This article was downloaded by: [128.143.1.14] On: 03 November 2016, At: 11:42 Publisher: Institute for Operations Research and the Management Sciences (INFORMS) INFORMS is located in Maryland, USA



Operations Research

Publication details, including instructions for authors and subscription information: <u>http://pubsonline.informs.org</u>

Behavioral Anomalies in Consumer Wait-or-Buy Decisions and Their Implications for Markdown Management

Manel Baucells, Nikolay Osadchiy, Anton Ovchinnikov

To cite this article:

Manel Baucells, Nikolay Osadchiy, Anton Ovchinnikov (2016) Behavioral Anomalies in Consumer Wait-or-Buy Decisions and Their Implications for Markdown Management. Operations Research

Published online in Articles in Advance 03 Nov 2016

. http://dx.doi.org/10.1287/opre.2016.1547

Full terms and conditions of use: <u>http://pubsonline.informs.org/page/terms-and-conditions</u>

This article may be used only for the purposes of research, teaching, and/or private study. Commercial use or systematic downloading (by robots or other automatic processes) is prohibited without explicit Publisher approval, unless otherwise noted. For more information, contact permissions@informs.org.

The Publisher does not warrant or guarantee the article's accuracy, completeness, merchantability, fitness for a particular purpose, or non-infringement. Descriptions of, or references to, products or publications, or inclusion of an advertisement in this article, neither constitutes nor implies a guarantee, endorsement, or support of claims made of that product, publication, or service.

Copyright © 2016, INFORMS

Please scroll down for article-it is on subsequent pages



INFORMS is the largest professional society in the world for professionals in the fields of operations research, management science, and analytics.

For more information on INFORMS, its publications, membership, or meetings visit http://www.informs.org



Behavioral Anomalies in Consumer Wait-or-Buy Decisions and Their Implications for Markdown Management

Manel Baucells

Darden School of Business, University of Virginia, Virginia 22903 baucellsm@darden.virginia.edu

Nikolay Osadchiy

Goizueta Business School, Emory University, Atlanta, Georgia 30322 nikolay.osadchiy@emory.edu

Anton Ovchinnikov

Smith School of Business, Queen's University, Kingston, ON K7L 3N6, Canada, anton.ovchinnikov@queensu.ca

The decision to buy an item at a regular price or wait for a possible markdown involves a multidimensional trade-off between the value of the item, the delay in getting it, the likelihood of getting it, and the magnitude of the price discount. Such trade-offs are prone to behavioral anomalies by which human decision makers deviate from the discounted expected utility model. We build an axiomatic preference model that accounts for three well-known anomalies and produces a parsimonious generalization of discounted expected utility. We then plug this behavioral model into a Stackelberg-Nash game between a firm that decides the price discount and a continuum of consumers who decide to wait or buy, anticipating other consumers' decisions and the resultant likelihood of product availability. We solve the markdown management problem and contrast the results of our model with those under discounted expected utility. We analytically show that accounting for the behavioral anomalies can result in larger markdowns and higher revenue. Finally, we calibrate our model via a laboratory experiment and validate its predictions out-of-sample.

- *Keywords*: retail operations; markdown management; probability and time tradeoff; behavioral anomalies; behavioral operations.
- Subject classifications: decision analysis: applications; inventory/production: policies: marketing/pricing; utility/preference: applications.

Area of review: Operations and Supply Chains.

History: Received October 2014; revisions received December 2015, May 2016; accepted June 2016. Published online in *Articles in Advance* November 3, 2016.

1. Introduction

Many consumers around the globe regularly make the following decision, i.e., buy an item now at the tag price, p, or wait until time t when the product will be marked down to p(1 - d), but may only be available with probability q. Understanding how consumers make such *wait-or-buy decisions* is crucial for retailers to properly optimize markdowns which, in turn, is critical for retailers' profitability (Agrawal and Smith 2009).

Historically, retailers optimized markdowns as if consumers were myopic (Kalish 1983), that is, fail to take into account the opportunity of future discounts. In this case, retailers were inclined to offer large markdowns (e.g., 50% or more) to expand the market. Such practice effectively taught consumers to wait, reduced retailers' profitability, and led to a body of literature on pricing with strategic consumers. Much of this literature uses the rational framework of discounted expected utility (DEU), and suggests that markdowns should be reduced or even eliminated in favor of everyday low prices (Besanko and Winston 1990, Ortmeyer et al. 1991, Aviv and Pazgal 2008). Observed behaviors, however, deviate from myopia and DEU. The goals of this paper are to build on the observed deviations, which we call *anomalies*, propose a new behaviorally-grounded model of wait-or-buy decisions, and study its implications for mark-down management.

From the consumer's viewpoint the wait-or-buy decision is a multidimensional trade-off between the value of the item, the price discount, the likelihood of getting the item, and the delay in getting it. DEU, while directionally correct, treats each dimension linearly and separately. In Section 2 we present evidence that consumers view these trade-offs in a nonlinear and interdependent way. For instance, consumers' sensitivity to the risk of not obtaining an item depends on the time delay and the magnitude of the discount; conversely, sensitivity to a time delay depends on the price discount and risk.

We capture this interdependency through the notion of *psychological distance* that (Baucells and Heukamp 2012) associated with the prospects of "buy now for sure" (zero

distance) and "perhaps buy later" (time and risk distance). In Section 3 we propose preference conditions that account for three robust behavioral anomalies, i.e., the common ratio effect describing the nonlinearity in sensitivity to risk; the common difference effect, describing the nonlinearity in sensitivity to time; and subendurance, describing how sensitivity to time depends on the magnitude of payoff. Models such as prospect theory or hyperbolic discounting account for these anomalies separately, but there appears to be no preference model that accounts for these anomalies simultaneously, which is necessary to describe a wait-or-buy decision. Rather than simply guessing such a model (and this is one of the key contributions of our paper) we set preference conditions (axioms) that capture these anomalies and derive the new model. We call our model discount, probability, and time trade-off (dPTT). The dPTT model is a parsimonious generalization of DEU.

Using dPTT, we optimize a retailer's markdown (a.k.a., price discount). To do so, in Section 4 we consider a standard Stackelberg-Nash game between the firm and a continuum of consumers. The firm announces the price discount. Given the price discount, consumers (Stackelberg followers) anticipate other consumers' decisions, as well as the resultant probability of product availability, which is endogenously determined in the Nash equilibrium. Anticipating this equilibrium, in Section 5 the retailer (Stackelberg leader) solves the markdown optimization problem.

