with a nearly 50% gain in expected profit over state-contingent pricing, with 126% larger optimal inventories, as shown in Table 9. Customers’ welfare gains also mirror the optimal markdown outcomes, but the optimal policy modestly increases net consumer surplus by raising the welfare of the high valuation segment at a small detriment to the bargain hunters. This further demonstrates the allocative role of markdown pricing.

The intuition underlying these profit results is shown using the simplified model in Appendix F. This analysis illustrates that randomization can be advantageous to the seller in the presence of heterogeneous monitoring costs, as randomized markdown times disadvantage the high valuation segment in obtaining units at a markdown. With time discounting but no monitoring costs, the seller instead prefers deterministic markdown times, i.e., an announced price schedule.

Note that these intuitions and practices may be modified in different settings of practical relevance.³⁴ While we find monotonically decreasing price levels optimal for markdowns when facing a body of experienced customers, these prescriptions may change when facing a more substantial

³³ These considerations of long-run profitability are distinct from oft-cited concerns about deep discounts damaging a retailer’s brand image and customer perception to lasting detriment.

³⁴ We thank an anonymous referee for this suggestion.

flow of new customers, when demand or products are recurring, or when products are strong substitutes. When demand recurs or regenerates as with new customers or nondurables, randomizing increasing or nonmonotonic price paths may benefit the firm as in [Hendel and Nevo \(2013\)](#). Substitution in particular may incentivize markdowns with broader operational considerations in mind, for instance ahead of new product introductions. On the other hand, an option to add inventory mid-season may de-emphasize the distinctions across these settings, as the retailer may use the option to “reset” the balance of demand and inventory in response to realized variations.

6.2. The Informational Content of Customer Histories

We address how a retailer may glean knowledge about its customers from their online histories, which we find to be highly predictive of their types. We consider the following experiment: after observing a customer over a season, we compute a belief about her type based on information about her past behavior. In particular, we explore using a simple metric, the customer’s purchase-to-visit ratio (PVR), defined as the ratio of her purchase count over her online visit count each over the past season, and her purchase history to perform the update. For comparison, we benchmark this against the alternatives of [\(i\)](#) using only purchase histories and [\(ii\)](#) using both purchase and online histories. In each case, the prior is based on the segments’ population shares, and we assume that the types’ true behavior is known as estimated.

After performing this update, we give the posterior variance over each individual customer’s type in [Table 11](#). For the customer drawn randomly from our sample, the expected variance in the type indicator is 0.021 when using the purchase history only. Cross-sectionally, the post-update variance remains nontrivial for roughly a quarter of customers. In contrast, the expected variance in the type indicator variable is about 0.004 when either using the PVR with a purchase history or using both the purchase and online histories. More importantly and contrastingly, cross-sectionally only about a couple percent of customers are subject to any nontrivial uncertainty in type.

Not only does using PVR essentially extract the value of the full histories, we find that the level of residual uncertainty about customers’ types becomes negligible. Motivated by the operational simplicity of PVR in achieving essentially complete information about whether a customer is a bargain hunter, we turn to the design and profitability of targeting customers with promotions.

6.3. Targeted Promotions

The classic paper by Rossi et al. (1996) uses canned-tuna sales data to demonstrate the value of household-level purchase histories in target couponing. In our setting, we must account for both dynamics and monitoring when valuing targeted promotions. It is well known that strategic consumer behavior alters the implications of dynamics even with a single offered price — for instance, the basic reasoning of Lazear (1986) that expected revenue from two periods dominates that from one no longer holds. Monitoring costs deepen the departure from such clean reasoning, since consumers may visit neither regularly nor uniformly. On the other hand, online histories and PVR are highly informative relative to the purchase histories exploited in Rossi et al. (1996).

To value targeting in our dynamic setting, we first define a targeted pricing policy as one in which three prices are offered at each time — one for exogenously arriving one-time demand, including foot traffic, one for the high valuation segment, and one for the bargain hunters. For a randomized markdown policy with targeting, markdowns occur simultaneously across these segments; for state-contingent, dynamic pricing with targeting, a revised price for each segment may be offered in each state. We envision that such policies would be implemented in practice through targeted promotional offers. As before, we jointly optimize the seller’s pricing and inventories.

While targeted promotions allow the seller to segment demand by price, the segments still share availability risk from the common inventory pool. It is unclear ex ante whether segmenting demand makes state-contingent flexibility in pricing more valuable. Our findings are shown in Table 12. Targeted promotions add 6% in profit over either of the baseline policies but even under complete information generate no material effect on the comparative value of committing to prices. The seller still constrains inventories to make its state-contingent prices credible. More broadly, welfare reflects a net loss in consumer surplus in each case, as the retailer captures more. (Even then, targeted randomized markdowns remain welfare-improving over both state-contingent policies.) High valuation customers face higher prices and availability risk, reducing their surplus, while the retailer still sells slightly more inventory overall by offering lower prices to the bargain hunters throughout the season. The targeted policies are shown for a cardigan in Figure 13.

As an important caveat, consumers may adjust their behavior or resources in response to a retailer targeting its promotions using a simple metric like PVR. On the other hand, high valuation customers find it substantially more costly to expend the time to adjust, and they do not appear to have adjusted meaningfully in response to the operationalized discrimination imposed by randomized markdowns. Nonetheless, care should be taken to: (i) correctly assess the true benefit of PVR-based promotions in practice; and (ii) employ a metric robust to the easiest forms of manipulation (e.g., PVR with days instead of sessions visited in its denominator to avoid manipulation from a customer repeatedly refreshing her browser sessions).

6.4. The Effects of Shifting Monitoring Costs

Since at least [Besanko and Winston \(1990\)](#), the predominant benchmark for understanding the consequences of strategic consumer behavior has been the seller’s loss in profits or revenues due to consumers being strategic rather than myopic. The magnitude of this “strategic consumer” loss has been quantified in empirical work including [Soysal and Krishnamurthi \(2012\)](#) (−8.8% revenue with the appropriate inventory stocking adjustment, and −34.6% without) for seasonal apparel. By virtue of this intuitively clear and measurable comparison, consumers’ strategic behavior is commonly assumed to hurt the seller. To our knowledge, the only prior counterpoint is [Su \(2007\)](#)’s note that enabling bargain hunters to delay their purchases can help the seller’s revenues, because high valuation consumers expect availability risk from bargain hunters purchasing at markdown.

As shown in [Table 14](#), our findings expand upon the implications of inventory scarcity. Under the randomized markdown policy, monitoring costs impede the high valuation consumers’ response to random markdowns when compared to bargain hunters. Raising (decreasing) the bargain hunters’ monitoring costs mitigates (exacerbates) this disadvantage, causing the high valuation segment to tend to purchase later (earlier). However, the profit swings may be attenuated by inventory planning: appropriately scaling back inventory would preserve some availability risk, exerting pressure to buy even when bargain hunters’ monitoring costs are high.

We also find that incrementally decreasing the monitoring costs of the high valuation segment grows the seller’s expected profits. Part of this outcome is a demand effect — the affected consumers

both engage with the retailer and purchase more often. Part is strategic, as the affected consumer can afford to behave more aggressively in all states, all else equal, although she may elect to do so selectively rather than uniformly. To the extent equilibrium purchasing rates escalate to become more intense and costly at markdowns without affording the high valuation consumer a commensurate gain in her chance of purchasing at a discount, she will be incentivized to smooth her visits over time by becoming more aggressive earlier as well since doing so is now less costly.

In fact, our retailer has been an aggressive, early mover in developing a sophisticated mobile app experience for its brands, offering immersive and well-curated browsing and shopping. The design seamlessly integrates mobile and non-mobile online sessions and in-store shopping — features include cross-platform access to shopping carts and barcode scanning to access reviews and recommendations. The iOS app in particular has posted strong and steady download rankings throughout its history, indicating substantial adoption. Our analysis supports the qualitative contention that these moves are likely to have benefited the firm.

7. Concluding Remarks

We study, theoretically and empirically, retail operations in a dynamic setting where operational considerations affect both the retailer’s pricing strategies and consumers’ behavior. We introduce the randomized markdown policy, which combines price commitment with the exploitation of consumers’ heterogeneous monitoring costs in the presence of inventory availability risk. Importantly, we account for joint decision-making for the firm’s pricing and inventory.

Using a novel dataset tracking customer-level purchasing and online histories, we find that consumers are substantially heterogeneous in their monitoring, with their opportunity costs for an additional online visit per month ranging from \$2 to over \$25 and inversely correlated with their price elasticities. Bargain hunters visit the retailer’s website frequently and opportunistically while making purchasing decisions. Novel implications are motivated by consumers’ active monitoring behavior; neither delay costs nor rational consumer behavior alone suffices in this regard. In the absence of monitoring, the seller in fact prefers a deterministic, pre-announced price schedule.

We find that the randomized markdown policy nets a 81% gain in profit over state-contingent, dynamic pricing when the retailer’s inventories and pricing policy are jointly decided. Employing simple measures such as the price-to-visit ratio, a seller can infer with near certainty a repeat customer’s cost to monitor and may add up to 6% in additional profit by offering targeted prices. Accounting for the retailer’s joint pricing and inventory decision is crucial — optimal inventories are 133% larger with price commitments, with the important welfare consequence that consumers gain as much as the retailer does from the accompanying growth in surplus. Reflecting the improved allocation from pricing with commitment, 89% of the new consumer surplus accrues to high valuation consumers. Lastly, facilitating strategic monitoring may benefit the seller by intensifying customers’ availability risk given the seller’s appropriate choice of finite inventory.

Table 1 Summary of Observed Customer Attributes

	% Known	Full Sample	Loyalty Enrolled	Online Visitors	No Online Visit
<u>Age (years)</u>	<u>71.7%</u>				
Mean		40.6	40.2	40.5	40.7
Median		40	39	39	40
Q25		29	29	29	28
Q75		51	50	51	52
<u>Gender</u>	<u>94.8%</u>				
Female		87.0%	91.6%	88.3%	86.0%
Male		7.8%	3.6%	7.5%	8.0%
<u>Loyalty Enrollment</u>	<u>100%</u>	57.7%	100%	59.0%	56.7%
<u>Distance to Nearest Retail Store (mi.)</u>	<u>98.9%</u>				
Median		5.6	5.2	6.3	5.2
Q25		2.6	2.5	2.8	2.5
Q75		13.3	11.5	16.6	11.4
<u>Sample Size N</u>		25,965	14,991	11,564	14,401
<u>% of Full Sample</u>		100%	57.7%	44.5%	55.5%

“% Known” is the percentage with the attribute observed. “Loyalty Enrolled” refers to enrollees in the loyalty program within the sample period. “Online Visitors” have at least one online visit or purchase during the sample period. “No Online Visit” refers to the complement.

Table 2 Summary of Customer Activity

Online Channel:		
	Total	Per-Capita Mean
Online Visits	765,070	66.2 [†]
Online Orders	35,958	3.11 [†]
		3.55 [‡]
% of Total Observed Net Revenue	21.7%	—
Orders with Promotion(s) Applied	11,332	0.98 [†]
		1.12 [‡]
Any Type	31.5%	
Shipping	18.8%	
User-Applied	2.3%	
System	14.0%	
Retail Channel:		
	Total	Per-Capita Mean
In-Store Orders	159,672	7.30 [◇]
% of Total Observed Net Revenue	78.3%	—
Sample Size $N = 25965$		

Mean for customers: 11,564 with online activity ([†]); 10,129 placing at least one online order ([‡]); 21,884 purchasing at least once in-store ([◇]). An “Online Visit” is an Internet session logged on the retailer’s brand website, including multiple sessions within a single day. An “Order” is a transaction that may encompass multiple purchases. Multiple promotions may apply to an Order. “Net Revenue” is revenue net of returns. “Online activity” refers to Online Visits and Orders. A “Shipping” promotion reduces or waives an Order’s shipping charge. A “User-Applied” promotion is applied by the customer using a promotion code. A “System” promotion is an offer available to all customers at the time of the Order. All figures are for the sample period: July 2010 to June 2012.

Table 3 List and Transaction Prices for Cardigans

List Price Range	% of Sample in Range
\$50-\$100	56%
\$100-\$200	41%
\$200-\$350	3%
Transaction Price as % of List Price	% of Sample in Range
100%	28.6%
80 to less than 100%	3.0%
60 to less than 80%	6.1%
40 to less than 60%	27.8%
20 to less than 40%	26.8%
Less than 20%	7.7%
Share of Transactions within 0.5% of Three (or Fewer) Chronologically Decreasing Price Levels per Product	71.4%
Sample Size $N = 98$, with 7385 purchases observed in total	

The sample consists of the cardigans launched in calendar 2011.

Table 4 Customers' Online Views of Cardigans' Detailed Product Pages

	Mean Number of Views of the Product's Detailed Page	Mean Interval between Repeat Views (in days)
Online Customer	1.57	14.98
Customer Purchasing Online (5.96% of Online Customers)	2.34	9.28

Sample covers the 98 cardigans launched in calendar 2011. For uniformity, customer-cardigan pairs are restricted to 12,790 where the customer is brand-acquainted (i.e., viewed the cardigan at least once and interacted with the brand within the six months prior to and including the cardigan's launch). 762 (5.96%) of these led to purchase. The "Mean Interval between Repeat Views" derives from 3290 customer-product pairs (25.72%) with more than one view of the cardigan's detailed page.

Table 5 Panel Regression Results

Dependent Variable: Paid Amounts as Fraction of List Prices		
	Random Effects (IV)	Random Effects (IV)
<u>Controls:</u>		
Number of Promotions Applied:	Y	-
Number of Promotions Applied (IV):	-	Y
Month Dummies:	Y	Y
Customer Attributes:	Y	Y
List-Price Purchase Total by Channel (IV):	Y	Y
Distance to Store of Residence:	Y	Y
<u>Estimated Coefficient:</u>		
PVR (with IV)	0.195**** (0.049)	0.146** (0.061)
Hausman Test P-value	1	0.383
R^2	0.154	0.179
Significance levels → **** - 0.001	*** - 0.01	** - 0.05 * - 0.1

Unbalanced panel is for customer-month pairs with purchase observed in Jan. 2011 – Jun. 2012. Customer attributes include gender and age, each when known, and loyalty program status. IV regressions use as instruments the 5- and 6-month lags of the PVR and the list-price purchase total by channel, in addition to the 5- and 6-month lags of online visits, online visits with searches, and online visits with product views. For the promotions applied, we use lags 3-6. All lagged variables are customer-specific. The non-IV regressions use 28,745 observations (across 7905 customers), for which the PVR exists; the IV regressions use 14,544 observations (across 3316 customers), for which additionally the lagged PVR exist.

Table 6 Maximum Likelihood Estimates for the Structural Model

Customer Segments:	Bargain	High	P-value for (nonzero) difference across types:	Exogenous
	Hunters	Valuation		Demand
Elasticity $\hat{\alpha}$	0.026 (0.001)	0.010 (0.001)	†	0.021 (0.000)
Monitoring cost \hat{r}_1	0.000 (0.003)	0.301 —	—	
Monitoring cost \hat{r}_2	1.602 (0.003)	7.601 (0.001)	†	
Heterogeneous valuation $\hat{\gamma}$	0 —	0.872 (0.000)	†	-0.637 (0.002)
Mean other-purchase payoff \hat{w}	-2.666 (0.001)	-2.259 (0.001)	†	
Population Share	17.98% (0.35%)	82.02% (0.35%)	—	
Nuisance Parameters:				
<u>Correlated Random Effects for \hat{v}:</u>			<u>Depreciation:</u>	
Intercept	-5.192 (0.006)		Hazard rate $\hat{\lambda}_D$	0.010 (0.001)
Inventory	0.004 (0.000)		Realized disutility \hat{w}	-1.813 (0.003)
List Price	0.004 (0.000)		<u>Retailer's Discounting:</u>	
Seasonal (Jul.-Oct.)	0.031 (0.004)		Low hazard rate $\hat{\lambda}^{R-Low}$	0.018 (0.002)
Seasonal (Nov.)	0.433 (0.004)		High hazard rate $\hat{\lambda}^{R-High}$	0.021 (0.005)
Seasonal (Dec.)	0.281 (0.004)			
Variance	0.022 (0.006)			

N=98 cardigan products. We use † to denote a p-value less than 0.001.

Table 7 Customers' Opportunity Costs of Monitoring

	Bargain Hunters	High Valuation
Additional US dollar compensation per visit required to increase monitoring rate by one online visit per month $\frac{1}{30} \cdot \frac{\hat{r}_2}{\hat{\alpha}}$	\$2.02 (\$0.065)	\$25.35 (\$3.52)

Sample size $N = 98$ cardigan products.

Table 8 Empirical Distribution of Purchases, Visits, and Revenues by Customer Segment


	List Price	Sale Price	Clearance	% Revenues
<u>Bargain Hunters:</u>				18.0%
Purchases Share	16.2%	17.4%	21.1%	
Expected Visits per Customer				
Present at Season Start	29.93	15.72	14.77	
Visits Share	60.9%	61.3%	65.3%	
<u>High Valuation:</u>				82.0%
Purchases Share	83.8%	82.6%	78.9%	
Expected Visits per Customer				
Present at Season Start	4.21	2.18	1.72	
Visits Share	39.1%	38.8%	34.7%	

Sample size $N = 98$ cardigans. Expected visits per customer exclude post-purchase visits but are otherwise unadjusted for market exit (i.e., are not re-weighted over customers remaining) and depreciation. The posterior distribution of customers' type assignments conditional on the observed data and estimated model is used.

Table 9 Comparing the Pricing Policies' Expected Profits and Welfare (Inventory Jointly Optimized)

<u>Policy</u>	<u>Expected Profit</u>	<u>Optimal Inventory (Cost)</u>
Optimal Randomized Markdowns	\$71,251 100%	\$155,450 100%
Randomized Markdowns at Observed Prices	\$58,602 82.3%	\$150,870 97.1%
State-Contingent (Dynamic) Pricing	\$39,349 55.2%	\$66,640 42.9%
Consumer Surplus for:		
<u>Policy</u>	<u>High Valuation Customers</u>	<u>Bargain Hunters</u>
Optimal Randomized Markdowns	+\$28,686	+\$3,705
Randomized Markdowns at Observed Prices	+\$26,456	+\$3,808
State-Contingent (Dynamic) Pricing	Baseline	Baseline

The 98 cardigans launched in calendar 2011 are used. Policies optimize over inventory and pricing, taking into account demand from both customer types and the flow of one-time (including foot traffic) visitors. Inventory is shown as the sum of marginal costs for the quantities ordered. Expected profit and inventory figures are also given as percentages of their values under the optimal randomized markdowns.

Figure 10 Expected Profits under the Pricing Policies for Cardigan #33

The solid line plots the sum of marginal costs incurred at each inventory quantity. The dashed curves plot expected discounted revenue curves under each pricing policy. At any inventory quantity, the vertical distance between the two is the expected discounted profit. The square marks the optimal inventory and profit under the optimal randomized markdown policy. The circle corresponds to state-contingent, dynamic pricing and its optimal inventory choice.

Table 11 Posterior Variance of the Customer’s Type (1 for High Valuation Type) after Bayesian Update

	Using:		
	Full Data	PVR + Purchase Histories	Purchase Histories
Expected Variance	0.004	0.004	0.021
Empirical Distribution of Variance			
75% Quantile	0	0	0.010
90% Quantile	0	0	0.069
95% Quantile	0	0	0.157
99% Quantile	0.192	0.188	0.244


The sample consists of $N = 12,790$ customer-product pairs for 98 cardigans launched in calendar 2011. The distribution of variance is for a customer randomly selected from our sample of customer-cardigan pairs, with the Bayesian update performed after observing a season of behavior. The update takes our demand model as estimated. With purchase histories, updates use the time and market state of observed purchases, if any. When using the purchase-to-visit ratios (PVR), the updates additionally use each customer’s observed PVR. For the “Full Data” benchmark, the retailer uses the time and market states of all visits and purchases. Posterior distributions are obtained via MCMC, with each updating scheme constraining the support of the customers’ visit times. (E.g., a visit is implied at the time of each observed purchase. A realized PVR implies constraints on the number of visits made.) MCMC jointly samples the assignment of the market’s customers into types.

Table 12 Expected Profits and Welfare from Targeted Pricing Policies

Policy	Expected Profit	Optimal Inventory (Cost)
Optimal Randomized Markdowns with Targeted Prices	\$75,247 +5.6%	\$160,620 +3.3%
State-Contingent (Dynamic) Pricing with Targeting	\$41,523 +5.5%	\$69,962 +5.0%

Policy	Consumer Surplus for:	
	High Valuation Customers	Bargain Hunters
Optimal Randomized Markdowns with Targeted Prices	-\$9,289	+\$573
State-Contingent (Dynamic) Pricing with Targeting	-\$7,035	+\$521

The counterfactual uses the 98 cardigans launched in calendar 2011. Inventory and pricing are jointly optimal, taking into account both customer types and the flow of one-time (including foot traffic) visitors. Inventory is shown as the sum of marginal costs for the quantities ordered. For comparison against baseline policies, we provide the percentage changes from targeting relative to the baseline's expected profits and optimal inventories.

Figure 13 Expected Profits with Targeting for Cardigan #33

The expected discounted revenue curves of Fig. 10 are in solid; dashed lines indicate those for targeted prices. The new square marks the optimal inventory and profit under the optimal randomized markdown policy with targeting. The circle corresponds to state-contingent, dynamic pricing with targeting and its optimal inventory choice.

Table 14 Expected Profits under Counterfactual Higher Monitoring Costs by Segment

Policy	Higher Monitoring Costs for:	
	High Valuation Consumers	Bargain Hunters
Optimal Randomized Markdowns (Expected Profit)	\$69,046 −3.1%	\$70,533 −1.0%
(Optimal Inventory)	\$152,970 −1.6%	\$153,830 −1.0%
State-Contingent (Dynamic) Pricing (Expected Profit)	\$37,633 −4.4%	\$38,122 −3.1%
(Optimal Inventory)	\$64,482 −3.2%	\$64,077 −3.8%

The counterfactual uses the 98 cardigans launched in calendar 2011. Inventory and pricing are jointly optimized. Inventory is shown as the sum of marginal costs for the quantities ordered, with percentage-change comparisons against optima under existing monitoring costs. The bargain hunters counterfactual raises the bargain hunter's monitoring costs incurred at any chosen visit rate to equal those of a high valuation type in utility terms. The high valuation counterfactual raises the type's monitoring costs incurred at any visit rate by 10%.

References

- Aguirregabiria, Victor, Pedro Mira. 2007. Sequential estimation of dynamic discrete games. *Econometrica* **75**(1) 1–53.
- Akşin, Zeynep, Barış Ata, Seyed Morteza Emadi, Che-Lin Su. 2013. Structural estimation of callers' delay sensitivity in call centers. *Management Sci.* **59**(12) 2727–2746.
- Allon, Gad, Awi Federgruen, Margaret Pierson. 2011. How much is a reduction of your customers' wait worth? An empirical study of the fast-food drive-thru industry based on structural estimation methods. *Manufacturing Service Oper. Management* **13**(4) 489–507.
- Arcidiacono, Peter, Robert A Miller. 2011. Conditional choice probability estimation of dynamic discrete choice models with unobserved heterogeneity. *Econometrica* **79**(6) 1823–1867.
- Aviv, Yossi, Yuri Levin, Mikhail Nediak. 2009. Counteracting strategic consumer behavior in dynamic pricing systems. Serguei Netessine, Christopher S Tang, eds., *Consumer-Driven Demand and Operations Management Models*. Springer, 323–352.
- Aviv, Yossi, Amit Pazgal. 2008. Optimal pricing of seasonal products in the presence of forward-looking consumers. *Manufacturing Service Oper. Management* **10**(3) 339–359.
- Belobaba, Peter P. 1989. Application of a probabilistic decision model to airline seat inventory control. *Oper. Res.* **37**(2) 183–197.
- Besanko, David, Wayne L Winston. 1990. Optimal price skimming by a monopolist facing rational consumers. *Management Sci.* **36**(5) 555–567.

- Besbes, Omar, Ilan Lobel. 2015. Intertemporal price discrimination: Structure and computation of optimal policies. *Management Sci.* **61**(1) 92–110.
- Bray, Robert L, Haim Mendelson. 2012. Information transmission and the bullwhip effect: An empirical investigation. *Management Sci.* **58**(5) 860–875.
- Bray, Robert L, Haim Mendelson. 2015. Production smoothing and the bullwhip effect. *Manufacturing Service Oper. Management* **17**(2) 208–220.
- Bulow, Jeremy I. 1982. Durable-goods monopolists. *J. Political Econom.* **90**(2) 314–332.
- Cachon, Gérard P, Pnina Feldman. 2015. Price commitments with strategic customers: Why it can be optimal to discount more frequently . . . than optimal. *Manufacturing Service Oper. Management* **17**(3) 399–410.
- Cachon, Gérard P, Robert Swinney. 2009a. The impact of strategic consumer behavior on the value of operational flexibility. Serguei Netessine, Christopher S Tang, eds., *Consumer-Driven Demand and Operations Management Models*. Springer, 371–395.
- Cachon, Gérard P, Robert Swinney. 2009b. Purchasing, pricing, and quick response in the presence of strategic consumers. *Management Sci.* **55**(3) 497–511.
- Cachon, Gérard P, Robert Swinney. 2011. The value of fast fashion: Quick response, enhanced design, and strategic consumer behavior. *Management Sci.* **57**(4) 778–795.
- Caldentey, René, Ying Liu, Ilan Lobel. 2015. Intertemporal pricing under minimax regret. *SSRN Working Paper* (No. 2357083).
- Cameron, A Colin, Pravin K Trivedi. 2005. *Microeconometrics: Methods and Applications*. Cambridge University Press.
- Caro, Felipe, Jérémie Gallien. 2010. Inventory management of a fast-fashion retail network. *Oper. Res.* **58**(2) 257–273.
- Chen, Yiwei, Vivek F Farias. 2015. Robust dynamic pricing with strategic customers. *Proc. Sixteenth ACM Conf. Econom. Computation*. ACM, 777–777.
- Chevalier, Judith, Austan Goolsbee. 2009. Are durable goods consumers forward-looking? Evidence from college textbooks. *Quarterly J. Econom.* **124**(4) 1853–1884.
- Chib, Siddhartha. 2001. Markov chain Monte Carlo methods: Computation and inference. James Heckman, Edward Leamer, eds., *Handbook of Econometrics*, vol. 5. North Holland, 3569–3649.
- Coase, Ronald H. 1972. Durability and monopoly. *J. Law & Econom.* **15** 143.
- Cohen, Maxime C, Ngai-Hang Z Leung, Kiran Pancharangam, Georgia Perakis, Anthony Smith. 2014. The impact of linear optimization on promotion planning. *SSRN Working Paper* (No. 2382251).
- Correa, José, Ricardo Montoya, Charles Thraves. 2016. Contingent preannounced pricing policies with strategic consumers. *Oper. Res.* **64**(1) 251–272.

- Dasu, Sriram, Chunyang Tong. 2010. Dynamic pricing when consumers are strategic: Analysis of posted and contingent pricing schemes. *Eur. J. Oper. Res.* **204**(3) 662–671.
- Dempster, Arthur P, Nan M Laird, Donald B Rubin. 1977. Maximum likelihood from incomplete data via the EM algorithm. *J. Royal Statist. Soc. Series B (Methodological)* **39**(1) 1–38.
- Dobzinski, Shahar, Shaddin Dughmi. 2013. On the power of randomization in algorithmic mechanism design. *SIAM J. Computing* **42**(6) 2287–2304.
- Doraszelski, Ulrich, Kenneth L Judd. 2012. Avoiding the curse of dimensionality in dynamic stochastic games. *Quantitative Econom.* **3**(1) 53–93.
- Dubé, Jean-Pierre, Jeremy T Fox, Che-Lin Su. 2012. Improving the numerical performance of static and dynamic aggregate discrete choice random coefficients demand estimation. *Econometrica* **80**(5) 2231–2267.
- Dutta, Prajit K, Rangarajan K Sundaram. 1998. The equilibrium existence problem in general Markovian games. *Organizations with Incomplete Information*, chap. 5. Cambridge University Press, 159–207.
- Echenique, Federico. 2000. Comparative statics by adaptive dynamics and the correspondence principle. *UC Berkeley Department of Economics Working Paper* (No. E00-273).
- Echenique, Federico. 2002. Comparative statics by adaptive dynamics and the correspondence principle. *Econometrica* **70**(2) 833–844.
- Ellison, Glenn, Sara Fisher Ellison. 2009. Search, obfuscation, and price elasticities on the Internet. *Econometrica* **77**(2) 427–452.
- Elmaghraby, Wedad, Altan Gülcü, Pınar Keskinocak. 2008. Designing optimal preannounced markdowns in the presence of rational customers with multiunit demands. *Manufacturing Service Oper. Management* **10**(1) 126–148.
- Elmaghraby, Wedad, Pınar Keskinocak. 2003. Dynamic pricing in the presence of inventory considerations: Research overview, current practices, and future directions. *Management Sci.* **49**(10) 1287–1309.
- Etzion, Hila, Edieal Pinker, Abraham Seidmann. 2006. Analyzing the simultaneous use of auctions and posted prices for online selling. *Manufacturing Service Oper. Management* **8**(1) 68–91.
- Fershtman, Chaim, Ariel Pakes. 2012. Dynamic games with asymmetric information: A framework for empirical work. *Quarterly J. Econom.* **127**(4) 1611–1661.
- Gallego, Guillermo, Garrett Van Ryzin. 1994. Optimal dynamic pricing of inventories with stochastic demand over finite horizons. *Management Sci.* **40**(8) 999–1020.
- Gorodnichenko, Yuriy, Viacheslav Sheremirov, Oleksandr Talavera. 2015. Price setting in online markets: Does it click? *Federal Reserve Bank of Boston Working Paper* (No. 15-1).
- Gourieroux, Christian, Alain Monfort. 1997. *Simulation-Based Econometric Methods*. Oxford University Press.

-
- Gul, Faruk, Hugo Sonnenschein, Robert Wilson. 1986. Foundations of dynamic monopoly and the Coase conjecture. *J. Econom. Theory* **39**(1) 155–190.
- Hendel, Igal, Aviv Nevo. 2013. Intertemporal price discrimination in storable goods markets. *Amer. Econom. Rev.* **103**(7) 2722–51.
- Honka, Elisabeth. 2014. Quantifying search and switching costs in the US auto insurance industry. *RAND J. Econom.* **45**(4) 847–884.
- Hörner, Johannes, Larry Samuelson. 2011. Managing strategic buyers. *J. Political Econom.* **119**(3) 379–425.
- Hortaçsu, Ali, Chad Syverson. 2004. Search costs, product differentiation, and welfare effects of entry: A case study of S&P 500 index funds. *Quarterly J. Econom.* **119**(4) 403–456.
- Jerath, Kinshuk, Serguei Netessine, Senthil K Veeraraghavan. 2010. Revenue management with strategic customers: Last-minute selling and opaque selling. *Management Sci.* **56**(3) 430–448.
- Kim, Sang Won, Marcelo Olivares, Gabriel Y Weintraub. 2014. Measuring the performance of large-scale combinatorial auctions: A structural estimation approach. *Management Sci.* **60**(5) 1180–1201.
- Kumar, Sumit. 2005. Supply chain strategies in the apparel industry: The case of Victoria’s Secret. Ph.D. thesis, Massachusetts Institute of Technology.
- Lazear, Edward P. 1986. Retail pricing and clearance sales. *Amer. Econom. Rev.* **76**(1) 14–32.
- Li, Jun, Nelson Granados, Serguei Netessine. 2014. Are consumers strategic? Structural estimation from the air-travel industry. *Management Sci.* **60**(9) 2114–2137.
- Liu, Qian, Garrett J van Ryzin. 2008. Strategic capacity rationing to induce early purchases. *Management Sci.* **54**(6) 1115–1131.
- Louis, Thomas A. 1982. Finding the observed information matrix when using the EM algorithm. *J. Royal Statist. Soc. Series B (Methodological)* **44**(2) 226–233.
- Maitra, Ashok. 1968. Discounted dynamic programming on compact metric spaces. *Sankhyā: Indian J. Statist., Series A* **30**(2) 211–216.
- Maskin, Eric, Jean Tirole. 2001. Markov perfect equilibrium: I. observable actions. *J. Econom. Theory* **100**(2) 191–219.
- Mattioli, Dana. 2012. Lululemon’s secret sauce — scarcity built a \$10.4 billion retailer — but will it prevent Lulu from growing? *The Wall Street Journal* (Mar. 22).
- McAfee, R Preston, Thomas Wiseman. 2008. Capacity choice counters the Coase conjecture. *Rev. Econom. Stud.* **75**(1) 317–331.
- McFadden, Daniel. 1989. A method of simulated moments for estimation of discrete response models without numerical integration. *Econometrica* **57**(5) 995–1026.
- Milgrom, Paul, John Roberts. 1990. Rationalizability, learning, and equilibrium in games with strategic complementarities. *Econometrica* **58**(6) 1255–1277.

- Milgrom, Paul, Chris Shannon. 1994. Monotone comparative statics. *Econometrica* **62**(1) 157–180.
- Moraga-González, José Luis, Matthijs R Wildenbeest. 2008. Maximum likelihood estimation of search costs. *Eur. Econom. Rev.* **52**(5) 820–848.
- Nair, Harikesh. 2007. Intertemporal price discrimination with forward-looking consumers: Application to the US market for console video-games. *Quantitative Marketing Econom.* **5**(3) 239–292.
- Netessine, Serguei, Christopher S Tang. 2009. *Consumer-Driven Demand and Operations Management Models*. Springer.
- Olivares, Marcelo, Christian Terwiesch, Lydia Cassorla. 2008. Structural estimation of the newsvendor model: An application to reserving operating room time. *Management Sci.* **54**(1) 41–55.
- Rossi, Peter E, Robert E McCulloch, Greg M Allenby. 1996. The value of purchase history data in target marketing. *Marketing Sci.* **15**(4) 321–340.
- Rust, John. 1994. Structural estimation of Markov decision processes. Robert Engle, Daniel McFadden, eds., *Handbook of Econometrics*, vol. 4, chap. 51. North Holland, 3081–3143.
- Seiler, Stephan. 2013. The impact of search costs on consumer behavior: A dynamic approach. *Quantitative Marketing Econom.* **11**(2) 155–203.
- Shen, Zuo-Jun Max, Xuanming Su. 2007. Customer behavior modeling in revenue management and auctions: A review and new research opportunities. *Production Oper. Management* **16**(6) 713–728.
- Sobel, Joel. 1984. The timing of sales. *Rev. Econom. Stud.* **51**(3) 353–368.
- Soysal, Gonca P, Lakshman Krishnamurthi. 2012. Demand dynamics in the seasonal goods industry: An empirical analysis. *Marketing Sci.* **31**(2) 293–316.
- Stigler, George J. 1961. The economics of information. *J. Political Econom.* **69**(3) 213–225.
- Stokey, Nancy L. 1979. Intertemporal price discrimination. *Quarterly J. Econom.* **93**(3) 355–371.
- Su, Che-Lin, Kenneth L Judd. 2012. Constrained optimization approaches to estimation of structural models. *Econometrica* **80**(5) 2213–2230.
- Su, Xuanming. 2007. Intertemporal pricing with strategic customer behavior. *Management Sci.* **53**(5) 726–741.
- Su, Xuanming, Fuqiang Zhang. 2009. On the value of commitment and availability guarantees when selling to strategic consumers. *Management Sci.* **55**(5) 713–726.
- Swinney, Robert. 2012. Inventory pooling with strategic consumers: Operational and behavioral benefits. *Working Paper* .
- Talluri, Kalyan T, Garrett J Van Ryzin. 2006. *The Theory and Practice of Revenue Management*, vol. 68. Springer.
- Taylor, Terry A. 2006. Sale timing in a supply chain: When to sell to the retailer. *Manufacturing Service Oper. Management* **8**(1) 23–42.