Our main result is that, compared with the DEU benchmark, accounting for the behavioral anomalies can lead to larger optimal markdowns and larger revenues. We further explore markdown timing and quantity optimization. According to dPTT, and relative to DEU, markdowns should be offered sooner and capacity rationing could be optimal under broader conditions. Our main result agrees with Smith and Achabal (1998), i.e., retailers should offer larger markdowns based on the dependency between demand and the remaining available inventory; and Özer and Zheng (2016) who argue that dynamic pricing is even more valuable than previously thought due to the nonpecuniary behavioral factors such as consumers' regret and misperceptions of product availability. We reach the same conclusion based on the consumers' psychological perceptions of time, risk, and price discount.

To understand our result, recall that the optimal markdown is selected to balance the marginal revenue from selling more units at the markdown price with the marginal cost of diverting consumers from buying at the tag price. The behavioral anomalies we study affect this balance in two ways. First, subendurance implies that consumers are less patient for small markdowns and more patient for large ones. At the markdown level optimal for DEU, the former effect dominates. Thus, the retailer exploits this impatience and offers larger markdowns without sacrificing sales at the tag price. Second, the nonlinearities in risk and time perception imply that consumers are more sensitive to psychological distance than the DEU model assumes when distances are small, but are less sensitive when distances are large. Increasing the markdown increases demand, which increases product availability risk and therefore the psychological distance. Thus, consumers who, at the DEU-optimal markdown, were "buying now" (i.e., had a zero distance) are very sensitive to increased distance; consequently, most continue to buy now. Likewise, most of those who were waiting continue to wait as they are less sensitive to the increased distance since waiting implies a positive psychological distance to begin with. The subendurance and nonlinearity effects complement each other and allow a retailer to offer larger markdowns and gain additional revenue.

To assess realistic values of markdown increases and revenue gains, in Section 6 we elicit model parameters through a laboratory experiment. To ensure the quality of the data, we used binary questions and choice lists (Holt and Laury 2002) in which participants face a battery of wait or buy choices under different price discount and risk scenarios. We observe that subjects' responses reveal clear indifference points at which subjects switch from buying to waiting. We therefore use indifference points to fit our model. To elicit responses, we use an incentive-compatible method called the Prior Incentive (Prince), which is a refined version of the randomized incentive scheme. In Prince, one choice is randomly pre-selected and given to a participant in a sealed envelope before the experiment. If the subject is selected, then the choice in the envelope is played for real after the experiment. Prince has been shown to improve the quality of elicitation (Johnson et al. 2014).

After estimating the parameters, we optimize the retailer's markdown, factoring in the associated consumer wait-or-buy equilibrium. The resulting optimal markdown under dPTT is 18.6% as opposed to 10.7% prescribed by DEU, and the resulting revenue is 1.5% larger. These differences are robust with respect to various parameters and errors in their estimation.

It is important to put the dPTT results into perspective. First, while the dPTT optimal markdown is almost twice as large as the DEU, it is still much smaller than the myopic value of 50%. Intuitively, a 10% markdown is too small. The DEU model overreacts to strategic waiting, and dPTT corrects that by recommending a markdown large enough to attract new consumers, yet small enough to control strategic waiting. Second, while the 1.5% revenue gain may seem small, a typical retailer operates with a net margin of approximately 3% (Damodaran 2015); increasing revenue by 1.5% without affecting costs translates into noticeable profits. Finally, much of the firm's revenue can be obtained by charging a single price, in which case the nuances of wait-or-buy behavior are irrelevant. Comparing only the markdown revenue, dPTT increases it by 25%-50%, clearly a substantial improvement in the effectiveness of markdowns.

We are not the first to consider behavioral motives in pricing and markdown decisions (see Özer and Zheng 2012, Chapters 2.1 and 3.1.2, for an excellent review). Therefore, two points deserve additional discussion, i.e., our reasons for choosing the specific behavioral anomalies, and how our results compare to previous studies that also consider behavioral factors. Our choice is based on the following criteria. The anomalies must: (i) be directly relevant to the wait-or-buy decision, (ii) have abundant and unambiguous experimental support, and (iii) affect the wait-or-buy decision ex ante. Many robust anomalies such as overconfidence, projection or confirming evidence bias do not seem a priori as relevant to the wait-or-buy decision. Because consumers can opt out of a loss by not purchasing, loss aversion is also not relevant.

Liu and van Ryzin (2008) found that risk aversion (in a form of concave transformation of payoffs) allows the retailer to increase revenue by rationing capacity and creating scarcity risk. In experiments, however, individuals rarely exhibit risk aversion by transforming payoffs only; rather, they seem to distort probabilities. Furthermore, consumers are, on average, risk-seeking over low-probability gains and high-probability losses and risk-averse over high-probability gains and low-probability losses (Tversky and Kahneman 1992). Our model captures all these patterns, qualitatively replicates the results of Liu and van Ryzin, and thus expands the justification for capacity rationing.

Nasiry and Popescu (2011) study the role of reference price, anchoring, and loss aversion to conclude that a constant price policy may often be preferred to dynamic pricing. The key assumption is that the reference price is a combination of the latest and lowest observed prices. The experimental evidence on reference price formation is still scarce, but recent research (Baucells et al. 2011) suggests that the lowest price has little influence, thus strengthening the case for dynamic pricing. Yet, their message of caution against deep discounts is consistent with dPTT.

Özer and Zheng (2016) found that markdowns can lead to revenue gains of comparable magnitude to ours if consumers exhibit regret and have probability misperceptions. While consumers do experience regret, the evidence as to whether they incorporate regret in their decisions ex ante is weak (Starmer and Sugden 1993). Probability misperceptions are also common, but can be incorporated as if consumers distort probability (which they do), or as if consumers distort the psychological distance that combines probability and delay (which we do). Both approaches capture similar behavior and, not surprisingly, yield qualitatively similar results.

Additional related behavioral features that have been studied include anecdotal reasoning (Huang and Liu 2014), uncertain product value (Swinney 2011), stockpiling and inertia (Su 2010, 2009), and reference dependence (Popescu and Wu 2007, Tereyagoglu et al. 2014). The innovativeness of our approach is that we provide a preference model of *how* consumers decide to wait or buy, in addition to studying the implications of such a model. Our paper adds to the empirical literature on strategic consumers, e.g., Mak et al. (2014), Kim and Dasu (2014), Li et al. (2014), and Osadchiy and Bendoly (2010). However, rather than using an experiment to motivate theory, our experiment is designed to calibrate a preference model built on a set of previously observed anomalies.