Varian, Hal R. 1980. A model of sales. *Amer. Econom. Rev.* **70**(4) 651–659.

Villas-Boas, J Miguel, Russell S Winer. 1999. Endogeneity in brand choice models. *Management Sci.* **45**(10) 1324–1338.

White, Halbert. 1982. Maximum likelihood estimation of misspecified models. *Econometrica* **50**(1) 1–25.

Electronic Companion

Appendix A: Customer Monitoring and Purchasing Behavior

A.1. Customers' Monitoring and Discounts Obtained

We supply additional descriptive and panel data results on how customers monitor and obtain discounts.

In Figure EC.1, we summarize customers' discounts obtained from the brand, by presenting the cross-sectional distribution of customers' lifetime and lifetime-average post-discount yields. (The first is a customer's lifetime amount paid to the brand over the sum of her purchases' list prices. The latter is her lifetime average, over her individual purchases, of the amount she paid for the purchase over its list price.) Notably, we observe a relatively more dispersed bi-modal distribution across customers in their lifetime post-discount yields, suggesting that customers who tend to wait for markdowns do so more reliably for their more expensive (by list price) purchases. Such behavior would be congruent with monitoring customers' incentives.

Figure EC.2 highlights that customers who obtain discounts more frequently tend to monitor more intensively. Each subfigure of Figure EC.2 illustrates a customer's monitoring behavior conditional on the percentage of her transactions made at a discount (of at least 10% from the list price). Figure EC.2a makes clear that the customers who frequently obtain discounts exhibit relatively higher session volumes. Figure EC.2b affirms this relationship at the monthly (instead of lifetime) level, demonstrating that these customers do in fact monitor more intensively rather than simply having been with the brand longer (hence recording more sessions over a longer time). The conspicuous dip in monitoring for customers with exactly 100% of their transactions made at a discount reflects the volume of one-time (or otherwise very low-incidence) visitors who make impulse purchases of discounted items while passing through; notably, conditioning on customers having made a minimum number of transactions with the brand largely dissipates this effect. Lastly, Figure EC.2c depicts this relationship in terms of the customers' observed purchase-to-visit ratios (PVR).

Additional panel results support our primary findings. Pooled OLS results in Table EC.3 support our random-effects panel findings, consistent with the assumption of random effects. As discussed in Section 3, employing lags as instruments selects for a relatively experienced subpopulation of customers who have been visiting the brand over an extended period of time (beyond a season). Arguably, we are most interested in the behavior of these repeat customers who are most likely to time their purchases. We nonetheless confirm that the difference across our IV and OLS results is not driven by sample selection, by presenting OLS without instruments for exactly the subpopulation for which lagged instruments are available (Table EC.4). We find no indication that selection drives the IV results, with our results in fact suggesting that any bias runs in the opposite direction.

The pooled OLS results in Table EC.3 are presented both with and without using our instrumental variables (IV). Similarly contrasting estimates are obtained without instruments for the random effects panel results shown in Table 5, evidencing in both cases a sizeable endogeneity bias. For a customer in the panel, such biases arise because we relate the customer's post-discount yield, which measures her discounts obtained on her purchases made, to her concurrent PVR in the context of a single month as the unit of observation. In contrast, our structural model identifies engaged customers and tracks their behavior and decision process

over a product’s full season. For an observation in our panel setting, whether a customer is interested in the products available (and possibly marked down that month) affects both her monitoring and purchasing behavior. Moreover, PVR is mechanically biased upwards, to extents varying for the customer types, during the months of actual purchasing rather than of primarily waiting.

Our IV approach remedies this issue by predicting the customer’s PVR from her behavior well prior to a product’s selling season. Because this is prior to knowing whether the product even appeals to the customer, this prediction better indicates her typical monitoring behavior and does not suffer such biases. Put differently, it accomplishes our aim of relating her typical monitoring behavior, captured by her predicted rather than actual PVR, to her purchasing outcomes.

A.2. Product Substitution

We provide background from the data, then describe our statistical analysis and its finding that the sample’s cardigans are not strong substitutes for one another. Next, we discuss how substitution, if present, may affect our main results. Importantly, we may underestimate the heterogeneity in customers’ price elasticities, which would lead us to understate the true value of PVR and targeting. Lastly, in the presence of substitution, randomizing markdown times may in fact make markdown pricing more robust to competition.

From the data, we provide additional background on how customers monitor and purchase multiple (distinct) cardigans from the brand. For visits in which a customer does view at least one cardigan’s detail page, on average the customer views 1.06 cardigans (and 1.48 items in total from the browsing category), indicating that it is relatively uncommon for multiple cardigans to be viewed concurrently in detail. In over 80% of such visits, only a single cardigan’s detail page is viewed. Online active customers purchase 1.7 cardigans on average over the data’s full two-year sample period, subject to an upward skew. 73% of these customers purchase either one or zero cardigans. Less than 8% purchase more than five. Customers typically monitor somewhat less frequently following a purchase: descriptively, with no adjustment made to control for customers’ types or number of observations by type, the median (75th quantile) inter-visit interval rises from 6 to 8 (19 to 30) days following any online purchase from the brand.

As discussed in Section 4.1, three observations shape how we consider product substitution. First, with detailed, continuous-time customer-level data, the customer’s purchase-or-wait decisions are made in the context of a single visit, where her decision to make other purchases is not a reliable indicator of substitution, as she may still purchase in a future visit within even a day or less. Second, browsing behavior is similarly rationalizable with the customer monitoring multiple products. Both of these observations arise because a customer may purchase (single) units of multiple products in our setting, as opposed to a setting where each customer has single-unit demand in sum over all available products. Lastly, because products are introduced continually and the retailer does not engage in seasonal discounts (which receive careful attention in prior work), we focus instead on whether customers tend to find certain products to be closer substitutes than others, justifying special attention to the prices and markdowns of these products.

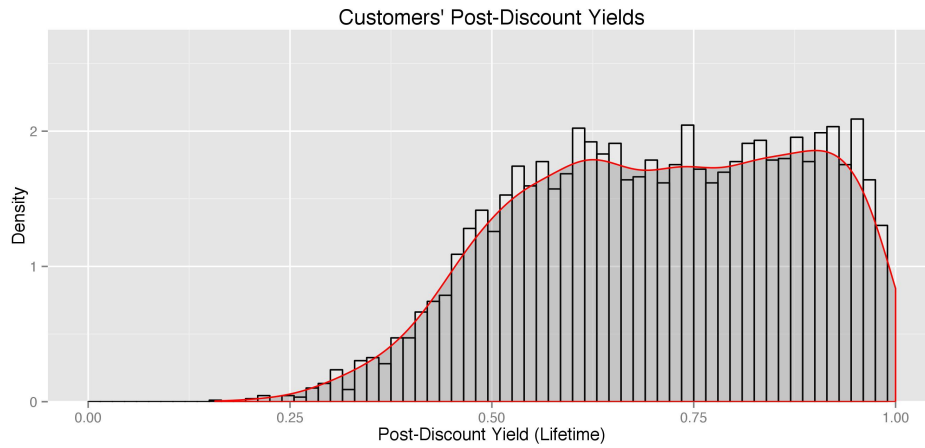
We address the possibility of substitution in two ways. First, we test directly for whether a customer exits the cardigan’s market upon purchasing another cardigan instead. We design the test to control for the customer’s overall number of sweater (i.e., browsing category) purchases, which importantly accounts for

the cross-sectional heterogeneity in customers' purchasing volumes. Second, we discuss how the presence of substitution may change or qualify our results.

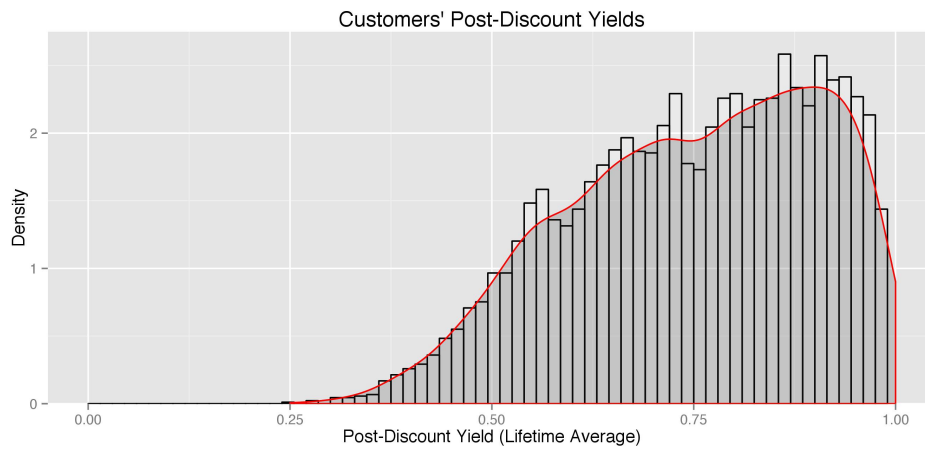
Accordingly, we first consider whether purchasing another cardigan reduces a customer's subsequent probability of buying the cardigan of interest, relative to having purchased another item in the same online browsing category (hence involving, and allowing us to control for, similar browsing and monitoring exposures). If cardigans are indeed closer substitutes, we would expect to see such a reduction in future buying. We select from our sample the three cardigans of highest transaction volume without any overlap in their seasons. Controlling for each customer's number of sweater (including cardigan) purchases in the first month post-launch, hence her overall volume of purchasing behavior, we see whether observing more cardigans among her purchased sweaters translates into a lower probability of her purchasing the cardigan of interest in the future. As presented in Table EC.5, this test does not show any evidence of market exit by substitution among cardigans.

Finally, we discuss how the presence of substitution could change or qualify our findings and results. First, we would underestimate the price elasticity of bargain hunters, who tend to buy later in a cardigan's season. Over the course of a selling season, customers would encounter opportunities to substitute away from the good. Because (A) our model would not account for the customers who did in fact exit by substitution, and (B) having exited, such customers would not respond to the good's markdowns by purchasing, we would consequently underestimate customers' price elasticities. While this observation applies to all customers, our model allows for the cardigan's utility value to depreciate, and how the typical customer substitutes away would be treated as a source of such depreciation. However, because bargain hunters are a smaller subgroup that is both relatively more likely to substitute away for price and more active toward the cardigan's season's end (following a greater number of opportunities to substitute), our estimates would still underestimate how sensitive they are to price. In sum, we would conservatively understate the true degree of customer heterogeneity that explains the behavior observed in our data. Importantly, we then expect our findings to understate the true value of targeting promotions using behavioral metrics such as PVR.

Second, while substitution raises the issue of competing retailers, randomizing markdown times may effectively make markdown pricing more robust to competition. In the literature on temporary sales with competition, Varian (1980) and Sobel (1984) highlight that that randomization (in the form of mixed strategies) is the key ingredient to allow for competitive equilibria in which no retailer finds it profitable to undercut its competition by becoming an everyday low price retailer (and becoming known as such to consumers in the presence of search costs). As derived in these papers, these equilibria are symmetric, so that they are sustainable even as customers associate sales patterns with brands, and each retailer profits enough from maintaining its pricing uncertainty that it would not benefit from deterministically undercutting its rivals using everyday low pricing. For the setting of temporary sales, they also show that this robustness is absent in equilibria featuring deterministic pure strategies. We predict that randomizing markdown times similarly facilitate markdown policies for competing retailers but leave a detailed analysis to future research.

Figure EC.1 Customers' Post-Discount Yields

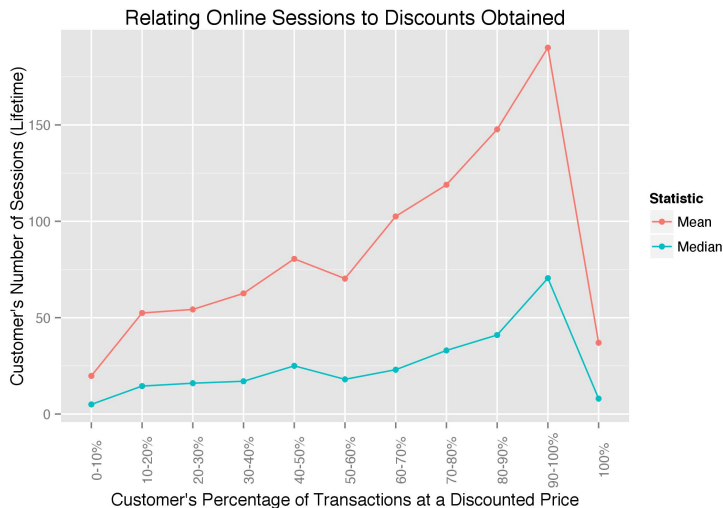
(a) Lifetime Post-Discount Yield



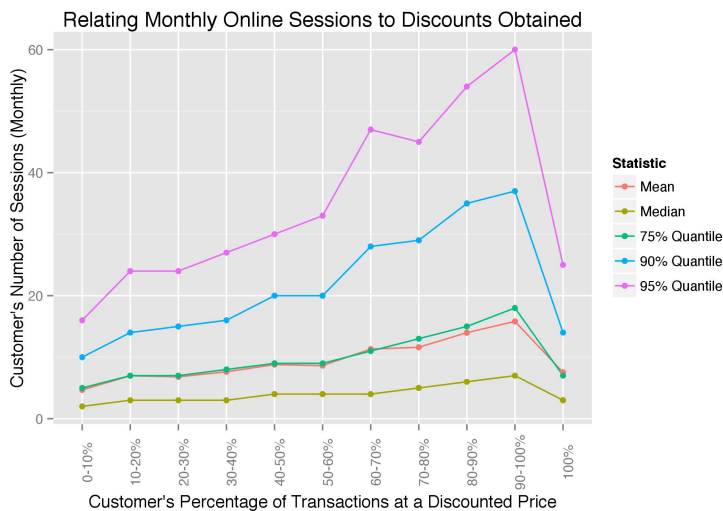
(b) Lifetime-Average Post-Discount Yield

Histogram and kernel density plots present the lifetime post-discount yield for 6,055 customers who visit online at least once and make at least 10 purchases during the sample period (and with post-discount yield less than one). A customer's lifetime post-discount yield equals the total amount she paid over the list-price sum of her purchases, and her lifetime-average post-discount yield averages her post-discount yields from each individual purchase she made.

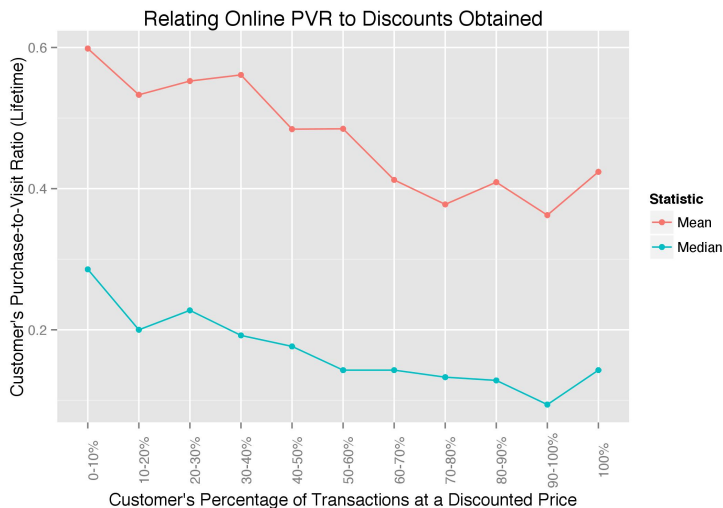
Figure EC.2 Relating Customers' Online Visits to Discounts Obtained (Jul. 2010 – Jun. 2012)



(a) 9840 Customers with Positive Visits and Transactions Observed



(b) 79,104 Customer-Month Pairs with Visit Observed for 9840 Customers with Purchase in Sample



(c) 9840 Customers with Positive Visits and Transactions Observed

Table EC.3 Pooled OLS Regression Results

<u>Dependent Variable:</u>					
Paid Amounts as Fraction of List Prices					
<u>Controls:</u>					
Number of Promotions Applied					
Observed:	Y	Y	Y	Y	-
IV:	-	-	-	-	Y
Month Dummies:	-	Y	Y	Y	Y
Customer Attributes:	-	-	Y	Y	Y
List-Price Purchase Total by Channel (IV):					
Distance to Store of Residence:	-	-	-	Y	Y
<u>Estimated Coefficient:</u>					
PVR (No IV)	0.021****	0.020****	0.015****	0.015****	-
	(0.002)	(0.002)	(0.002)	(0.002)	
PVR (with IV)	0.166***	0.169****	0.149****	0.150****	0.141****
	(0.0133)	(0.013)	(0.015)	(0.015)	(0.022)
R^2 (No IV)	0.024	0.037	0.053	0.053	-
Significance levels →	**** - 0.001	*** - 0.01	** - 0.05	* - 0.1	

Unbalanced panel is for customer-month pairs with purchase observed in Jan. 2011 – Jun. 2012. Customer attributes include gender and age, each when known, and loyalty program status. Refer to Table 5 for a description of the instruments used. The non-IV regressions use 28,745 observations (across 7905 customers), for which the PVR exists; the IV regressions use 14,544 observations (across 3316 customers), for which additionally the lagged PVR exist.