Admittedly, the dPTT model yields a refinement of the DEU results. The experimental basis for our model and its economic impact, however, are strong. Can we trust its conclusions? We run an out-of-sample experiment (Section 6.6) offering the DEU-recommended markdowns versus the dPTT markdowns, and find that revenue is higher under dPTT. Thus, contributions of our paper offer the full circle: Starting from observed behavioral anomalies, we construct a preference model that precisely captures them. Then, we embed this model into the markdown optimization problem and solve it. Finally, we calibrate the model via an experiment and verify its predictions out-of-sample, deriving relevant implications for markdown management.

2. The DEU Model and the Behavioral Anomalies it Fails to Explain

Much of the existing literature on markdown management (e.g., Besanko and Winston 1990, Aviv and Pazgal 2008, Liu and van Ryzin 2008, Zhang and Cooper 2008) uses the DEU model to solve the wait-or-buy problem. Let u denote the willingness to pay in the "buy now" case, which we call *the benefit of consumption*. By default, "buy now" ensures the purchase. According to DEU, "opt out" yields 0 utility, "buy now" yields u - p, and "wait" yields

$$U(p, d, q, t) = [u - p(1 - d)]qe^{-rt}.$$
 (1)

Here, r > 0 denotes the time discount rate, and q the probability that the product is available at time t. According to DEU, the consumer will opt out if u < p(1 - d), wait if $p(1 - d) \le u < H^{\text{DEU}}$, and buy now if $u \ge H^{\text{DEU}}$, where $H^{\text{DEU}} = p(1 - (1 - d)qe^{-rt})/(1 - qe^{-rt})$. The solution is intuitive: When u is high, the penalty for availability risk and/or discounting is high, prompting the consumer to buy now. In contrast to, a high price discount makes "wait" attractive, but this attractiveness is dampened if t is large or q is small. The cut-off point always satisfies $H^{\text{DEU}} \ge p$ because for consumers with $p(1 - d) \le u < p$ the only reasonable option is to wait.

We believe that DEU is directionally correct, i.e., consumers like price discounts and dislike availability risk and delay. However, DEU fails to account for three behaviorally important effects.

The first is the *common ratio effect* in risk preferences, by which the effect of the probability q is not linear in the mind of the consumer. A 20% change in the probability of the product being available from 100% to 80% has a much higher relative impact than the same 20% change from 10% to 8%. Consumers seem to be less sensitive to probability ratios when probabilities become small. The second is the *common difference effect* in time preference (a.k.a., hyperbolic discounting). Changing the delay from 0 (no delay) to 4 weeks has a higher relative impact than the same 4-week change from 26 to 30 weeks. Consumers seem to be less

Table 1 shows experimental evidence for these three anomalies. Pattern 1–2 replicates the common ratio effect and pattern 3–4 reproduces the common difference effect. These two are at the core of a considerable literature on nonlinear probability weighting (Allais 1953, Wakker 2010) and hyperbolic discounting (Laibson 1997, O'Donoghue and Rabin 1999, DellaVigna and Malmendier 2004). Pattern 5–6 illustrates subendurance. DEU is highly incompatible with these patterns. Indeed, pattern 1–2 is incompatible with linear probability weighting, and pattern 3–4 is incompatible with exponential discounting. Pattern 5–6 requires that time discounting be affected by the outcome dimension, while DEU assumes that the time discount rate, r, is fixed.

Our goal is to propose a DEU modification that better approximates how consumers feel about the trade-offs between price, price discounts, probabilities, and delays (i.e., all the patterns in Table 1). Rather than guessing a utility model, we propose preference conditions characterizing such a model.

Our model builds on the axiomatic preference framework for probability and time trade-offs by Baucells and Heukamp (2012). They showed that the total psychological distance can be viewed as a sum of time and risk distances, where time t can be used for the time distance "as is," while probability q should be log-transformed to $\ln(1/q)$ because probabilities multiply rather than add, and higher q's correspond to smaller distances. Thus, as they do, we assume that the risk and time distances are substitutes, that their "exchange rate," r, may depend on the outcome, and that consumers exhibit diminishing sensitivity to distance, i.e., the probability and delay penalty is a concave function of distance. That is, consumers are disproportionately sensitive to a small change from full and immediate availability to partial availability or small delay. Conversely, consumers are less sensitive to additional delays, or to increases in availability risk, if the prospect is in the future or not certain to begin with.

Extending the Baucells and Heukamp (2012) framework, the outcome in our model has two dimensions, i.e., price and price discount. We assume that it is the price discount that drives the subendurance effect. Simply put, consumers will be more patient if the price discount is high, which implies that r is a decreasing function of d. This assumption is consistent with Tversky and Kahneman (1981)'s observation that consumers are willing to travel 10 minutes to grab a 33% price discount on a calculator that costs \$15, but not willing to travel the same 10 minutes to grab a 5% price discount on a jacket that costs \$100. That the dollar discount is the same, \$5, shows that what drives customer acceptance of the delay of 10 minutes depends on the price discount percentage, not the dollar value. The importance of a percentage price discount itself has also been suggested by Thaler (1985) and corroborated by Darke and Freedman (1993). The assumption is psychologically plausible, as the price discount d is comparable across purchases. A more general model where r would depend on d and u seems unnecessarily complicated. The current model is also easier to calibrate, as d is directly observable, whereas *u* is not.

3. Price Discount, Probability, and Time Trade-off

Next, we propose a set of axioms (preference conditions) that characterize a preference relation capable of explaining the behavioral anomalies described above. Let τ denote the current calendar date. At time τ , the consumer exhibits preferences between pairs in $\mathscr{X}_{\tau} = [0, \infty) \times [0, 1] \times [0, 1] \times [\tau, \infty)$, where a typical element will be written as $x = (p_x, d_x, q_x, t_x) \in \mathscr{X}$. Here, p_x represents the tag price, d_x the price discount percentage, q_x the probability of the good being available, and $t_x \ge \tau$ is the purchase date. Each consumer desires one item. The benefit of consumption does not depend on τ . The availability of the item is revealed at time t_x .