Table EC.4 Pooled OLS Regression Results (No IV) Restricted to IV Subpopulation

<u>Dependent Variable:</u>					
Paid Amounts as Fraction of List Prices					
<u>Controls:</u>					
Number of Promotions Applied:					
		Y	Y	Y	Y
Month Dummies:		-	Y	Y	Y
Customer Attributes:		-	-	Y	Y
List-Price Purchase Total by Channel (IV):					
Distance to Store of Residence:		-	-	-	Y
<u>Estimated Coefficient:</u>					
PVR		0.017****	0.017****	0.011***	0.011***
		(0.003)	(0.003)	(0.004)	(0.004)
R^2		0.020	0.035	0.053	0.053
Significance levels →		**** - 0.001	*** - 0.01	** - 0.05	* - 0.1

Unbalanced panel is for customer-month pairs with purchase observed in Jan. 2011 – Jun. 2012. Customer attributes include gender and age, each when known, and loyalty program status.

Table EC.5 Probit Regressions to Test Cardigan Substitution

Dependent Variable: Customer Purchases Cardigan Subsequent to First Month						
	Estimate	P-value	Estimate	P-value	Estimate	P-value
<u>First Month:</u>						
Number of Cardigans Purchased	0.03 (0.12)	0.82	0.09 (0.14)	0.53	-0.10 (0.20)	0.63
Number of Sweaters Purchased	0.33** (0.15)	0.03	0.36 (0.23)	0.12	0.46**** (0.14)	0.00
Customer Attributes (Excluding Age)			Y		Y	
Customer Age					Y	
Sample Size	83		83		60	
Conditional on Sweater Purchase						
	Estimate	P-value	Estimate	P-value	Estimate	P-value
<u>First Month:</u>						
Number of Cardigans Purchased	0.05 (0.16)	0.77	0.07 (0.20)	0.70	-0.06 (0.30)	0.84
Number of Sweaters Purchased	0.16**** (0.04)	0.00	0.15** (0.08)	0.04	0.29**** (0.09)	0.00
Customer Attributes (Excluding Age)			Y		Y	
Customer Age					Y	
Sample Size	270		270		188	
Conditional on Purchase of Top						

Significance levels → **** - 0.001 *** - 0.01 ** - 0.05 * - 0.1

Samples are restricted to repeat customers who do not purchase the cardigan of interest during its first month post-launch. Three cardigans are used from our sample, launched in Jan. 2011, July 2011, and Jan. 2012, respectively, with the latter two dates corresponding to the cardigans with the largest number of purchases in our sample. The Jan. 2011 cardigan has the most recorded purchases of cardigans launched that month. The seasons for these cardigans do not overlap. Product-specific intercepts are included in all regressions, with robust standard errors for clustering by product. (Not adjusting for clustering does not result in qualitative changes for the p-values.) “Sweaters” refer to the products in the common browsing category as cardigans and include cardigans. Estimates and p-values for the cardigans coefficient remain similar when using the number of first-month purchases from the women’s tops category instead of sweaters only. Customer attributes are age (if known), loyalty enrollment, gender, time since first purchase with brand (zero if yet to occur by the relevant cardigan’s launch), and distance to nearest store. For each cardigan, the repeat customers are those with a recorded purchase or visit with the brand in the six months prior to and including the launch date and who view the cardigan’s detail page at least once. Results are also similar, with typically higher p-values for the coefficient on cardigans purchased, when using the full sample of 751 repeat customers for the three cardigans without conditioning on a first-month purchase taking place.

Appendix B: Markov Stationary Equilibrium

In this subsection we summarize our structural model and define our Markov stationary equilibrium concept.

Products: Let $j \in \mathbb{J}$ denote a product in the category of interest. The setting for each product's market is defined by:

- (i) The product's selling horizon $\mathbb{T} := [0, \infty)$.
- (ii) The initial inventory. We define the inventory $I_t, t \in \mathbb{T}$ as the quantity $I_0 > 0$ at time $t = 0$ and then reduced incrementally upon customers' purchases, without replenishment or backlogging. Remaining inventory is available to all customers.
- (iii) The set of potential price levels which are pre-determined and known to all players:

$$\{p^{\text{List}} = p^0, p^1, \dots, p^L = p^{\text{Min}} \in \mathbb{R}^{L+1} : p^k > p^{k+1} \text{ for } k = 0, \dots, L-1\} \text{ with } L \text{ finite.}$$

- (iv) The initial population of customers for the product. N customers are drawn i.i.d. from a heterogeneous, finite-type population (common to all products) consisting of the two customer types $\theta \in \{1, 2\}$, each with the associated parameter values $\vec{\theta}_k = \{\gamma^{(k)}, \alpha^{(k)}, w^{(k)}, r_1^{(k)}, r_2^{(k)}\}$ for $k = 1, 2$, respectively. For a customer i , we may also denote her relevant parameter values as γ_i, α_i, w_i , and \vec{r}_i . A customer is of type $\theta = 1$ with probability π , and her type persists throughout the season. At each time $t \in \mathbb{T}$, we denote the number of customers of each type remaining in the product market by N_t^1 and N_t^2 , respectively.

- (v) The product's base (utility) valuation $v_j \in \mathbb{R}$. Given the product's characteristics, Z_j , and the correlation-parameter vector, β , we assume that $v_j = Z_j \beta + \eta_j$, where, across the product category \mathbb{J} , η_j are drawn independently from the normal distribution with mean zero and standard deviation parameter σ_{CRE} . In our application, Z_j is defined to include the product's list price p^{List} , its initial inventory I_0 , and seasonal dummy variables. We motivate this correlated random effects specification in Section 5.

Next, the customers and the retailer are modeled as follows.

Customers: A customer's action space, $A_t^{(i)}$, is given by the product space of:

- (i) The customer's endogenous visit hazard rate at time t , $\lambda_t^{(i)} \in [0, M]$; and
- (ii) The customer's discrete choice upon the event of a visit at time t , $c_t^{(i)} \in C_t^{(i)}$, where her choice set $C_t^{(i)}$ depends on whether she has exited the market prior to t :

$$C_t^{(i)} = \begin{cases} \{\text{Product Purchase, Other Purchase, No Purchase}\} & \text{if no market exit prior to } t \\ \{\text{Other Purchase, No Purchase}\} & \text{if market exit prior to } t \end{cases}.$$

Accordingly, $a_t^{(i)} := (\lambda_t^{(i)}, c_t^{(i)})$ denotes the action of customer i at time t . The customer exits the market at the time of a purchase, i.e., at the time t when she visits and $C_t^{(i)} = \text{Product Purchase}$.

The retailer: The retailer's action, the discount hazard rate λ_t^R is chosen from the action space A_t^R , given by the state-dependent, discrete set:

$$A_t^R = \begin{cases} \{\lambda^{R-High}, \lambda^{R-Low}\} & \text{if } p_t > p^{\text{Min}} \text{ and } I_t > 0 \\ \{0\} & \text{otherwise} \end{cases},$$

where $\lambda^{R-High}, \lambda^{R-Low} \in \mathbb{R}_+$, $\lambda^{R-High} \geq \lambda^{R-Low}$ are the exogenously fixed, high and low discount hazard rates.

Finally, we define the dynamic, stochastic game between the customers and the retailer as follows:

The game: The state of the product market at time t is defined as

$$X_t := \{I_t, N_t^1, N_t^2, p_t, D_t\}, \quad \forall t \in \mathbb{T},$$

where D_t is an indicator variable for whether the product has depreciated in utility value as of time t . Let \mathbb{X} denote the state space for market states. At each time $t \in \mathbb{T}$, the history of the product market states up to time t , $h_t^X := \{X_r : r \in [0, t)\}$, is assumed to be common knowledge for all customers and the retailer. A play of the dynamic, stochastic game for product j 's market is defined as a measurable function

$$h : \mathbb{T} \rightarrow \mathbb{X} \times \mathbb{R}^{3 \cdot N} \times \prod_{i \in N} A_t \times A_t^R,$$

which we express as $h(t) := (X_t, \epsilon_t, a_t)$. Expanding beyond just the history of market states h_t^X , we also let $h_t^{(i)} := \{X_r, \epsilon_s^{(i)} : r \in [0, t), s \in [0, t]\}$ denote the private state history for customer i , relevant during her visit at time t .

Finally, we let H denote the space of all plays and H^{-i} , H^{-R} its subspaces excluding the actions of customer i and the retailer, respectively. Then, a strategy for player i is given by the measurable function $s_i : \mathbb{T} \times H^{-i} \rightarrow A_t^{(i)}$, with the informational restrictions that:

- (1) For every $h, h' \in H^{-i}$ and $t \in \mathbb{T}$ such that $h_t^X = h'_t{}^X$, then the prescribed $\lambda_t^{(i)}$ is equal with probability 1 under each of $s_i(t, h)$ and $s_i(t, h')$; and
- (2) For every $h, h' \in H^{-i}$ and $t \in \mathbb{T}$ such that $h_t^{(i)} = h'_t{}^{(i)}$, then the prescribed $c_t^{(i)}$ is equal with probability 1 under each of $s_i(t, h)$ and $s_i(t, h')$.

A strategy for the retailer is defined analogously. We use S_i to denote the space of strategies for player i , with $S := S_R \times \prod_{i \in N} S_i$. Finally, note that in defining S we restrict our attention to the set of pure strategies. Each strategy profile, $s \in S$, designates a well-defined probability distribution over plays.

Players' utilities: Each player has an instantaneous discount rate, ρ . For customer i , her ρ -discounted payoff from play $h \in H$ is given by the following discounted sum of infinite-horizon utility flows and shocks:

$$- \int_0^\infty e^{-\rho \cdot t} \cdot \delta(\lambda_t^{(i)}) dt + \sum_{m=1}^\infty e^{-\rho \cdot T_m^{(i)}} \cdot \Psi_i \left(X_{T_m^{(i)}-}, c_t^{(i)} \right),$$

where $T_m^{(i)} \in \mathbb{T}$ is the random time of customer i 's m -th visit, and her cost function δ is positive, increasing, and convex. Here, $X_{\tau-}$ denotes the left limit of the sample path of state process X at $\tau \in \mathbb{T}$. The instantaneous utility payoff Ψ_i associated with customer i 's m -th visit, occurring at random time $T_m^{(i)}$, is given by:

$$\Psi_i \left(X_{T_m^{(i)}-}, c_{T_m^{(i)}}^{(i)} \right) = \begin{cases} \epsilon_{T_m^{(i)}}^{(i),0} & \text{if } c_{T_m^{(i)}}^{(i)} = \text{No Purchase} \\ \gamma_i + v_j - \omega \cdot D_{T_m^{(i)}} - \alpha_i \cdot p_{jT_m^{(i)}-} + \epsilon_{T_m^{(i)}}^{(i),1} & \text{if } c_{T_m^{(i)}}^{(i)} = \text{Product Purchase} \\ w_i + \epsilon_{T_m^{(i)}}^{(i),2} & \text{if } c_{T_m^{(i)}}^{(i)} = \text{Other Purchase} \end{cases}$$

Note that the expected-payoff horizon is infinite, which accounts for the customer's behavior post-exit (occurring by her purchase or the cardigan's stock out), where she visits only to monitor and make purchases not including the cardigan. This is also reflected in how we specify her post-exit, in-visit discrete choice (see item (ii) under the Paragraph "Customers" above), to be limited to her options for other purchase or no purchase, as the cardigan is no longer available.

For the retailer, the ρ -discounted payoff from play $h \in H$ is given by the following discounted product revenue:

$$\sum_{i \in N} \sum_{m=1}^{\infty} e^{-\rho \cdot T_m^{(i)}} \cdot p_{j T_m^{(i)}} \cdot \mathbf{1}_{\left\{c_{T_m^{(i)}}^{(i)} = \text{Product Purchase}\right\}}.$$

All players seek to maximize their expected sums of ρ -discounted payoffs, where the expectation is taken over the distribution of plays subject to the strategy profile. In states X_t for which the value of product j has not yet depreciated, i.e., $D_t = 0$, a constant hazard rate of depreciation, λ_D , applies.

One-time, price-elastic demand: We specify exogenous demand (from one-time foot and web visitors) as a hazard rate of product purchases, given by the product of the (observed) seasonal arrival rate of such customers and their price-dependent, per-arrival purchase probability specified as logit, with the mean purchase option payoff $v_j + \gamma_{\text{Outside}} - \omega \cdot D_t - \alpha_{\text{Outside}} \cdot p_{jt}$ and no purchase normalized to zero. As mentioned earlier, we also include newly arriving online visitors who first encounter the brand mid-season after a cardigan’s launch, since their behavior differs substantially from the engaged body of existing customers (including whether and how much they continue to interact).

Post-exit states: To simplify exposition, we largely neglect to detail post-exit states and behavior. Once a customer exits the market for a cardigan, either by her purchasing it or by stock out, her in-visit choice set is restricted as described above in the Paragraph “Customers” and no longer depends on the currently observed state. In plain words, she visits only to monitor and make other purchases. Consequently, it is very simple to obtain the customer’s value function and stationary policy (decoupled from the strategic behavior of others) which are uniform across all post-exit states. At various points, we use V_0 to denote this value function post-exit.

Markov stationary equilibrium: We define a Markov stationary equilibrium for the game as a strategy profile in S : (i) that is a fixed point of the best-response correspondence (informally, each strategy in the profile best responds to the strategy profile); and (ii) where each strategy is Markov stationary on the state space \mathbb{X} . The best-response correspondence is defined as standard for Markov perfect equilibria, for players maximizing their expected ρ -discounted payoffs.

Equilibrium Properties: We refer the reader to Section 4.5 for our existence result and an overview of the properties we discuss here.

As shown in the proof of Proposition 1, any Markov stationary strategy can be characterized as a point in the $|\mathbb{X}|$ -fold, compact, Euclidean space, $[0, M]^{|\mathbb{X}|} \subset \mathbb{R}^{|\mathbb{X}|}$, when such strategy is expressed as the $|\mathbb{X}|$ -tuple of the customer’s visit rates over the state space \mathbb{X} . Given the customer’s type, $\vec{\theta}_i$, we can equivalently express her Markov stationary strategy as the $|\mathbb{X}|$ -tuple of her effective product purchase rates, i.e., the product of: (i) her visit hazard rate; and (ii) her product-purchase probability conditional upon a visit, over the state space \mathbb{X} , again as a point in the compact space $[0, M]^{|\mathbb{X}|}$.

Let us denote a strategy for customer i in this space by $s_i = (s_i^1, \dots, s_i^{|\mathbb{X}|}) \in [0, M]^{|\mathbb{X}|}$, and the Cartesian product of the other customers’ strategies by s_{-i} . On this lattice, define also a partial ordering \geq on the set $[0, M]^{|\mathbb{X}|}$, such that for $s_i, \hat{s}_i \in [0, M]^{|\mathbb{X}|}$, $s_i \geq \hat{s}_i$ if and only if $s_i^k \geq \hat{s}_i^k$, $\forall k = 1, \dots, |\mathbb{X}|$. Put simply, $s_i \geq \hat{s}_i$ if customer i purchases (weakly) more aggressively (i.e., adopts a higher effective purchase rate) in each state under s_i . For Cartesian products, such as s_{-i} , a corresponding partial ordering is defined when \geq holds component-wise.

PROPOSITION 2. *Fix the retailer’s markdown hazard rate policy, and consider the resulting dynamic game among the customers. Furthermore, consider only the class, \mathbb{C} , of strategy profiles, $s \in [0, M]^{|\mathbb{X}| \times N}$, such that each player is weakly more aggressive (higher in effective purchase rate) in the successor states that immediately follow the exit (with purchase) of another customer:*

$$\mathbb{C} := \left\{ s \in [0, M]^{|\mathbb{X}| \times N} : s_i^{X'} \geq s_i^X \text{ for all } i, X, X' \in X \cup S_{\text{Exit}}(X) \right\},$$

where $S_{\text{Exit}}(X)$ is the set of states in \mathbb{X} that, relative to the given state X , decrements inventory by one and a customer segment N^{Type} by one. Then for any customer $i \in N$, her best response s_i^* is increasing in s_{-i} , i.e., given $s_{-i}, \hat{s}_{-i} \in \mathbb{C}_{-i}$:

$$\{s_i^*(s_{-i}), s_i^*(\hat{s}_{-i})\} \subset \mathbb{C}_i$$

$$s_{-i} \geq \hat{s}_{-i} \implies s_i^*(s_{-i}) \geq s_i^*(\hat{s}_{-i}).$$

Proposition 2 establishes that, under certain conditions, the customers’ best responses are strategic complements. As formalized in our next proposition, strategic complementarity assures the existence of a stable equilibrium to which best-response dynamics across players converge. Two disclaimers persist, each of which poses a potential obstacle to computational tractability. First, this guarantee lacks an associated rate of convergence, hence as a computational matter convergence is not guaranteed. Second, our strategic complementarity does not encompass the retailer.

Nonetheless, we do not find computational convergence to pose an issue in practice, even with a very large number of customers. In notable contrast, the standard-setting MPEC approach of [Su and Judd \(2012\)](#) and [Dubé et al. \(2012\)](#) proves unworkable for our setting due to the number of variables and constraints.³⁵ The seminal work of [Doraszelski and Judd \(2012\)](#) outlines the computational advantages of the continuous-time formulation of dynamic, stochastic games over their discrete-time counterparts. Our nested fixed point strategy for computational estimation couples our best-response contraction with the advantages of strategic complementarity in a continuous-time framework.

Furthermore, the set of equilibria satisfy the following properties.