A word on notation: We write (d, x_{-d}) or (q, t, x_{-qt}) to denote the vectors (p_x, d, q_x, t_x) or (p_x, d_x, q, t) , respectively. Throughout, "decreasing" implies "nonincreasing," otherwise, we use "strictly decreasing". The same holds for "increasing" or "concave." Finally, $\mathbf{0} = (0, 0, 0, 0)$.

Table 1.Rows 1–2 are taken from Baucells and Heukamp (2010, Table 1). Rows 3–4
are taken from Keren and Roelofsma (1995, Table 1) (1 fl or Dutch Gulden in
1995 = \$0.6). Rows 5–6 from Baucells et al. (2009).

Prospect A	vs.	Prospect B	Response (%)	Ν
 (9€, for sure, now) (9€, with 10%, now) 	vs.	(12€, with 80%, now)	58 vs. 42	142
	vs.	(12€, with 8%, now)	22 vs. 78	65
3. (100 fl, for sure, now)	VS.	(110 fl, for sure, 4 weeks)	82 vs. 18	60
4. (100 fl, for sure, 26 weeks)	VS.	(110 fl, for sure, 30 weeks)	37 vs. 63	60
 5. (5€, for sure, 1 month) 6. (100€, for sure, 1 month) 	vs.	(5€, with 90%, now)	43 vs. 57	79
	vs.	(100€, with 90%, now)	81 vs. 19	79

3.1. Axioms

Let \succeq_{τ} denote a preference ordering over pairs in \mathscr{X}_{τ} as expressed by a consumer from the point of view of τ . Our first condition guarantees the existence of a continuous function, $V_{\tau}(p, d, q, t)$, which represents such preferences.

AXIOM 1. For each $\tau \ge 0$, \succeq_{τ} is a complete, transitive, and continuous ordering on \mathscr{X}_{τ} .

The next condition states that preferences are not a function of calendar time, but a function of time relative to τ . It translates reference dependence (outcomes are not evaluated in absolute, but relative to a reference point) into the time dimension.

AXIOM 2 (TIME INVARIANCE). For all $x, y \in X$, $0 \leq \tau \leq t_x, t_y$, and $\Delta \ge 0$,

$$x \sim_{\tau} y$$
 if and only if $(t_x + \Delta, x_{-t}) \sim_{\tau+\Delta} (t_y + \Delta, y_{-t})$.

Time invariance implies that $V_{\tau}(p, d, q, t) = V_0(p, d, q, t - \tau)$. Hence, specifying the preferences from the viewpoint of $\tau = 0$ automatically determines the preferences from all time viewpoints. Henceforth, when we omit the subscript τ from \mathcal{X} , V, and \succeq , it means that $\tau = 0$.

Next, we impose monotonicity and solvability conditions. Null purchases, those having q = 0, are interpreted as no purchases and are deemed indifferent. The directional effects of price, price discounts, time, and probability are the same as in DEU. Finally, while the item is desirable for free, there is a finite price one is willing to pay.

AXIOM 3 (MONOTONICITY AND SOLVABILITY). Let

 $\mathscr{X}^0 = \{x \in \mathscr{X} : q_x = 0\}.$ For all $x \in \mathscr{X}$,

A3.0 if $x \in \mathcal{X}^0$, then $x \sim \mathbf{0}$;

A3.p let $p < p_x$. If $x \notin \mathcal{X}^0$, then $(p, x_{-p}) \succ x$;

- A3.q let $q > q_x$. If x > 0, then $(q, x_{-q}) > x$; and if x < 0, then $(q, x_{-q}) < x$;
- A3.t let $t < t_x$. If $x \succ \mathbf{0}$, then $(t, x_{-t}) \succ x$; and if $x \prec \mathbf{0}$, then $(t, x_{-t}) \prec x$; and
- A3.u there exist a $u \in (0, \infty)$ such that $(u, 0, 1, 0) \sim \mathbf{0}$.

Because A3.u is imposed for immediate sure purchases with no price discount, u does not depend on the dimensions of x, but only on the item itself. Specifically, by time invariance, u does not change with the passage of calendar time. For seasonal products, where preferences seem to depend on time, our model can be extended by introducing a state variable, e.g., hours of daylight, so that u depends on the state, but not on time per se. That is, u_{state} would still be time invariant, but *state* could depend on time. Otherwise, the effect of impatience cannot be disentangled.

Next, we assume that, for immediate purchases, price and price discount are rationally encoded in a way that only the effective price matters. The condition precludes framing effects, i.e., changes in preferences associated with simultaneously increasing the tag price and the price discount while keeping the same effective price. We acknowledge that framing effects may matter, but they are not the focus of our exercise. Note that the condition is imposed only on immediate purchases.

AXIOM 4 (EFFECTIVE PRICE CONDITION). For all $x, y \in \mathcal{X}$ such that $q_x = q_y$ and $t_x = t_y = 0$,

$$x \succeq y$$
 if and only if $p_x(1-d_x) \leq p_y(1-d_y)$.

The next condition links risk and time preferences (Baucells and Heukamp 2012). It describes the intuitive notion that "time is intrinsically uncertain." Thus, if a delay of $\Delta = 1$ month is exchangeable with a probability factor of $\theta = 80\%$, then a delay of $\Delta = 2$ months is exchangeable with a probability factor of $\theta^2 = 64\%$, and this exchange holds independently of the base level of time and probability. For any delay $\Delta > 0$, one can find a reduction in probability $\theta < 1$ (without the delay) that offsets the effect of Δ . The condition states that once this trade-off is established at some probability and time base level, it holds for all probability and time base levels. The condition does not extend to different price discounts because the probability and time trade-off may depend on d_x .

AXIOM 5 (PROBABILITY AND TIME TRADE-OFF). For all $x \in \mathcal{X}$, $p \ge 0$, θ , $q \in [0, 1]$, and Δ , $t \in [0, \infty]$,

$$(p_x, d_x, q_x, t_x + \Delta) \sim (p_x, d_x, q_x\theta, t_x)$$
 if and only if
 $(p, d_x, q, t + \Delta) \sim (p, d_x, q\theta, t).$

The indifference in Axiom 5 allows us to define the "exchange rate" between probability and time, r(d), such that $r(d) \cdot \Delta = \ln(1/\theta)$. By Axiom 5, r(d) does not depend on q, t, and p, but may depend on d (not so in the case of DEU). We refer to r(d) as the probability discount rate.

Axiom 5 is compatible with, but logically independent from, the three preference patterns exhibited in Table 1. To explain pattern 5–6, it is necessary to let the probability discount rate depend on the outcome dimension. To reduce degrees of freedom, and as explained in Section 2, we assume that what drives probabilistic patience is the price discount only.

Axiom 6 (Price Discount Subendurance). For all $x \in \mathcal{X}$, $\theta \in [0, 1)$, $\Delta \in [0, \infty]$, and $d > d_x$,

$$if x \succ \mathbf{0} and (t_x + \Delta, x_{-t}) \sim (q_x \theta, x_{-q}),$$

$$then (p_x, d, q_x, t_x + \Delta) \succeq (p_x, d, q_x \theta, t_x)$$

$$if x \prec \mathbf{0} and (t_x + \Delta, x_{-t}) \sim (q_x \theta, x_{-q}),$$

$$then (p_x, d, q_x, t_x + \Delta) \preceq (p_x, d, q_x \theta, t_x).$$

Axiom 6 implies that consumers become more patient for higher price discounts, i.e., r(d) is decreasing in d. This also ensures that a higher price discount makes an attractive product even more attractive.

The following condition considers pattern 1-2 (where a reduction in probabilities renders the prospect with the better outcome more attractive). It reflects a loss of sensitivity to risk distance. Together with Axiom 5, it implies a loss of sensitivity to time distance, i.e., pattern 3–4.

AXIOM 7 (SUBPROPORTIONALITY). Let $x, y \in \mathcal{X}$ with $q_x \leq q_y$, $t_x = t_y$. For all $\theta \in [0, 1]$,

if
$$x \sim y \succ \mathbf{0}$$
, then $(\theta q_y, y_{-q}) \preceq (\theta q_x, x_{-q})$; and
if $x \sim y \prec \mathbf{0}$, then $(\theta q_y, y_{-q}) \succeq (\theta q_x, x_{-q})$.

The next condition produces a simple structure by assuming separability between the price and nonprice dimensions (same as in DEU), but only for prospects received now at no discount. This is known as the hexagon condition, and it is a specialization of the *corresponding trade-off condition* (Keeney and Raiffa 1976, Theorem 3.2) for the case of two attributes.

AXIOM 8 (RESTRICTED PROBABILITY-PRICE SEPARABILITY). For all $x \in X$ with $t_x = 0$ and $d_x = 0$, $p, p', p'' \ge 0, q, q', q'' \in [0, 1]$, if three of the following indifferences holds, the fourth one holds as well.

$$(p,q',x_{-pq}) \sim (p',q,x_{-pq}), \quad (p',q',x_{-pq}) \sim (p'',q,x_{-pq}), (p,q'',x_{-pq}) \sim (p',q',x_{-pq}), \quad (p',q'',x_{-pq}) \sim (p'',q',x_{-pq}).$$

3.2. Representation

The representation below $[x \succeq_{\tau} y \text{ iff } V_{\tau}(p_x, d_x, q_x, t_x) \ge V_{\tau}(p_y, d_y, q_y, t_y)]$ will involve three continuous functions: a value function, $v: \mathbb{R} \to \mathbb{R}_+$, strictly increasing and with v(0) = 0; a psychological distance function, $s: \mathbb{R}_+ \to \mathbb{R}_+$, strictly increasing with s(0) = 0, s(1) = 1, and $s(\infty) = \infty$; and a probability discount function, $r: [0, 1] \to (0, \infty)$.

PROPOSITION 1. Preference ordering \succeq_{τ} on \mathscr{X}_{τ} satisfies A1–A8, if and only if, for some value function, v, two concave psychological distance functions, s^+ and s^- , and some decreasing and bounded probability discount function, r, \succeq_{τ} is represented by

$$V_{\tau}(p, d, q, t) = \begin{cases} v(u - p(1 - d)) \cdot e^{-s^{+}(\sigma)}, & u - p(1 - d) \ge 0; \\ v(u - p(1 - d)) \cdot e^{-s^{-}(\sigma)}, & u - p(1 - d) < 0, \end{cases}$$
(2)

where $\sigma = \ln(1/q) + r(d)(t-\tau)$ is the psychological distance of the prospect $(p, d, q, t) \in \mathscr{X}_{\tau}$.

All proofs are presented in the E-Companion (available as supplemental material at https://doi.org/10.1287/ opre.2016.1547). We call (2) the *dPTT model*, for (price) discount-probability-time trade-off. We note several properties of the representation.

1. dPTT collapses into DEU if v, s^+ , and s^- are the identity function and r(d) is constant.

2. For immediate purchases, dPTT agrees with a prospect theory like formulation in which v is a value function and $w(q) = e^{-s(-\ln q)}$ is a subproportional probability weighting function.

3. For future purchases with no availability risk, dPTT agrees with a hyperbolic discounting model in which $f(t) = e^{-s(t)}$ is a substationary time discount function.

4. Intuitively, if consumers exhibit diminishing sensitivity to risk distance (7), and risk and time distance are substitutes (5), then they should exhibit diminishing sensitivity to time distance (a.k.a., substationarity, or a delay makes the prospect with the better outcome more attractive, as seen in pattern 3–4 of Table 1). Also, monotonicity with respect to price A3.p and the effective price condition (4) imply monotonicity with respect to price to price discounts.

PROPOSITION 2. A1-A7 imply

(A7.t) Substationarity. Let $x, y \in \mathcal{X}$ with $t_x \ge t_y, d_x \ge d_y$, and $q_x = q_y$. For all $\Delta \ge 0$,

if
$$x \sim y \succ \mathbf{0}$$
 then $(t_y + \Delta, y_{-t}) \preceq (t_x + \Delta, x_{-t});$ and

if $x \sim y \prec \mathbf{0}$ then $(t_y + \Delta, y_{-t}) \succeq (t_x + \Delta, x_{-t})$.

(A3.d) For all $x \in \mathcal{X}$, if $x \succ \mathbf{0}$ and $d > d_x$, then $(d, x_{-d}) \succ x$.

5. dPTT has a minimum time consistency. Null purchases will be deemed indifferent at $\tau = 0$ and at any subsequent time $\tau > 0$. Moreover, with the passage of time, favorable deals will remain favorable, and unfavorable deals will remain unfavorable. This is important for the markdown problem because if a product is attractive today, but the consumer decides to wait, then the product will remain attractive in the future if it is available.

PROPOSITION 3. Assume A1–A4 and let $\mathbf{0}_{\tau} = (0, 0, 0, \tau)$, for some given $\tau \ge 0$.