PROPOSITION 3. *Fix the retailer’s markdown hazard rate policy, and consider the resulting dynamic game among the customers. Assume that all equilibria are in the complete lattice, \mathbb{C} , defined in Proposition 2. Let τ be a parameter such that the (Markov stationary) best response functions defined in the proof of Proposition 1 are weakly increasing in τ for all players. Then:*

- (i) *There exist extremal (smallest and largest) equilibria that are each increasing in τ .*
- (ii) *Starting from any pre-change equilibrium, best-response dynamics lead to a weakly larger equilibrium in response to an increase in τ .*

³⁵ For our model, the AMPL/Knitro optimization solver obtained over 3 million variables and constraints (each) coming out of its pre-solve stage. For comparison, [Su and Judd \(2012\)](#) put forth 100,000 as a reasonable size limitation for feasibility on each, provided appropriate sparsity and inputs for the solver, as of the time of their writing.

To characterize these results informally, observe that each customer’s prospective payoffs are affected only by the other customers’ effective Markov stationary purchase rates. Proposition 2 formalizes the idea that replacing a customer (or a customer type) with a more aggressive customer (or type) is met with more aggressive purchasing behavior from all other customers.³⁶ This strategic complementarity effect is the outcome of customers’ purchasing behavior affecting prospective availability risk due to the product’s finite inventory.

Proposition 3 formally translates strategic complementarity into the basis for monotone comparative statics and for reliably computing equilibria under best-response dynamics. Like Echenique (2002), whose results we invoke, we rely on characterizing the players’ best-response functions directly rather than relying on the payoff structure of the game.³⁷ In fact, our game is neither supermodular as defined in Milgrom and Roberts (1990) nor a game with strategic complementarities as defined in Milgrom and Shannon (1994).

We discuss how the foregoing lattice structure allows us to address the possibility of multiple equilibria. First, the lattice structure of Proposition 3 imposes the existence of a “maximal” and a “minimal” equilibria under the partial ordering of the lattice, and each of these is guaranteed to be stable. Of course, these two equilibria are one and the same in the special case of a unique equilibrium. Therefore, by computing each of the “maximal” and “minimal” equilibria and simply confirming these to be the same equilibrium, we have a reliable way to ensure uniqueness even absent an ex ante theoretical guarantee. Lastly, by virtue of the (guaranteed) stable and monotone dynamics, we can reach each of these equilibria by initializing the strategy profile with sufficiently aggressive or non-aggressive behavior across all players, then iterating best-response dynamics to an extreme equilibrium. Using this approach to confirm uniqueness, we successfully rule out any instances of multiple equilibria in our estimation and counterfactuals.

Using induction and the Envelope Theorem, we can characterize the effect of changes to many of the parameters describing customer behavior. Proposition 4 and Table EC.6 summarize the effects of such changes on the customer’s best response. By leveraging Proposition 2, we immediately obtain monotone comparative statics to describe customers’ equilibrium responses for a wide range of parameter perturbations where the customer’s best-response purchasing rate becomes unambiguously more or less aggressive. In those cases, we know from Proposition 3 that the feedback loop from the best responses of the other customers will reinforce the direction of the original shift (toward more or less aggressive behavior).

PROPOSITION 4. *Define for all market states $x \in \mathbb{X}$:*

$$B_i^*(x) := Pr_{\lambda^*} \{i \text{ purchases the product at time } \tau \mid X_{\tau-} = x, i \text{ visits at time } \tau, s_{-i}\},$$

³⁶ By a “more aggressive” customer, we again mean a customer who purchases at a higher rate in each state. Note that this is not synonymous with a customer “with a lower cost to monitor.” For an intuitive illustration, consider the extreme case of replacing a typical customer in the market with a new hypothetical customer capable of near-constant monitoring. Such a customer would tend to purchase either immediately upon markdown or shortly prior to stock out. Before such states are reached, she is unlikely to purchase at all relative to the typical customer she has replaced. Therefore, this example fails to fit the paradigm of replacing a customer with a more aggressive type as we have defined it: to qualify, she would also need to purchase at least as aggressively as the replaced customer in the earlier states. Nonetheless, if we were to compute this example as a counterfactual, our algorithms would still exploit the structural lattice properties that arise from the game’s underlying strategic complementarity (in the sense we have defined).

³⁷ See also Echenique (2000).

as the per-visit probability of purchasing the product in the market state x under customer i 's unique Markov stationary, best-response strategy λ^* , given the Markov stationary strategy profile s_{-i} for the rest of the players. Then, fixing any Markov stationary s_{-i} , for all market states $x \in \mathbb{X}$ we have:

(1) $V(x) \geq V^0$, where V^0 is the uniform value of customer i 's value function in her post-exit states, where the inequality is strict for pre-exit states x .

(2) $\frac{\partial}{\partial v} B_i^*(x) \geq 0$, $\frac{\partial}{\partial \gamma} B_i^*(x) \geq 0$, $\frac{\partial}{\partial v} \lambda_x^* \geq 0$, and $\frac{\partial}{\partial \gamma} \lambda_x^* \geq 0$, where the inequalities are strict for pre-exit states.

(3) $\frac{\partial}{\partial \alpha} B_i^*(x) \leq 0$ and $\frac{\partial}{\partial \alpha} \lambda_x^* \leq 0$, where the inequalities are strict for pre-exit states.

(4) $\frac{\partial}{\partial \omega} B_i^*(x) \geq 0$ and $\frac{\partial}{\partial \omega} \lambda_x^* \geq 0$ for pre-depreciation states x , $\frac{\partial}{\partial \omega} B_i^*(x) \leq 0$ and $\frac{\partial}{\partial \omega} \lambda_x^* \leq 0$ for post-depreciation states x , where the inequalities are strict for pre-exit states.

(5) $\frac{\partial}{\partial \lambda_D} B_i^*(x) \geq 0$ and $\frac{\partial}{\partial \lambda_D} \lambda_x^* \geq 0$ for pre-depreciation states x (and exactly zero for post-depreciation states x).

(6) $\frac{\partial}{\partial r_{1,i}} B_i^*(x) \geq 0$ and $\frac{\partial}{\partial r_{2,i}} B_i^*(x) \geq 0$, where the inequalities are strict for pre-exit states.

(7) Assuming that for all $y \in \mathbb{X}$, the customer's visit rate λ_y^* is weakly smaller than the sum of exogenous transition rates and all customers' purchasing rates, then $\frac{\partial}{\partial r_{1,i}} \lambda_x^* < 0$.

The effects of basic parameter shifts on the customer's best response are characterized by claims 2-5 in Proposition 4. These generally agree with intuition — for instance, increasing a customer's valuation naturally serves to increase both her visit rate and her probability of purchase. Increases in her monitoring cost unambiguously increase the customer's probability of purchasing within any single visit (claim 6). Meanwhile, an increase in the monitoring cost's linear coefficient, r_1 , decreases the customer's visit rate in all states (claim 7). For the monitoring cost's quadratic-term coefficient, r_2 , we cannot obtain a similar result that holds unambiguously for each state. Intuitively, r_2 exerts an incentive to smooth out visits over time and states. Whether the customer chooses to increase or decrease her visit rate in any single state will depend heavily on her visit rates in its possible successor states, to and from which smoothing is potentially viable. Since the customer could hypothetically replicate her post-exit, stationary strategy in her pre-exit states for the same expected payoff stream, we would reason that her value function must be higher pre-exit. Statement 1 of Proposition 4 formalizes this intuition.

To touch briefly on the intuitive argument for identification of the model, we note that starting from the baseline of existing dynamic discrete-choice models, our data additionally provide the customer's state-dependent visit frequencies, i.e., the data records the time of and between the visits in which she makes her discrete choices. For a customer, her observed baseline frequency of monitoring and how responsive she is to her (changing) expected payoffs from visiting at a given time (state) are novel components of our data that naturally serve to additionally identify the two parameters of her monitoring costs. Proposition 4 and Table EC.6 provide more detail on the comparative statics.

Table EC.6 Monotone Comparative Statics of Parameter Changes on the Customer's Own Best Response

	Effect on the Visit Rate	Effect on the In-Visit Purchase Probability
Increase valuation (v or γ)	+	+
Increase price elasticity (α)	–	–
Increase depreciation magnitude (ω)	Pre-depreciation: + Post-depreciation: –	+ –
Increase depreciation hazard (λ_D)	+	+
Monitoring cost linear coefficient (r_1)	–	+
Monitoring cost quadratic coefficient (r_2)	Not Uniform	+

Refer to Proposition 4 for precise statements of conditions and applicable states.

Appendix C: Proofs

Preliminaries: Deriving the Continuous-Time Bellman Equation

We discuss informally the instantaneous Bellman equation solved by the customer to determine her best response. Define customer i 's gain function:

$$J(x, s) := \mathbb{E} \left[- \int_0^\infty e^{-\rho \cdot t} \cdot \delta(\lambda_t^{(i)}) dt + \sum_{m=1}^\infty e^{-\rho \cdot T_m^{(i)}} \cdot \Psi_i \left(X_{T_m^{(i)}-}^x, c_t^{(i)} \right) \right],$$

where the process X_t^x denotes the market-state, stochastic process with the initial state $x \in \mathbb{X}$ at time $t = 0$ (implicitly, under the strategy profile s). We can further define the corresponding value function:

$$V_i(x) := \sup_{s_i \in \mathcal{S}_i} J(x, s_i \times s_{-i}),$$

holding fixed the strategies of all other players.

For each state $X_t \in \mathbb{X}$, we can derive the instantaneous Bellman equation using the standard tools for analyzing continuous-time jump processes. Instead of a formal derivation, we present it informally for intuition. Note that for a sufficiently small time increment Δt we have:

$$\begin{aligned} V_i(X_t) &= \max_{\lambda \in [0, M], c \in C_t^{(i)}} \left\{ -\delta(\lambda) \cdot \Delta t + (1 + \rho \cdot \Delta t)^{-1} \right. \\ &\quad \left. \times \mathbb{E}_{\lambda, c} \left[V_i(X_t + \Delta X) + \Psi_i(X_{\tau-}, c) \cdot \mathbf{1}_{\{\text{Customer } i \text{ visits at } \tau \in (t, t + \Delta t)\}} \right] \right\} \\ &= \max_{\lambda \in [0, M]} \left\{ -\delta(\lambda) \cdot \Delta t + (1 + \rho \cdot \Delta t)^{-1} \right. \\ &\quad \left. \times \mathbb{E}_\lambda \left[\max_{c \in C_t^{(i)}} \mathbb{E}_c \left[V_i(X_t + \Delta X) + \Psi_i(X_{\tau-}, c) \cdot \mathbf{1}_{\{\text{Customer } i \text{ visits at } \tau \in (t, t + \Delta t)\}} \mid \lambda \right] \right] \right\}. \end{aligned} \quad (\text{EC.1})$$

Technically, the value function will also depend on whether customer i has exited the market, but it is straightforward to treat the post-exit cases separately. Next, multiply both sides by $(1 + \rho \cdot \Delta t)$, then subtract V_i from each and let $\Delta t \rightarrow 0$ and ignore terms of order dt^2 , to obtain:

$$\begin{aligned} \rho \cdot V_i(X_t) \cdot dt &= \max_{\lambda \in [0, M]} \left\{ -\delta(\lambda) \cdot dt \right. \\ &\quad \left. + \mathbb{E}_\lambda \left[\max_{c \in C_t^{(i)}} \mathbb{E}_c \left[dV_i + \Psi_i(X_{\tau-}, c) \cdot \mathbf{1}_{\{\text{Customer } i \text{ visits at } \tau \in (t, t + \Delta t)\}} \mid \lambda \right] \right] \right\}. \end{aligned} \quad (\text{EC.2})$$

We can explicitly write out Ψ and unpack dV treating the state transitions as a standard jump process (applying the Ito Lemma for jump-diffusion processes) to reach the Bellman equation to be solved by the customer i for all states X in the state space to select her endogenous monitoring hazard rate $\lambda_X^* \in [0, M]$:

$$\begin{aligned} \rho \cdot V_i(X) &= \max_{\lambda \in [0, M]} \left\{ -\delta_i(\lambda) + \sum_{y \in \mathcal{S}(X)} q(X, y) \cdot (V_i(y) - V_i(X)) \right. \\ &\quad \left. + \lambda \cdot \mathbb{E} \left[\max \left\{ V_i(X) + \epsilon_t^{(i),0}, \gamma_i + v_j - \omega \cdot D_t - \alpha_i \cdot p_{jt} + V_i^0 + \epsilon_t^{(i),1}, w_i + V_i(X) + \epsilon_t^{(i),2} \right\} \right] - \lambda \cdot V_i(X) \right\}, \end{aligned} \quad (\text{EC.3})$$

and integrating out the unobserved-state contributions, $\epsilon_t^{(i),\cdot}$, completes an informal derivation of the instantaneous Bellman equation as given in Equation 5.

Proof of Proposition 1

First, we justify that restricting equilibrium strategies to those that are Markov and stationary is without any loss of generality when considering the existence of an equilibrium in Markov strategies. It suffices: (i) to show that for each Markov, stationary $s_{-i} \in S_{-i}$, customer i 's best-response correspondence $s_i^* : S_{-i} \rightarrow 2^{S_i}$ is well-defined at s_{-i} and includes a Markov, stationary strategy; (ii) then, to define $\bar{S} \subset S$, as the strategy space subset of Markov, stationary strategies, and $\bar{s}_i^* : \bar{S}_{-i} \rightarrow 2^{\bar{S}_i}$, as the best-response correspondence over this set; and (iii) finally, to show the existence of a fixed point of \bar{s}^* . Note that by (i), the correspondence \bar{s}_i^* of (ii) is well-defined and that its best responses maximize the expected discounted utility payoff of customer i from among the set of all available strategies, i.e.,

$$\forall s_{-i} \in \bar{S}_{-i}, \bar{s}_i^*(s_{-i}) \neq \emptyset, \bar{s}_i^*(s_{-i}) \subset s_i^*(s_{-i}).$$

Therefore, the fixed point of (iii) both comprises only stationary, Markov strategies and is a best-response equilibrium without imposing any further, a priori restriction on the strategy space S with respect to which best responses are defined.

Step (i): Take any Markov stationary $s_{-i} \in S_{-i}$. Then, a best-response strategy, if one exists, is a solution to agent i 's (single-agent) dynamic optimization problem. The finite, Markov stationary, state-transition hazard rates for agent i 's problem are denoted by q below without any notational reference to their dependence on s_{-i} . Furthermore, define $\mathbb{B}(\mathbb{X})$ to be the class of (bounded) real-valued functions on the finite, market-state space \mathbb{X} , and define on it the metric d as follows:

$$\forall f_1, f_2 \in \mathbb{B}(\mathbb{X}), d(f_1, f_2) := \sup_{X \in \mathbb{X}} |f_1(X) - f_2(X)|.$$

Then, note that $(\mathbb{B}(\mathbb{X}), d)$ is a complete metric space. In addition, define the operator T mapping $\mathbb{B}(\mathbb{X})$ into itself such that, for $f \in \mathbb{B}(\mathbb{X})$, $Tf : \mathbb{X} \rightarrow \mathbb{R}$ is given by:³⁸

$$\begin{aligned} (Tf)(X) &= \left(\rho + \sum_{X' \in \mathbb{X}, X' \neq X} q(X, X') + M \right)^{-1} \\ &\times \max_{\lambda \in [0, M]} \left\{ -r_1 \cdot \lambda - \frac{r_2}{2} \cdot \lambda^2 + \sum_{X' \in \mathbb{X}, X' \neq X} q(X, X') \cdot f(X') \right. \\ &\left. + \lambda \cdot [\Gamma + \log(1 + \exp\{w_i\}) + \exp\{\gamma_i + v - \omega \cdot D_t - \alpha_i \cdot p_t + f(0) - f(X)\}] + M \cdot f(X) \right\}. \quad (\text{EC.4}) \end{aligned}$$

To define T as a valid mapping from $\mathbb{B}(\mathbb{X})$ into itself as above, check that the operator is in fact closed in $\mathbb{B}(\mathbb{X})$ by simply noting that the bracketed maximand on the right-hand side is continuous in λ over its compact domain for any $X \in \mathbb{X}$, hence attaining a real-valued maximum over λ for each $X \in \mathbb{X}$.³⁹ Moreover,

³⁸ We have simplified notation slightly by using “0” to uniformly denote post-market-exit states.

³⁹ Here, we rely on the assumption of the market-state space as being finite. Absent this assumption, the general approach involves defining $\mathbb{B}(\mathbb{X})$ to be the class of bounded, upper-semicontinuous functions over an appropriate state space \mathbb{X} , and then showing (roughly stated): (i) that, provided the transition hazard function q is continuous (in an appropriately sufficient sense) in the current state for each successor, the maximand defining Tf is upper-semicontinuous as a function of X and λ ; and (ii) provided the maximand is also bounded, maximizing over $\lambda \in [0, M]$ yields a function over \mathbb{X} that is upper-semicontinuous and bounded. Therefore, with appropriate assumptions on q

T is easily seen to be a monotone operator, by splitting the final term $M \cdot f(X)$ into $(M - \lambda) \cdot f(X)$ and $\lambda \cdot f(X)$, and distributing the latter inside the bracketed logarithm.