• If $x \sim 0$, then $x \sim_{\tau} \mathbf{0}_{\tau}$; if $x \succ_{0} \mathbf{0}$, then $x \succ_{\tau} \mathbf{0}_{\tau}$; and if $x \prec_{0} \mathbf{0}$, then $x \prec_{\tau} \mathbf{0}_{\tau}$, $\tau \leq t_{x}$.

• If $(p, d', q, \tau) \succeq_0 (p, d, 1, 0) \succ_0 \mathbf{0}$ and $d' \ge d$, then $(p, d', 1, \tau) \succ_{\tau} \mathbf{0}_{\tau}$.

6. The term $s(\ln(1/q) + r(d)t)$ implies that the risk and time distance are substitutes, and that consumers exhibit diminishing sensitivity to distance (of either type). Thus, adding distance of one type reduces sensitivity to additional distance of either type. This observation affects our markdown problem as follows: since the option of waiting always exhibits time distance, consumers will not be very sensitive to probability reductions, even for *q* close to one.¹

3.3. The Parametric dPTT Model

Our theoretical results are all proven for the *general* model above, but for the numerical illustrations and empirical calibration we use a *parametric* form described below. Since the consumer can opt out from undesirable prospects, we can restrict attention to $u \ge p(1 - d)$ and omit the superscript "+" from *s*. We propose the following parametric specification:

$$w(u - p(1 - d)) = u - p(1 - d),$$

 $s(\sigma) = \sigma^{\beta}, \quad 0 < \beta \le 1, \text{ and}$
 $r(d) = \rho \ e^{\mu(d_0 - d)}, \quad \rho > 0, \ \mu \ge 0, \ d_0 \in [0, 1].$

• The linear form for v is widely used in management science and economics models.

• The power form for *s* is associated with Prelec (1998)'s probability weighting function, $w(q) = e^{-(-\ln q)^{\beta}}$ and Ebert and Prelec (2007)'s time discount functions, $f(t) = e^{-(t)^{\beta}}$. These specifications of *w* and *f* fit the experimental data well (Ebert and Prelec 2007, Booij et al. 2010) in the risk-only and time-only domains, respectively.

• The function r(d) is new. Our rationale is that r(d) is a parsimonious function that is bounded, decreasing, and strictly positive. The parameter μ captures subendurance (the higher the μ , the more impatience with respect to small price discounts). The parameter d_0 represents the reference discount around which subendurance emerges. That is, the subendurant dPTT consumer (with $\mu > 0$) is less patient for discounts lower than d_0 and more patient for discounts over d_0 . The parameter ρ represents the "baseline" time discount rate without subendurance, i.e., that of a DEU decision maker (with $\mu = 0$) or, equivalently, that of a dPTT decision maker facing a price discount of d_0 which, by definition, is a discount at which our model is calibrated with DEU.

The resulting model, called *parametric dPTT model*, is given by:

$$V_{\tau}(p,d,q,t) = [u - p(1-d)]$$

$$\cdot \exp\{-(\ln(1/q) + \rho e^{\mu(d_0 - d)}(t-\tau))^{\beta}\}, \quad (3)$$

with $u \ge p(1-d)$, $\beta \in (0, 1]$, $\rho > 0$, $d_0 \in [0, 1]$, and $\mu \ge 0$. Setting $\beta = 1$ and $\mu = 0$ yields DEU. Values of $\beta < 1$ will induce diminishing sensitivity to psychological distance in risk and time, and values of $\mu > 0$ will induce more subendurance.

Wait or Buy Decisions Under the dPTT Model

We now turn our attention to a market environment in which numerous consumers with homogeneous dPTT preferences best respond to a selling mechanism designed by the retailer.

4.1. The Selling Mechanism

The game involves one retailer (seller) and a continuum of consumers with a total mass of $\lambda > 0$. Consumers exhibit identical dPTT parameters and differ only in the benefit of consumption u. The value of u is private information independently drawn from a distribution with cdf F(u). We assume that F is continuous with support $[0, \bar{u}], \bar{u} > p$. Without loss of generality, we set $\bar{u} = 1$. Both λ and F are common knowledge. Throughout, \bar{F} denotes 1 - F and U[0, 1] denotes the uniform distribution on [0, 1].

The retailer has an initial inventory Q of a homogeneous, perishable, and infinitely divisible product that cannot be replenished and needs to be depleted over a two-period selling season. Without loss of generality, we let time 0 be period 1 and some given t > 0 be the "markdown" period 2. At time 0, the product is priced at the (exogenously given) tag price $p \in [0, 1]$; the retailer's markdown management problem is to decide on the discount percentage $d \in [0, 1]$ to be applied to all units of unsold inventory at time *t*. That the retailer commits to *d* in period 1 is a typical assumption in dynamic pricing literature (Aviv and Pazgal 2008, Liu and van Ryzin 2008).

At the beginning of period 1, each consumer observes (Q, p, d) and chooses to "opt out," "wait" until time *t* and buy at price p(1-d) but face availability risk, or "buy now" at price *p*. All consumers act simultaneously and do not observe each other's choices. Let λ_1 and λ_2 be the mass of consumers who decide to "buy now" and "wait," respectively.

Clearance is modeled as an instantaneous event and calculated using a fluid model.² If $\lambda_1 \leq Q$, then all consumers who "buy now" do so at price *p*. If $\lambda_1 > Q$, then units are allocated following a lottery with each customer facing equal probability of receiving an item (a proxy for random arrivals and first-come, first-served allocation). The remaining inventory is for those customers who choose to "wait." Similarly, the probability of obtaining an item in period 2 is the inventory that remains available divided by the number of consumers that decided to wait. In summary, if $\lambda_1 > 0$ and $\lambda_2 > 0$, then the *probability of obtaining the item* in periods 1 and 2, respectively, is equal to

$$q_{1} = \min\left(\frac{Q}{\lambda_{1}}, 1\right) \text{ and}$$

$$q_{2} = \min\left(\frac{\max(Q - \lambda_{1}, 0)}{\lambda_{2}}, 1\right),$$

$$(4)$$

and otherwise equal to the limit of these expressions as $\lambda_1 \rightarrow 0$ or $\lambda_2 \rightarrow 0$ (i.e., equal to zero if there is no quantity available, and equal to one otherwise). Observe that customers in period 1 have priority, leading to $q_1 \ge q_2$ and $(1-q_1)q_2 = 0$ (if $q_1 < 1$, then $q_2 = 0$). By contrast to Özer and Zheng (2016), where consumers misperceive probability of availability, consumers in our model use $s(\ln(1/q) + r(d)(t-\tau))$, a nonlinear distortion of probability and time distance.

The payoffs of a *u*-consumer associated with "opt out," "wait," and "buy now," respectively, are

$$V_0(p, 0, 0, 0) = 0,$$

$$V_0(p, d, q_2, t) = [u - p(1 - d)]e^{-s(\ln(1/q_2) + r(d)t)}, \text{ and }$$

$$V_0(p, 0, q_1, 0) = [u - p]e^{-s(\ln(1/q_1))}.$$

All payoffs are calculated from the point of view of time 0. By Proposition 3, those who waited will carry out the intended purchase at time t.

The payoff for the retailer is given by the expected revenue:

$$R = p \cdot \min\{\lambda_1, Q\} + e^{-\omega t} \cdot p(1-d)$$
$$\cdot \min\{\lambda_2, \max(Q - \lambda_1, 0)\}, \tag{5}$$

where $\omega < \rho$ represents the retailer's opportunity cost (timevalue of money, obsolescence, etc.) We neglect those production costs that are already incurred.

Structurally, the selling mechanism is that of a Stackelberg game, with the retailer as the leader and the consumers as the followers. As is common in the analysis of such games, we first discuss the reaction of the followers, and then in Section 5 discuss the problem of the leader.

4.2. Consumer's Best Response and Nash Equilibrium

Suppose a *u*-consumer observes (Q, p, d). How should she react? We show that the best response can be characterized by a threshold, $H \in [p, 1]$.

• If $0 \le u < p(1-d)$, then "opt out" is a dominant strategy (buy is never profitable).

• If $p(1-d) \le u < p$, then "wait" is a dominant strategy ("buy now" is never profitable).

• If $p \le u \le 1$, then the consumer needs to form some expectation of q_1 and q_2 . For any such expectation, note that $V_0(p, d, q_2, t)$ ("wait") and $V_0(p, 0, q_1, 0)$ ("buy now") are linear functions of u, with the "wait" payoff having a smaller slope and a higher intercept. Hence, there is a unique threshold $H \in [p, 1]$ such that the consumer will wait if $p(1-d) \le u \le H$, and buy now if $H < u \le 1$.

Next, we infer the existence of, and restrict attention to, a symmetric equilibrium in pure strategies in which all consumers use the same threshold $H \in [p, 1]$. To characterize equilibrium, we assume that all consumers use $H \in [p, 1]$, calculate any one consumer's best response by means of the threshold $\mathcal{B}(H) \in [p, 1]$, and impose the equilibrium condition $\mathcal{B}(H) = H$. We denote by H^* any such solution. We begin by calculating $\mathcal{B}(H)$.

PROPOSITION 4. Assume all consumers use the threshold H. Then,

$$\lambda_1 = \lambda \overline{F}(H)$$
 and $\lambda_2 = \lambda (F(H) - F(p(1-d))).$ (6)

The best response of any one consumer is to opt out if u < p(1-d), to wait if $p(1-d) \leq u \leq \mathcal{B}(H)$, and to buy if $u > \mathcal{B}(H)$, where

$$\mathscr{B}(H) = \min\left\{p \cdot \frac{e^{-s(\ln(1/q_1))} - (1-d)e^{-s(\ln(1/q_2) + r(d)t)}}{e^{-s(\ln(1/q_1))} - e^{-s(\ln(1/q_2) + r(d)t)}}, 1\right\},$$
(7)

and q_1 and q_2 are as in (4). Moreover, $\mathfrak{B}(H)$ is increasing in H.

Because $\mathscr{B}(H): [p, 1] \rightarrow [p, 1]$ is a continuous mapping from a closed and convex set into itself, it admits at least one fixed point, $\mathscr{B}(H^*) = H^*$, ensuring that *a symmetric equilibrium in pure strategies exists*. The existence of a pure strategy equilibrium is crucial. The definition and existence of equilibrium in mixed strategies would be problematic in our setup because consumers treat probabilities in a nonlinear fashion. That $\mathscr{B}(H)$ is increasing in *H* hinges on positive externality: The more customers who "wait," the higher the product availability in period 2. That $dq_2/dH > 0$ is a bit counterintuitive at first, but clear in hindsight. Consider a small increase in *H*, i.e., a few consumers switch from "buy now" to "wait." Some units that would have been purchased with probability q_1 are now purchased with probability $q_2 \leq q_1$. On a first-order approximation, the demand in period 2 increases by q_2 , the supply increases by q_1 , and the net effect is an increase in availability.

We distinguish three regimes, depending on supply. In the first regime, all customers purchase with probability one. In the second, there is only rationing among those that decide to wait. In the third, all customers may experience rationing.

PROPOSITION 5. Given (λ, F) and (Q, p, d), there are three regimes:

(i) Abundant supply. If $Q \ge \lambda \overline{F}(p(1-d))$, then $q_1 = q_2 = 1$, and the best response threshold is constant for all H and given by

$$\mathscr{B} = \min\left\{p \cdot \frac{1 - (1 - d)e^{-s(r(d)t)}}{1 - e^{-s(r(d)t)}}, 1\right\}.$$
(8)

Note that $\mathcal{B} > p$. There is a unique equilibrium given by $H^* = \mathcal{B}$.

(ii) Intermediate supply. If $\lambda \overline{F}(p) < Q < \lambda \overline{F}(p(1-d))$, then $q_1 = 1$, $q_2 \in (0, 1)$, and the best response threshold is given by

$$\mathscr{B}(H) = \min\left\{p \cdot \frac{1 - (1 - d)e^{-s(\ln(1/q_2) + r(d)t)}}{1 - e^{-s(\ln(1/q_2) + r(d)t)}}, 1\right\}.$$
 (9)

There is at least one equilibrium solving $\mathcal{B}(H^*) = H^*$ with $H^* > p$.

(iii) Limited supply. If $Q \leq \lambda \overline{F}(p)$, then $\mathcal{B}(H) = p$ on $H \in [p, F^{-1}(1 - Q/\lambda)]$. We have that $H^* = p$ is always an equilibrium, but other equilibria with $H^* > p$ may exist.

Note that $\mathscr{B}(H) = p$, if and only if, $q_2 = 0$. Thus, the equilibrium $H^* = p$ implies a congestion of "buy now" customers facing $q_1^* = Q/\lambda \bar{F}(p) \leq 1$, and those that must wait face $q_2^* = 0$. Any equilibrium with $H^* > p$ exhibits $q_1^* = 1$ and $q_2^* > 0$.