We now show that T is a contraction mapping on $(\mathbb{B}(\mathbb{X}), d)$, implying then that a unique fixed point exists by the Banach fixed-point theorem. Take arbitrary $f_1, f_2 \in \mathbb{B}(\mathbb{X})$. By definition, $f_1 \leq f_2 + d(f_1, f_2)$, and by T being monotone, we obtain $Tf_1 \leq T(f_2 + d(f_1, f_2))$. By the definition of T , we obtain that $\forall X \in \mathbb{X}$ we have:

$$\begin{aligned} T(f_2 + d(f_1, f_2))(X) &= \left(\rho + \sum_{X' \in \mathbb{X}, X' \neq X} q(X, X') + M \right)^{-1} \\ &\quad \times \max_{\lambda \in [0, M]} \left\{ -r_1 \cdot \lambda - \frac{r_2}{2} \cdot \lambda^2 + \sum_{X' \in \mathbb{X}, X' \neq X} q(X, X') \cdot f_2(X') \right. \\ &\quad \left. + \lambda \cdot [\Gamma + \log(1 + \exp\{w_i\} + \exp\{\gamma_i + v - \omega \cdot D_t - \alpha_i \cdot p_t + f_2(0) - f_2(X)\})] \right. \\ &\quad \left. + M \cdot f_2(X) + \left(\sum_{X' \in \mathbb{X}, X' \neq X} q(X, X') + M \right) \cdot d(f_1, f_2) \right\} \\ &\leq Tf_2(X) + \frac{\sum_{X' \in \mathbb{X}, X' \neq X} q(X, X') + M}{\rho + \sum_{X' \in \mathbb{X}, X' \neq X} q(X, X') + M} \cdot d(f_1, f_2). \end{aligned} \tag{EC.5}$$

Finally, we obtain:

$$Tf_1 - Tf_2 \leq \beta \cdot d(f_1, f_2),$$

where since \mathbb{X} is finite we have:

$$\beta := \max_{X \in \mathbb{X}} \frac{\sum_{X' \in \mathbb{X}, X' \neq X} q(X, X') + M}{\rho + \sum_{X' \in \mathbb{X}, X' \neq X} q(X, X') + M} < 1.$$

Therefore, by the Banach fixed-point theorem, T is a contraction mapping with a unique fixed point in $\mathbb{B}(\mathbb{X})$. Let us denote this unique fixed point by $V^* \in \mathbb{B}(\mathbb{X})$, in which case $V^* = TV^*$ is equivalent to, $\forall X \in \mathbb{X}$:

$$\begin{aligned} \rho \cdot V^*(X) &= \max_{\lambda \in [0, M]} \left\{ -r_1 \cdot \lambda - \frac{r_2}{2} \cdot \lambda^2 + \sum_{X' \in \mathbb{X}, X' \neq X} q(X, X') \cdot (V^*(X') - V^*(X)) \right. \\ &\quad \left. + \lambda \cdot [\Gamma + \log(1 + \exp\{w_i\} + \exp\{\gamma_i + v - \omega \cdot D_t - \alpha_i \cdot p_t + V^*(0) - V^*(X)\})] \right\}. \end{aligned} \tag{EC.6}$$

The bracketed maximand on the right-hand side can be viewed as a function of X and λ on the product space $\mathbb{X} \times [0, M]$. We can obtain directly via a continuity argument (or alternatively by using the Theorem of the Maximum, or for generality's sake by applying a selection theorem such as the ‘‘Selection Theorem’’ in Section 3 of Maitra (1968)), the existence of a map $\hat{s}_i : \mathbb{X} \rightarrow [0, M]$ such that $\hat{s}_i(X)$ as a policy achieves such function's maximum $V^*(X)$ at X , for each $X \in \mathbb{X}$. Then \hat{s}_i is a Markov, stationary policy, satisfying the Bellman optimality condition given s_{-i} ; therefore it is a Markov, stationary best response of agent i to s_{-i} .

Step (ii): Follows directly from Step (i) above.

and \mathbb{X} , T is still a valid map from $\mathbb{B}(\mathbb{X})$ into itself, suggesting that we may proceed to obtain a fixed point of T and apply an appropriate selection theorem to prove that an optimal, stationary policy exists.

However, guaranteeing properties for the hazard rates q requires these properties to hold for the stationary, best-response policies of $-i$, raising a ‘‘self-referential’’ problem complicating the choice of strategy and state-space topology. See Dutta and Sundaram (1998) for a discussion of these issues in the context of discrete-time, Markovian games.

Step (iii): Lastly, we apply Kakutani's fixed-point theorem to establish the existence of an equilibrium in Markov, stationary strategies. Since Step (i) sets forth the unique Markov stationary best response, it imposes certain necessary conditions on the (Markov stationary) equilibrium strategies. Without loss of generality, we can restrict our attention to strategies satisfying these necessary conditions in looking for the fixed point equilibrium. In particular, we will not impose any restriction on the state-wise hazard rates that the customer may choose, but we will impose that her discrete choices satisfy expected utility maximization (taking as given both the opponents' strategy profiles and her own chosen visit hazard rates, λ , in each state) going forward. With this restriction, we can equivalently view the game as each customer selecting her hazard rates for each state, with her discrete choices implicit in her expected payoff to the full strategy profile. Note that as her opponents' strategies change, her strategy, $s_i \in \bar{S}_i$, under the original game also shifts even if we hold her choice of hazard rates constant. The change in expected payoffs associated with this own-strategy shift is pushed into the payoff function of the new game.⁴⁰ A fixed point in the new game corresponds to a unique fixed point in the original game.

Now each strategy in hazard rates can be uniquely represented as a (unique) point in the compact, convex, Euclidean cube, $[0, M]^{|\mathbb{X}|} \subset \mathbb{R}^{|\mathbb{X}|}$, namely $s_i(X_1), \dots, s_i(X_{|\mathbb{X}|})$, where $X_k, k = 1, \dots, |\mathbb{X}|$ denote the states in \mathbb{X} . Then we can write the best-response correspondence in the space of Markov stationary strategies, $\bar{s}^* : \bar{S} \rightarrow 2^{\bar{S}}$, by its analogous representation, $\bar{s}^* : [0, M]^{|\mathbb{X}| \cdot N} \rightarrow 2^{[0, M]^{|\mathbb{X}| \cdot N}}$, with $[0, M]^{|\mathbb{X}| \cdot N}$ clearly compact, convex, and non-empty. We have already shown the best-response to be non-empty, and the fact that it is convex-valued is trivial given the uniqueness of V^* .

The only remaining property to confirm is the upper-semicontinuity of the best-response correspondence \bar{s}^* ; for this, it suffices to demonstrate the upper-semicontinuity of i 's best-response mapping \bar{s}_i^* w.l.o.g.

As a preliminary matter, now let $V_i^*(\bar{s})$ be defined for each $\bar{s} \in \bar{S}$ as the respective value function $V^* \in \mathbb{B}(\mathbb{X})$ associated with i and \bar{s}_{-i} , as defined and derived in Step (i) of this proof. Now observe that V_i^* is bounded uniformly over \bar{S} and \mathbb{X} , below by $-\frac{r_i}{2 \cdot \rho} \cdot M^2$ and above by:

$$\frac{\bar{q}}{\rho} \cdot \log \left(1 + 3 \cdot \max_{X', X'' \in \mathbb{X}} \left[\exp \{ U_{it}(X', X'') \} \right] \right),$$

where:

$$\bar{q} := \sum_{X', X'' \in \mathbb{X}} q(X', X'') + \bar{M} \cdot M < \infty.$$

Now suppose $\bar{s}^k \in \bar{S}$ and $s_i^k \in \bar{S}_i$, $k \in \mathbb{N}$, such that:

$$\bar{s}^k \rightarrow \bar{s} \in \bar{S}, s_i^k \in \bar{s}_i^*(\bar{s}^k) \forall k, \text{ and } s_i^k \rightarrow s_i \in \bar{S}_i,$$

whereupon we wish to show $s_i \in \bar{s}_i^*(\bar{s})$. Consider the sequence of i 's value functions $V_i^*(\bar{s}^k)$. By the uniform bounds on V_i^* , there exists a convergent subsequence $V_i^*(\bar{s}^k)(X)$ for each $X \in \mathbb{X}$, and therefore for the finite-dimensional functions $V_i^*(\bar{s}^k)$ as well. Call its limit $\bar{V} \in \mathbb{B}(\mathbb{X})$, and henceforth consider only this convergent

⁴⁰ More formally, the new game's payoff function incorporates the Markov stationary equilibrium expected payoffs of the players' dynamic discrete choice game defined by holding the strategy profile's hazard rates fixed. The equilibrium existence proof is similar to the one presented here.

subsequence of the index k , while preserving the index k for purposes of notation. Clearly, for the state-transition hazard rates, $q(\bar{s}^k) \rightarrow q(\bar{s})$ pointwise (hence uniformly) over state pairs, with $q(\cdot)$ making explicit the rates' previously-suppressed dependence on other players' strategies. Finally, note that i 's instantaneous Bellman equation holds for each k , with $V_i^*(\bar{s}^k)$, $q(\bar{s}^k)$, and s_i^k . Taking the limit as $k \rightarrow \infty$ yields that i 's instantaneous Bellman equation holds for \bar{V} , $q(\bar{s})$, and s_i , which is our desired result. Q.E.D.

Proof of Proposition 2

Our proof proceeds by induction using the fact that we can proceed through the states of the state space, \mathbb{X} , in an ordering where no state is reached prior to having reached all of its successors. Suppose $s_{-i} \geq \hat{s}_{-i}$ (and $s_{-i} \not\leq \hat{s}_{-i}$ to rule out the trivial case). Observe that there exists a finite (of some length K) sequence of strategy profiles, $s_{-i}^{(K)}, \dots, s_{-i}^{(1)}$, such that $s_{-i} = s_{-i}^{(K)} \geq \dots \geq s_{-i}^{(1)} = \hat{s}_{-i}$ and each pair $s_{-i}^{(k+1)}$ and $s_{-i}^{(k)}$ differ by the behavior of a single opponent being strictly more aggressive in a single state, with such state being weakly further along the induction ordering than for $s_{-i}^{(j+1)}$ and $s_{-i}^{(j)}$ with $j \leq k$. In short, we decompose the difference in behavior between s_{-i} and \hat{s}_{-i} into the incremental increases in aggression by the single opponents in a given state, running through the states in the induction ordering. Note that each $s_{-i}^{(k)}$ is then guaranteed to be an element of \mathbb{C}_{-i} , and since the relation \geq is transitive, without loss of generality it suffices to prove our result for an increase in aggression for a single opponent in a single state. Therefore, suppose without loss of generality that $s_{-i} \geq \hat{s}_{-i}$ with a single opponent acting strictly more aggressively in a single state, $X \in \mathbb{X}$, under s_{-i} .

Let $V_i^*(\cdot; s_{-i})$ and $V_i^*(\cdot; \hat{s}_{-i})$ denote the best-response value functions of customer i given the opponent strategy profiles s_{-i} and \hat{s}_{-i} , respectively. Furthermore, let X^- denote the successor state of X reached when a customer of the more aggressive opponent's type purchases. Then, by the Bellman equations, we have:

$$\begin{aligned} \rho \cdot [V_i^*(X; s_{-i}) - V_i^*(X; \hat{s}_{-i})] &= \frac{r_2 \cdot [\lambda_X^*(s_{-i})^2 - \lambda_X^*(\hat{s}_{-i})^2]}{2} + q(X, X^-; s_{-i}) \cdot [V_i^*(X^-; s_{-i}) - V_i^*(X; s_{-i})] \\ &\quad - q(X, X^-; \hat{s}_{-i}) \cdot [V_i^*(X^-; \hat{s}_{-i}) - V_i^*(X; \hat{s}_{-i})]. \end{aligned}$$

Letting $\Delta q := q(X, X^-; s_{-i}) - q(X, X^-; \hat{s}_{-i}) > 0$ and using that $V_i^*(X^-; s_{-i}) = V_i^*(X^-; \hat{s}_{-i})$ we obtain:

$$\begin{aligned} [\rho + q(X, X^-; \hat{s}_{-i})] \cdot [V_i^*(X; s_{-i}) - V_i^*(X; \hat{s}_{-i})] &+ \frac{r_2 \cdot [\lambda_X^*(\hat{s}_{-i})^2 - \lambda_X^*(s_{-i})^2]}{2} \\ &= \Delta q \cdot [V_i^*(X^-; s_{-i}) - V_i^*(X; s_{-i})]. \end{aligned} \tag{EC.7}$$

Now, $s_i^*(s_{-i}) \geq s_i^*(\hat{s}_{-i})$ if and only if $V_i^*(X; s_{-i}) \leq V_i^*(X; \hat{s}_{-i})$. The important observation from Equation (EC.7) is that the left-hand side is non-positive if and only if $V_i^*(X; s_{-i}) \leq V_i^*(X; \hat{s}_{-i})$. Therefore, by Equation (EC.7), the condition that $V_i^*(X^-; s_{-i}) - V_i^*(X; s_{-i})$ be non-positive for all $s_{-i} \in \mathbb{C}_{-i}$ is sufficient to complete our proof.

Therefore, we now show that $V_i^*(X; s_{-i}) \geq V_i^*(X^-; s_{-i})$ for all $s_{-i} \in \mathbb{C}_{-i}$, by induction over the state space. Take and fix any $s_{-i} \in \mathbb{C}_{-i}$, which we now make implicit in the notation. Without loss of generality, let X^- denote the successor state of X reached upon a purchase by an opponent of the type resulting in the highest value $V_i^*(X^-; s_{-i})$ across the possible types. Then:

$$\begin{aligned} \rho \cdot [V_i^*(X) - V_i^*(X^-)] &= \frac{r_2 \cdot [\lambda_X^{*2} - \lambda_{X^-}^{*2}]}{2} + \sum_{y \in \mathcal{S}(X)} q(X, y) \cdot [V_i^*(y) - V_i^*(X)] \\ &\quad - \sum_{z \in \mathcal{S}(X^-)} q(X^-, z) \cdot [V_i^*(z) - V_i^*(X^-)]. \end{aligned}$$

For the induction step, we seek to prove that $V_i^*(X) \geq V_i^*(X^-)$ for X , while assuming under the induction hypothesis that this property holds for all successor states of X . First, note that the exogenous transition hazards (of discount, depreciation, or outside purchase) are the same in the states X and X^- . Labeling for state X the set of successor states reached by an exogenous transition as $\tilde{S}(X)$, we obtain:

$$\begin{aligned} & \left[\rho + \sum_{y \in \tilde{S}(X)} q(X, y) \right] \cdot [V_i^*(X) - V_i^*(X^-)] + \frac{r_2 \cdot [\lambda_{X^-}^{*-2} - \lambda_X^{*-2}]}{2} \\ &= \left[\sum_{y \in \tilde{S}(X)} q(X, y) \cdot V_i^*(y) - \sum_{z \in \tilde{S}(X^-)} q(X^-, z) \cdot V_i^*(z) \right] \\ &+ \sum_{w \in S(X) \setminus \tilde{S}(X)} q(X, w) \cdot [V_i^*(w) - V_i^*(X)] - \sum_{x \in S(X^-) \setminus \tilde{S}(X^-)} q(X^-, x) \cdot [V_i^*(x) - V_i^*(X^-)], \end{aligned}$$

where the right-hand side's first bracketed term is non-negative by the induction hypothesis. The remaining terms on the right-hand side capture the transition hazards imposed by the opponents' purchasing rates in the states X and X^- . For any type θ , let X_θ and X_θ^- be the respective successor states reached from X and X^- upon a purchase by the type θ . Denote as θ^M the type of customer whose purchase causes a transition from X to X^- . Then:

$$\begin{aligned} & \left[\rho + \sum_{y \in \tilde{S}(X)} q(X, y) + q(X, X^-) \right] \cdot [V_i^*(X) - V_i^*(X^-)] + \frac{r_2 \cdot [\lambda_{X^-}^{*-2} - \lambda_X^{*-2}]}{2} \\ &= \left[\sum_{y \in \tilde{S}(X)} q(X, y) \cdot V_i^*(y) - \sum_{z \in \tilde{S}(X^-)} q(X^-, z) \cdot V_i^*(z) \right] \\ &+ q(X^-, X_{\theta^M}^-) \cdot [V_i^*(X^-) - V_i^*(X_{\theta^M}^-)] \\ &+ \sum_{\theta \neq \theta^M} q(X, X_\theta) \cdot [V_i^*(X_\theta) - V_i^*(X)] - \sum_{\theta \neq \theta^M} q(X^-, X_\theta^-) \cdot [V_i^*(X_\theta^-) - V_i^*(X^-)]. \end{aligned}$$

Finally we consider purchases by each type θ other than θ^M . Taking any such type θ , by the definition of the class \mathbb{C}_{-i} , the type purchases more aggressively in state X^- than in X , i.e., $q(X, X_\theta) \leq q(X^-, X_\theta^-)$. Then defining $\Delta q_\theta := q(X^-, X_\theta^-) - q(X, X_\theta) \geq 0$:

$$\begin{aligned} & \left[\rho + \sum_{y \in \tilde{S}(X)} q(X, y) + q(X, X^-) + \sum_{\theta \neq \theta^M} q(X, X_\theta) \right] \cdot [V_i^*(X) - V_i^*(X^-)] + \frac{r_2 \cdot [\lambda_{X^-}^{*-2} - \lambda_X^{*-2}]}{2} \\ &= \left[\sum_{y \in \tilde{S}(X)} q(X, y) \cdot V_i^*(y) - \sum_{z \in \tilde{S}(X^-)} q(X^-, z) \cdot V_i^*(z) \right] \\ &+ q(X^-, X_{\theta^M}^-) \cdot [V_i^*(X^-) - V_i^*(X_{\theta^M}^-)] \\ &+ \sum_{\theta \neq \theta^M} \Delta q_\theta \cdot [V_i^*(X^-) - V_i^*(X_\theta^-)] + \sum_{\theta \neq \theta^M} q(X, X_\theta) \cdot [V_i^*(X_\theta) - V_i^*(X_\theta^-)]. \end{aligned} \tag{EC.8}$$

The left-hand side of Equation (EC.8) is non-negative if and only if $V_i^*(X) \geq V_i^*(X^-)$. The right-hand side is non-negative by the induction hypothesis. This completes the induction step, and the base case is easily shown. Q.E.D.

Proof of Proposition 3

Claim (i) is shown using Proposition 2 and invoking Tarski's Theorem. (A useful approach is to modify the concise proof for Theorem 6 in Milgrom and Roberts (1990) without relying on supermodularity. Note that our Markov stationary best response is unique, hence a function.) Claim (ii) relies on Proposition 2 and Theorem 3 of Echenique (2002). Q.E.D.

Proof of Proposition 4

As with Proposition 2, our proofs of Statements 1-4 proceed by induction over the state space, \mathbb{X} , to prove that a property of interest holds for the value function in all states.

For each Statement, our proof focuses on the induction step only, as each base case is easily confirmed. For state $X \in \mathbb{X}$, we denote the set of its successors (i.e., states that can be reached from X by a single state transition) by $S(X)$. Lastly, for notational convenience, define for any function $f \in \mathbb{B}(\mathbb{X})$ and any state $X \in \mathbb{X}$:

$$G_f(X) := \log(1 + \exp\{w\}) + \exp\{v + \gamma - \omega \cdot D_X - \alpha \cdot p_X + f^0 - f(X)\},$$

where f^0 is shorthand for the value of f at the market-exit state from X .

(1) *Proof:* It is readily confirmed that $\rho \cdot V^0 = \frac{r_2}{2} \cdot \lambda_0^{*2} = \frac{[\Gamma - r_1]^2}{2r_2} > 0$. Similarly, for any $X \in \mathbb{X}$:

$$\rho \cdot V(X) = \frac{r_2}{2} \cdot \lambda_X^{*2} + \sum_{y \in S(X)} q(X, y) \cdot [V(y) - V(X)]. \quad (\text{EC.9})$$

Therefore, supposing the induction hypothesis to hold over $S(X)$:

$$\rho \cdot (V^0 - V(X)) \leq \frac{r_2}{2} \cdot (\lambda_0^{*2} - \lambda_X^{*2}) - [V^0 - V(X)] \cdot \sum_{y \in S(X)} q(X, y) \quad (\text{EC.10})$$

$$\implies V^0 - V(X) \leq \frac{-2 \cdot (\Gamma - r_1) \cdot G_V(X) - G_V(X)^2}{2r_2 \cdot (\rho + \sum_{y \in S(X)} q(X, y))} \leq 0, \quad (\text{EC.11})$$

with the final inequality strict if X is a pre-exit state. Q.E.D.

(2) *Proof:* For pre-exit state X , since $\frac{\partial}{\partial v} B^*(X) = B^*(X) \cdot [1 - B^*(X)] \cdot [1 - \frac{\partial}{\partial v} V(X)]$ and $\frac{\partial}{\partial v} \lambda_X^* = \frac{1}{r_2} \cdot B^*(X) \cdot [1 - \frac{\partial}{\partial v} V(X)]$, our intended result follows if $\frac{\partial}{\partial v} V(X) < 1$. The induction step follows from the induction hypothesis after taking the partial derivative on both sides of the Bellman equation, and the base case is easily shown. The result for γ is analogous. Q.E.D.

(3) *Proof:* For pre-exit state X , since $\frac{\partial}{\partial \alpha} B^*(X) = B^*(X) \cdot [1 - B^*(X)] \cdot [-p_X - \frac{\partial}{\partial \alpha} V(X)]$ and $\frac{\partial}{\partial \alpha} \lambda_X^* = \frac{1}{r_2} \cdot B^*(X) \cdot [-p_X - \frac{\partial}{\partial \alpha} V(X)]$, our intended result follows if $\frac{\partial}{\partial \alpha} V(X) > -p_X$. For the induction step, by the Bellman equation, the induction hypothesis, and price levels being nonincreasing:

$$\frac{\partial}{\partial \alpha} V(X) > \frac{-\rho \cdot p_X - \lambda_X^* \cdot B^*(X) \cdot p_X - \sum_{y \in S(X)} q(X, y) \cdot p_y}{\rho + \lambda_X^* \cdot B^*(X) + \sum_{y \in S(X)} q(X, y)} \geq -p_X,$$

and the claim follows. Q.E.D.

(4) *Proof:* For pre-exit, post-depreciation states,

$$\frac{\partial}{\partial \omega} B^*(X) = -B^*(X) \cdot [1 - B^*(X)] \cdot [1 + \frac{\partial}{\partial \omega} V(X)]$$

$$\text{and } \frac{\partial}{\partial \omega} \lambda_X^* = -\frac{1}{r_2} \cdot B^*(X) \cdot [1 + \frac{\partial}{\partial \omega} V(X)].$$

Our intended result follows if $\frac{\partial}{\partial \omega} V(X) > -1$. The induction step follows from the induction hypothesis using the Bellman equation after applying the partial derivative on both sides.

For pre-exit, pre-depreciation states,

$$\frac{\partial}{\partial \omega} B^*(X) = -B^*(X) \cdot [1 - B^*(X)] \cdot \frac{\partial}{\partial \omega} V(X)$$

$$\text{and } \frac{\partial}{\partial \omega} \lambda_X^* = -\frac{1}{r_2} \cdot B^*(X) \cdot \frac{\partial}{\partial \omega} V(X).$$

Our intended result follows if $\frac{\partial}{\partial \omega} V(X) < 0$ follows from an induction hypothesis of $\frac{\partial}{\partial \omega} V(y) \leq 0$ for all successor states, $y \in S(X)$, for any pre-depreciation state, X , and if X is post-depreciation, that $\frac{\partial}{\partial \omega} V(X) \leq 0$ follows from the same induction hypothesis. These can be shown to hold. Q.E.D.

(5) *Proof:* For post-depreciation states, the result is immediately evident. For pre-exit, pre-depreciation states, we have

$$\frac{\partial}{\partial \lambda_D} B^*(X) = -B^*(X) \cdot [1 - B^*(X)] \cdot \frac{\partial}{\partial \lambda_D} V(X),$$

$$\text{and } \frac{\partial}{\partial \lambda_D} \lambda_X^* = -\frac{1}{r_2} \cdot B^*(X) \cdot \frac{\partial}{\partial \lambda_D} V(X).$$

Therefore, it suffices to show that $\frac{\partial}{\partial \lambda_D} V(X) < 0$. Letting X_D denote the depreciated successor state of X :

$$\frac{\partial}{\partial \lambda_D} V(X) = \frac{V(X_D) - V(X) + \sum_{y \in S(X)} q(X, y) \cdot \frac{\partial}{\partial \lambda_D} V(y)}{\rho + \lambda_X^* \cdot B^*(X) + \sum_{y \in S(X)} q(X, y)}.$$

Therefore, with the induction hypothesis, it is sufficient to show that $V(X) > V(X_D)$, which can be shown using straightforward arguments. Q.E.D.

(6) *Proof:* First note that:

$$\frac{\partial}{\partial r_1} B_i^*(X) = B_i^*(X) \cdot (1 - B_i^*(X)) \cdot \left[\frac{\partial}{\partial r_1} V^0 - \frac{\partial}{\partial r_1} V(X) \right], \quad (\text{EC.12})$$

$$\frac{\partial}{\partial r_2} B_i^*(X) = B_i^*(X) \cdot (1 - B_i^*(X)) \cdot \left[\frac{\partial}{\partial r_2} V^0 - \frac{\partial}{\partial r_2} V(X) \right]. \quad (\text{EC.13})$$

Since $B_i^*(X) \in (0, 1)$ for all pre-exit $X \in \mathbb{X}$, the sign of $\frac{\partial}{\partial r_1} B_i^*(X)$ depends solely on that of $\left[\frac{\partial}{\partial r_1} V^0 - \frac{\partial}{\partial r_1} V(X) \right]$ (and similarly for the partial with respect to r_2).

Taking the partial derivatives w.r.t. r_1 of both sides of the Bellman equation while using the Envelope Theorem, and then differencing after simple manipulation yields, under the induction hypothesis $\left[\frac{\partial}{\partial r_1} V^0 - \frac{\partial}{\partial r_1} V(y) \right] \geq 0$ holding over $y \in S(X)$:

$$\frac{\partial}{\partial r_1} V^0 - \frac{\partial}{\partial r_1} V(X) = \frac{\lambda_X^* - \lambda_0^* + \sum_{y \in S(X) \setminus 0} q(X, y) \cdot \left[\frac{\partial}{\partial r_1} V^0 - \frac{\partial}{\partial r_1} V(y) \right]}{\rho + \sum_{y \in S(X)} q(X, y) + \lambda_X^* \cdot B^*(X)} \geq 0. \quad (\text{EC.14})$$

The analogous steps w.r.t. r_2 yield:

$$\frac{\partial}{\partial r_2} V^0 - \frac{\partial}{\partial r_2} V(X) = \frac{\lambda_X^{*2} - \lambda_0^{*2} + \sum_{y \in S(X) \setminus 0} q(X, y) \cdot \left[\frac{\partial}{\partial r_2} V^0 - \frac{\partial}{\partial r_2} V(y) \right]}{\rho + \sum_{y \in S(X)} q(X, y) + \lambda_X^* \cdot B^*(X)} \geq 0, \quad (\text{EC.15})$$

which concludes the proof of the claim. Q.E.D.

(7) *Proof:* Since $r_2 \cdot \frac{\partial}{\partial r_1} \lambda_X^* = -1 + B^*(X) \cdot \left[\frac{\partial}{\partial r_1} V^0 - \frac{\partial}{\partial r_1} V(X) \right]$ for pre-exit states $X \in \mathbb{X}$, it is sufficient to show that $\frac{\partial}{\partial r_1} V^0 - \frac{\partial}{\partial r_1} V(X) < 1$ for all such X . Under the induction hypothesis and the sufficient condition assumed in the proposition's statement, the following steps suffice. Note that:

$$\begin{aligned} \frac{\partial}{\partial r_1} V^0 &= -\frac{\lambda_0^*}{\rho} \text{ and} \\ \rho \cdot \frac{\partial}{\partial r_1} V(X) &= -\lambda_X^* + \lambda_X^* \cdot B^*(X) \cdot \left[\frac{\partial}{\partial r_1} V^0 - \frac{\partial}{\partial r_1} V(X) \right] + \sum_{y \in S(X)} q(X, y) \cdot \left[\frac{\partial}{\partial r_1} V(y) - \frac{\partial}{\partial r_1} V(X) \right], \end{aligned}$$

which further imply that:

$$\left[\rho + \sum_{y \in S(X)} q(X, y) + \lambda_X^* \cdot B^*(X) \right] \cdot \left[\frac{\partial}{\partial r_1} V^0 - \frac{\partial}{\partial r_1} V(X) \right] = \lambda_X^* - \lambda_0^* + \sum_{y \in S(X)} q(X, y) \cdot \left[\frac{\partial}{\partial r_1} V^0 - \frac{\partial}{\partial r_1} V(y) \right],$$

and the claims follow. Q.E.D.

Appendix D: Estimating an Equilibrium under Asymmetric, Imperfect Information

We substantially relax the informational assumptions of our Markov stationary equilibrium and, while employing a substantially different framework and methodology, arrive at estimates consistent with those from our primary structural model.

Specifically, we no longer model customers as possessing knowledge of a cardigan’s market state, which includes its price level, inventory, and market demand. Instead, we provide that by visiting the retailer’s website customers can individually update their private knowledge about the cardigan’s current price level and observe whether it remains in stock; in addition, they make time-dependent inferences about: *(i)* the cardigan’s price level (between visits); *(ii)* market demand for the cardigan; and *(iii)* the cardigan’s probability of stock-out. Naturally, customers do know at all times whether a cardigan has depreciated in its utility value to them.

To tractably accommodate this relaxation, we are forced to compromise on customer rationality. In general dynamic games with asymmetric and incomplete information (e.g., perfect Bayesian), players possess uncertain beliefs not only about the game’s state history, which alone pose computational challenges, but also about the history of beliefs held by each of the other players (which include their own beliefs about others’ such beliefs, and so on ad infinitum). It is widely acknowledged that this theoretical literature has yet to yield dynamic equilibrium concepts applicable to empirical settings as rich as ours.

Relative to our Markov stationary equilibrium, two salient modifications distinguish our approach. First, we draw on [Fershtman and Pakes \(2012\)](#) to define an equilibrium concept founded on a form of bounded rationality, wherein a customer’s beliefs are defined over her information set that may be coarser than the state space required (e.g., by a MPE) to fully distinguish her expected payoffs. [Fershtman and Pakes \(2012\)](#) allow the player’s (i.e., customer’s) information set to additionally include “informationally relevant” (rather than payoff-relevant) states that are indirectly useful to her in inferring her expected payoffs. In equilibrium, a customer’s beliefs about her expected payoffs over her information set are required to be consistent with her “experience,” i.e., under equilibrium play, the customers accurately project (and are best-responding to) the expected payoffs they would receive in each state in the information set under the equilibrium strategy profile (which, along with best responses, are restricted by the assumption that customers’ strategies condition only on the states of the specified information set).

For our model, we specify that a customer’s information set at any given time includes: *(i)* the cardigan’s price level at the time of her last visit; *(ii)* whether, at the time of her last visit, she had observed this price level before; *(iii)* whether the cardigan remained in stock as of her last visit; *(iv)* whether the cardigan has depreciated; and *(v)* the time elapsed since her last visit. Importantly, the customer’s strategy is no longer stationary with respect to time, the passage of which may at times incentivize the customer to visit more frequently to keep her information current or less frequently if the elapsed time has substantially raised the probability of stock-out. Intuitively, we model the customer as tracking and conditioning her behavior on these states, while in equilibrium her beliefs for these tracked states must be consistent with the average outcomes she would receive in them. In practice, simulation is used to compute these averages.

Second, our estimation approach modifies Arcidiacono and Miller (2011), which applies the EM algorithm to obtain Aguirregabiria and Mira (2007)’s nested pseudo likelihood (NPL) estimator for MPE under unobserved heterogeneity. The NPL estimator is an M-estimator solving a constrained version of the optimization used to obtain the MLE. More precisely, for MPE the MLE solves:⁴¹

$$\left\{ \hat{\theta}_{\text{MLE}}, \hat{P}_{\text{MLE}} \right\} := \operatorname{argmax}_{(\theta, P) \in Z} \log L_n(\theta, P) \quad (\text{EC.16})$$

$$Z := \{(\theta, P) : \theta \in \Theta, P = \Upsilon(\theta, P)\}, \quad (\text{EC.17})$$

where θ is the parameter vector of interest, P the (candidate) equilibrium play, and Υ the alternative best-response mapping, with P and Υ formally defined in Aguirregabiria and Mira (2007). Intuitively, the MLE is obtained by maximizing the log likelihood subject to the equilibrium constraint in (EC.17), which imposes that equilibrium play (as the P on the left-hand side) is consistent with players best-responding (under the model specification and θ) to beliefs consistent with such equilibrium play (captured by Υ and its inputs, θ and P , on the right-hand side). In contrast, the NPL solves:

$$\left\{ \hat{\theta}_{\text{NPL}}, \hat{P}_{\text{NPL}} \right\} := \operatorname{argmax}_{(\theta, P) \in Y} \log L_n(\theta, P) \quad (\text{EC.18})$$

$$Y := \{(\theta, P) : P = \Upsilon(\theta, P), \theta = \operatorname{argmax}_{\theta' \in \Theta} \log L_n(\theta', P)\}. \quad (\text{EC.19})$$

The NPL estimator is typically obtained by iteratively applying updates from each of the two constraints in (EC.19) (in alternating sequence to convergence), by updating the constraint’s left-hand side variable to satisfy the constraint at its current right-hand side iterates (i.e., within each iteration, we alternately obtain the best response P given the current vector θ and prior iterate P , then solve for the maximizer θ given the updated P). We do not elaborate on conditions for NPL convergence, e.g. Lyapunov stability, discussed by the subsequent literature. Arcidiacono and Miller (2011) then apply the EM algorithm to execute the maximization step within each iteration of the NPL procedure described above.

We adopt the NPL approach but modify the best-response constraint, $P = \Upsilon(\theta, P)$, with the “experience”-based requirement that customers best respond to beliefs over our specified information set. While notationally equivalent, the modified constraint captured by Υ imposes that customers best respond: *(i)* over the set of strategies restricted to the customer’s information set (in particular, their actions cannot be conditioned on states not in the information set); and *(ii)* their beliefs’ expected payoffs in the states of the information set are consistent with equilibrium play P . Here, we suppress a detail specific to the case of unobserved heterogeneity: we solve the Υ constraint separately at each value of our model’s latent random variables or in practice for all draws from the MCMC E-Step described below. Then, the modified NPL estimator is a fixed point of the following EM-based iteration over k :

$$\hat{P}_{k+1} = \Upsilon(\hat{\theta}_k, \hat{P}_k) \quad (\text{EC.20})$$

$$\hat{\theta}_{k+1} = \operatorname{argmax}_{\theta \in \Theta} \mathbb{E}_{X|Y, \hat{P}_{k+1}, \hat{\theta}_k} \left\{ \log L(Y, X | \hat{P}_{k+1}, \theta) \right\}, \quad (\text{EC.21})$$

where Y and X denote the observed and latent variables, respectively, of our model. As in Arcidiacono and Miller (2011), we swap the order of conducting update (EC.20) and the E-Step portion of update (EC.21)

⁴¹ Since the context is well-known, we omit typical regularity conditions and formalities, e.g., Θ compact.

in our implementation without affecting the algorithm's fixed point. Again, we suppress that updating \hat{P} involves updating best responses at all drawn values of X . For our model, update (EC.20) essentially amounts to updating beliefs \hat{P}_{k+1} to match estimates of transition probabilities under the best responses to the previously proposed equilibrium-play transition probabilities, \hat{P}_k . We obtain nonparametric, sieve-logit estimates of these transition probabilities from simulating the market process under \hat{P}_k and best responses under $\hat{\theta}_k$ to \hat{P}_k .

To initialize the described algorithm with (candidate) equilibrium play \hat{P}_0 , we use sieve-based techniques to nonparametrically estimate customers' equilibrium policies as observed in the data for each state in the information set and then simulate out the implied transition probabilities. Since the NPL estimation procedure does not actually require a consistent initial point, we omit a detailed exposition here.

In sum, we have described a framework to tractably estimate a dynamic equilibrium while considerably weakening the assumptions made about the information available to customers under our Markov stationary equilibrium. Moreover, the methodology we develop here is substantially different. Despite these differences, Table EC.7 shows that our point estimates remain close to the MLE under our Markov stationary equilibrium.

Table EC.7 Nested Pseudo Likelihood Point Estimates for the Experience-based Equilibrium Model

	Bargain Hunters	High Valuation
Additional US dollar compensation per visit required to increase monitoring rate by one online visit per month $\frac{1}{30} \cdot \frac{\hat{r}_2}{\hat{\alpha}}$	\$2.42	\$22.35
Elasticity $\hat{\alpha}$	0.02	0.01
Monitoring cost \hat{r}_1	0.01	0.29
Monitoring cost \hat{r}_2	1.60	7.59
Heterogeneous valuation $\hat{\gamma}$	0	0.87
Mean other-purchase payoff \hat{w}	-2.67	-2.27
Population Share	20%	80%

Sample size $N = 98$ cardigan products.

Appendix E: Additional Robustness Checks and Counterfactuals

E.1. Number of Price Levels

We present findings to explain why a retailer may limit its markdown policy to a small number of price levels: the returns to scheduling additional price levels diminish rapidly to be outweighed by the concomitant complexity for both consumers and the retailer.

In our setting, the retailer’s randomized markdown policy accommodates three price levels: list, sale, and clearance. As discussed earlier, a cardigan’s sale and clearance price levels are set at relatively predictable percentages of its list price, permitting customers to form stable expectations about the prices they will see. The relative simplicity of this pricing scheme underlies its credibility, since customers would easily “catch on” in the event that the retailer begins to deviate from their expectations. In turn, Section 6 highlights that this credibility carries important inventory and profit implications for retailers.

At the same time, optimizing over a greater number of price levels would in principle raise the retailer’s profits. However, with too many price levels, a markdown policy may no longer be sustained as credible, since a complex policy makes it more difficult for customers: *(i)* to derive accurate expectations on the pricing policy from their experiences and observations; and *(ii)* then given expectations that capture the retailer’s policy, to discern and detect systematic deviations. Notably, if customers cannot reliably detect when the retailer deviates, the retailer’s commitment becomes less credible to them. Therefore, in adding a markdown price level, the retailer must trade off its potential profit gained against these complexity costs.

As shown in Table EC.8, the retailer’s marginal profit gained from adding an additional price level to the markdown policy diminishes rapidly, explaining why retailers such as ours commit to markdown schedules featuring only a small number of price levels. In contrast, inventory effects from altering the number of price levels involved are negligible in either case.

Table EC.8 Profit Gains from Adding a Markdown Price Level

Number of Price Levels	Profit Change	Inventory Change
Three instead of Two	+10.2%	+0.6%
Four instead of Three	+3.9%	+0.6%

The counterfactual uses the 98 cardigans launched in calendar 2011. Inventory and pricing are jointly optimized. The retailer’s state-independent markdown hazard rates are adjusted to preserve the expected time to clearance (i.e., the time at which the clearance price level comes into effect). “Inventory Change” is the change in the USD sum of the optimal inventory levels’ marginal costs.

E.2. Three Customer Segments

We estimate our model with three customer types instead of two, and our results corroborate our finding of an inverse correlation between customers’ price elasticities and monitoring costs. Moreover, we do not find the three-type model to better fit the data across likelihood-based criteria (which take into account the greater number of parameters). Computationally, accommodating three types appears to be the upper limit

on segmentation for which we can currently obtain point estimates. The estimation results for the case with three customer types are summarized in Table EC.9.

Table EC.9 MLE for Structural Model with Three Types

	Bargain Hunters 1	Bargain Hunters 2	High Valuation
Additional US dollar compensation per visit required to increase monitoring rate by one online visit per month $\frac{1}{30} \cdot \frac{\hat{r}_2}{\hat{\alpha}}$	\$2.32	\$5.92	\$25.88
Elasticity $\hat{\alpha}$	0.02	0.02	0.01
Monitoring cost \hat{r}_1	0	0.10	0.36
Monitoring cost \hat{r}_2	1.61	3.69	7.61
Heterogeneous valuation $\hat{\gamma}$	0	0.30	0.87
Mean other-purchase payoff \hat{w}	-2.66	-2.52	-2.23
Population Share	12%	16%	71%

Sample size $N = 98$ cardigan products. The price-elasticity point estimate is largest for “Bargain Hunters 1”.

E.3. Variation in Cardigan Sizes

This subsection uses the data to rule out size as a motivation for the retailer’s markdown policy. As recommended by our retailer and conforming with its usual practices, our analysis aggregates the brand’s SKUs to the product “style” level. Consequently, our model is designed to relate pricing, inventory, and consumer monitoring as observed for cardigan products that are defined in this way from our data. Whereas our paper focuses on this important interaction, a more granular exploration of inventory management could explicitly track and manage additional product attributes, such as size.

For size in particular, theory allows that a retailer’s markdown policy may be motivated to effectively charge lower prices for a product’s unpopular sizes, which could tend to remain available during the sale or clearance periods. While we do not accommodate sizes in our structural model, we can test whether sales of the less popular “extreme” sizes (i.e., ‘XS’ or ‘XL’) tend to occur later in time for the cardigans in our estimation sample.

We can test this by assigning for each cardigan, an integer rank (from 1 to N) to each of its observed N sale transactions under chronological order by date. After conducting this assignment, we are interested in the ranks assigned to purchases of extreme sizes. If these sizes tend to be purchased later in the selling season, the ranks for this subset of transactions will tend to be larger than if the tendency were absent. Consequently, the subset’s sum of ranks would likewise tend to be larger.

More formally, we carry out a Mann-Whitney-Wilcoxon rank-sum test under the null hypothesis of a common distribution of ranks (i.e., across all purchases regardless of the purchased size) against the alternative that the extreme sizes’ chronological ranks stochastically dominate (i.e., tend later than) those of the popular sizes. Under the null hypothesis, the extreme-size purchases’ rank-sum statistic is distributionally invariant

under the algebraic group comprising all rank permutations over the sample. This test and observation extend naturally to the case of many cardigans instead of one: under the same null hypothesis, the direct (i.e., Cartesian) product across cardigans of these cardigan-specific permutation groups is itself an algebraic group, under which the vector of cardigans’ rank sums (with each sum calculated from a separate assignment of ranks to that cardigan’s purchases only) is distributionally invariant. Our test statistic is then simply the sum of the cardigans’ individual extreme-size rank sums (each constructed as above) — this quantity tends larger under the alternative. The one-sided tests we conduct, which reject the null hypothesis for large values of the test statistic, are exactly level, even in finite sample.

Table EC.10 presents our results using the 94 cardigans from calendar 2011 for which sizes ‘XS’ and ‘XL’ were offered. For the one-sided alternative hypothesis of sales occurring later for these extreme sizes, we find a highly insignificant test result (in fact suggesting that if anything the sales of these sizes tend to occur earlier in time). We can also examine the distribution of p-values from the rank-sum tests conducted for each cardigan separately. Under the null hypothesis, these p-values would be distributed uniformly over the interval from zero to one, and at the upper end of the distribution the share of p-values appear to be consistent with this hypothesis. Interestingly, at the lower end there appears a clustering of cardigans for which the extreme sizes surprisingly tend to sell (out) earlier in the season.

Therefore, for our retailer’s cardigans we find it highly unlikely that size-based price discrimination explains the retailer’s markdown policy or our model’s empirical findings. Instead, for some cardigans the evidence suggests that the extreme ‘XS’ and ‘XL’ sizes may sell earlier than the standard sizes; whether and how this phenomenon might relate to detailed inventory management and availability are questions falling outside the scope of our paper.

Table EC.10 Mann-Whitney-Wilcoxon Rank-Sum Tests for Cardigan Sizes and Purchase Times

	Rank Sum	P-value
Overall Rank Sum	83,948	Greater than 95%
Cardigans’ Individual Rank-Sum Tests		
Share of One-Sided Tests Rejecting Null at:		
5% Level	3.2%	
10% Level	5.3%	
20% Level	18.1%	
50% Level	38.3%	
80% Level	66.0%	
90% Level	79.8%	

Sample size $N = 94$ cardigan products for which ‘XS’ and ‘XL’ sizes are offered. The omitted cardigans offer only the ‘S’, ‘M’, and ‘L’ sizes, or alternatively the intermediate sizes ‘S/M’ and ‘M/L’. Rank (by chronological date over all purchases) sums are calculated for the ‘XS’- or ‘XL’-sized subset of the purchases made for each cardigan, with these sizes constituting 1372 observations out of 7231 overall purchases of the cardigans. The ‘Overall Rank Sum’ is the sum of these individual rank sums across the 94 cardigans. Tests are of the null hypothesis of a common distribution against the alternative hypothesis that the ranks of the extreme-size subset of purchases stochastically dominate (i.e., that purchases for extreme sizes tend to occur later in time). P-values are calculated using 100,000 randomly drawn rank permutations per cardigan, and such tests are exactly level, even for finite samples (and finite random draws).

Appendix F: Comparing Random with Deterministic Markdown Times

We present a simple pricing model to compare a seller’s expected revenues under committed prices with and without randomization. Our key finding is that the seller profitably randomizes in the presence of monitoring costs, otherwise (even with customers that discount for time) preferring a deterministic price schedule.

Consider a total unit mass of forward-looking customers, composed of two types. A fraction π_H has a high valuation v_H for the product, while fraction $\pi_L = 1 - \pi_H$ has a lower valuation $v_L < v_H$. Each customer has a single-unit demand (so the aggregate demand also equals one) and decides at each discrete time period whether to purchase or wait (we also allow for mixed strategies). The timing of the game is as follows: the retailer commits to a pricing policy and then customers decide their strategies after observing the retailer’s action. Both customers and the retailer discount future payoffs at rate $\beta \in (0, 1)$ per period. Finally, the season’s finite starting inventory level is set to $I \in (\pi_H, 1)$, i.e., it is set at a level that is lower than the aggregate demand but higher than the demand of high valuation customers.

At the start of the season, the retailer may choose between two pricing policies, each of which features a single markdown. The first deterministic option involves committing to two prices, retail p_R and clearance p_C , and a positive integer period t . The retail price is offered to customers for the first t time periods, after which there is a single clearance period during which the clearance price is offered. When demand exceeds inventory, allocation is random. For congruence with clearance sales in our setting, we will impose the constraint that $p_C = v_L$ in order to sell all units.

The second pricing policy likewise involves two prices but the timing of the markdown is random. As before, the retailer commits to retail and clearance prices, in addition to a daily markdown probability q . Each period following the first is a retail period with probability $1 - q$ and the final clearance period with probability q . The retail price p_R is offered in retail periods, while the clearance price $p_C = v_L$ is offered for the final clearance period. Given a period t , we denote by $q_t := \frac{1}{t}$ the daily markdown probability such that the expected time to markdown is t .

Finally, to capture monitoring costs in a very simple way, we will suppose that the high valuation customer type has a probability p of being present on any given day following the first — a lower probability corresponds to costlier monitoring. Higher monitoring costs may also deter the customers’ overall level of engagement with the retailer. To capture this effect, we will take $g(p)$ to be the fraction of high valuation types participating in the market, as a function of p that is nondecreasing with $g(0) = 0$ and $g(1) = 1$.

We start with the following proposition showing that deterministic markdowns are preferred by the retailer when the high value type faces zero costs to monitor ($p = 1$):

PROPOSITION 5. *Suppose $p = 1$ and fix any t and initial inventory level I . The optimal announced pricing schedule with clearance time t , $A = \{p_R^{A*}, p_C = v_L, t\}$, yields a higher revenue than the optimal randomized pricing schedule with expected clearance time t , $B = \{p_R^{B*}, p_C = v_L, q_t\}$.*

Proof: It is straightforward to derive the equilibrium for each policy such that customers either purchase in the first period or wait for clearance. In particular, the equilibrium corresponding to each pricing policy shares the same structure: high value customers follow a mixed strategy in the first period after which those that did not buy then wait for the clearance period.

Define α_A^* and α_B^* to be the respective mixed strategy probabilities in equilibrium of deciding to buy rather than wait in the first period. Under the policy that features a deterministic markdown, i.e., policy A , $(1 - \alpha_A^*) \cdot \pi_H$ high value customers wait for the clearance period to obtain a unit at the clearance price with the individual allocation probability, $\frac{I - \alpha_A^* \cdot \pi_H}{1 - \alpha_A^* \cdot \pi_H}$. We can make an analogous statement for the policy B . The equilibrium conditions dictate that:

$$\frac{v_H - p_R^A}{v_H - v_L} \cdot \left[\frac{1}{\beta} \right]^t = \frac{I - \alpha_A^* \cdot \pi_H}{1 - \alpha_A^* \cdot \pi_H}, \quad (\text{EC.22})$$

$$\frac{v_H - p_R^B}{v_H - v_L} \cdot \frac{1 - \beta \cdot (1 - q_t)}{q_t \cdot \beta} = \frac{I - \alpha_B^* \cdot \pi_H}{1 - \alpha_B^* \cdot \pi_H}. \quad (\text{EC.23})$$

Observe in both cases that when more customers purchase earlier, the value of waiting declines. Therefore the seller can exert maximum pricing leverage (and charge the highest price) in the equilibrium where all high value customers purchase early. Therefore, setting $\alpha_A^* = \alpha_B^* = 1$, the retailer can charge:

$$p_R^{A*} = v_H - \frac{I - \pi_H}{1 - \pi_H} \cdot \beta^t \cdot (v_H - v_L), \quad (\text{EC.24})$$

$$p_R^{B*} = v_H - \frac{I - \pi_H}{1 - \pi_H} \cdot \frac{q_t \cdot \beta}{1 - \beta \cdot (1 - q_t)} \cdot (v_H - v_L). \quad (\text{EC.25})$$

From our argument above, it is sufficient to show that $p_R^{A*} \geq p_R^{B*}$ to complete our proof. From (EC.24) and (EC.25), a necessary and sufficient condition is the following inequality:

$$\frac{1 - \beta \cdot (1 - q_t)}{q_t \cdot \beta} \leq \left[\frac{1}{\beta} \right]^t \iff \frac{1}{q_t \cdot \beta} - \frac{1 - q_t}{q_t} - \left[\frac{1}{\beta} \right]^t \leq 0. \quad (\text{EC.26})$$

At $\beta = 1$, this inequality holds with equality. The partial derivative with respect to β of the left side of inequality (EC.26) is:

$$\frac{t}{\beta^2} \cdot \left(\left[\frac{1}{\beta} \right]^{t-1} - 1 \right) \geq 0,$$

over the joint domain of the relevant parameters, hence the necessary and sufficient inequality condition holds. Q.E.D.

In the presence of monitoring costs on the other hand, we obtain the following result (the proof is omitted as it follows from similar arguments as in the proof of Proposition 5).

PROPOSITION 6. *Fix any t and initial inventory level I . Suppose $p \in (0, 1]$. The optimal announced price schedule with clearance time t , $A = \{p_R^{A*}, p_C = v_L, t\}$, yields revenue given by:*

$$\Pi_A(p) = \pi_H \cdot \left[v_H - \frac{I - \pi_H}{1 - \pi_H} \cdot \beta^t \cdot (v_H - v_L) \right] + v_L \cdot (I - \pi_H).$$

The optimal randomized markdown schedule with expected clearance time t , $B = \{p_R^{B}, p_C = v_L, q_t\}$, yields revenue given by:*

$$\Pi_B(p) = g(p) \cdot \pi_H \cdot \left[v_H - \min \left\{ 1, \frac{I - g(p) \cdot \pi_H}{1 - \pi_H} \right\} \cdot p \cdot \frac{q_t \cdot \beta}{1 - \beta \cdot (1 - q_t)} \cdot (v_H - v_L) \right] + v_L \cdot \min \{ 1 - \pi_H, I - g(p) \cdot \pi_H \}.$$


Higher monitoring costs impose two countervailing effects on the seller's revenue under randomized markdowns. First, the costs diminish the high valuation customer's prospective reward from waiting for clearance, because of the risk of missing the timing. By making it less likely that the high valuation type will postpone

buying at any given price, this allows the seller to collect higher revenue in the first time period. On the other hand, higher monitoring costs reduce engagement as captured by the function g .


Accordingly, if the associated reduction in engagement is relatively weak, higher monitoring costs (again, represented by a lower monitoring probability p) can make randomized markdowns the seller's most profitable strategy. A numerical example is shown in Figure [EC.11](#), where panel [EC.11c](#) shows an intermediate case where customer engagement falls in response to monitoring costs but randomized markdowns are valuable.

Figure EC.11 Numerical Comparison of Revenues under an Announced Price Schedule and Randomized Markdowns


$t = 100, \beta = 0.999$ per day, $v_H = 10, v_L = 6, I = 0.85$ — **Randomized Markdowns in Solid-Line Plot**



(a) No loss of engagement — $g(p)=1$ for $p > 0$



(b) Strong engagement response to monitoring costs — $g(p) = p^{0.1}$



(c) Intermediate case — $g(p) = p^{1/30}$