4.3. Uniqueness and Equilibrium Selection

A general condition for equilibrium uniqueness is $\mathscr{B}'(H) < 1$ at all points where \mathscr{B} is differentiable. This condition is trivially met if the supply is abundant (\mathscr{B} is constant). The condition is also met if the supply is close to abundant. As shown in the proof of Proposition 4, both dq_2/dH and $\mathscr{B}'(H)$ are proportional to $(\lambda_2 + \lambda_1 - Q)$. Hence, if $Q \rightarrow \lambda \overline{F}(p(1-d)) = \lambda_1 + \lambda_2$, then q_2 and $\mathscr{B}(H)$ become insensitive to *H*, and the equilibrium is unique. A $\mathscr{B}(H)$ insensitive to *H* reduces the strategic burden on consumers.

If the supply is not close to abundant, however, q_2 might be quite sensitive to H and the equilibrium might not be unique. A lack of uniqueness is not due to behavioral effects. Osadchiy and Vulcano (2010) and, more recently, Correa et al. (2016) observed multiple equilibria under DEU and provided sufficient conditions for uniqueness. The dPTT model makes the best response function more nonlinear, and can exacerbate the nonuniqueness problem.

Our game is one of coordination. Intuitively, the equilibrium with highest q_2^* (i.e., highest H) improves the payoff for those who always wait, and leaves unaffected (under intermediate supply) or may improve (under limited supply) the payoff of those who always buy now. Hence, all the equilibria can be Pareto ranked, and the equilibrium with highest H^* is Pareto dominant.

PROPOSITION 6. Let H^* and H'^* be two equilibria. If $H^* > H'^*$, then H^* Pareto-dominates H'^* .

Naturally, our selection criteria is to choose the one equilibrium with the highest H^* .

We verify that the equilibria of the fluid model are good approximations of the equilibria of a stochastic demand model. To do so, we numerically analyzed a stochastic demand model in which the total number of customers is a Poisson random variable with rate λ . In that model, iterative calculations of the best response rapidly converge to a fixed point. The difference between the buy-now equilibria thresholds computed under the stochastic and fluid models diminishes as the demand and capacity are scaled up.

5. Markdown Management

Following the usual approach for analyzing Stackelberg games, we have assumed thus far that the quadruplet (Q, p, d, t) is fixed, and subsequently studied equilibrium consumer (follower) behavior, which is described by the mapping H^* : $[0, Q] \times [0, 1] \times [0, 1] \times [0, T] \rightarrow [p, 1]$. We have shown that a symmetric equilibrium exists; it is unique if supply is sufficiently abundant, or otherwise there is unique Pareto-dominant equilibrium. The goal of the seller (leader) is to find the selling arrangement (Q, p, d, t) that maximizes the revenue function. Since the main focus of this section is the markdown optimization d, we define the seller's markdown optimization problem as:

$$R(Q, p, t) \equiv \max_{d} R(d)$$

= $\max_{d} \{ p \min\{\lambda \bar{F}(H^*), Q\} + p(1-d)e^{-\omega t} \cdot \min\{\lambda(F(H^*) - F(p(1-d))), (Q - \lambda \bar{F}(H^*))^+ \} \}.$ (10)

Let d^{dPTT} be the solution to (10) for general $s(\sigma)$ and r(d), H^{dPTT} , i.e., the corresponding consumer equilibrium, and R^{dPTT} , i.e., the optimal revenue. In the absence of the behavioral anomalies, let d^{DEU} denote the solution to (10) for $s(\sigma) = \sigma$ and $r(d) = \rho$, H^{DEU} the corresponding consumer equilibrium, and R^{DEU} the optimal revenue. The difference between R^{dPTT} and R^{DEU} represents the revenue opportunity associated with the behavioral anomalies, and the difference

between R^{dPTT} and $R^{dPTT}(d^{DEU})$ is the revenue gain due to incorporating dPTT behavior into markdown optimization.

For the analytical results below, we assume F = U[0, 1]. Proposition 7 formulates our main result: A retailer who currently implements d^{DEU} (thinking that consumers follow DEU, or absent a tool to account for the actual behavior), while consumers exhibit the dPTT anomalies, could increase its revenue by offering larger markdowns.

PROPOSITION 7. For a given λ , p, t, and $u \sim U[0, 1]$, let $Q > \lambda(1 - \frac{1}{2}p(1 + e^{\omega t - \rho t}))$ and $p(e^{\rho t} - 1 + (1 - e^{\omega t - \rho t})/2) < e^{\rho t} - 1$. Then

$$\begin{split} \frac{\partial R^{\mathrm{dPTT}}}{\partial d} \bigg|_{d=d^{\mathrm{DEU}}} &> 0, \quad iff - r'(d^{\mathrm{DEU}}) \times s'(\sigma) \\ &< \frac{4(e^{s(\sigma)-\rho t}-1)(e^{s(\sigma)}-1)}{e^{s(\sigma)}t(1-e^{\omega t-\rho t})(2-e^{-\omega t}-e^{-\rho t})}. \end{split}$$

This result is subject to three technical conditions. The first condition requires the inventory Q to be sufficiently large (see Section 5.2 for more detail). The second precludes the "all-wait" equilibrium. The third states that the dPTT behavioral effects cannot be simultaneously too strong. Indeed, because $r'(d) \leq 0$ and $s'(\sigma) > 0$, the left-hand side (LHS) is a positive number. The right-hand side (RHS) is also positive if $s(\sigma) \ge \rho t$, which is true for small σ because s is increasing and concave with s(0) = 0, s(1) = 1. Hence, for the second condition to hold, the two derivatives cannot be simultaneously large at d^{DEU} . This will be true when $|r'(d^{\text{DEU}})|$ or $s'(\sigma)$ are not very large. In particular, since |r'(0)| is bounded, the condition holds if t is sufficiently large.

5.1. Impact of Behavioral Anomalies on Markdowns and Revenue

To quantify the impact of the dPTT behavior on markdowns and revenue, we adopt the parametric dPTT model (3) and conduct a numerical study. For the base case, we consider $\lambda = 1, u \sim U[0, 1], Q = 0.625, p = 0.5, and t = 3.$ The quantity Q ensures that the seller is in the abundant or intermediate supply regime. The price p = 0.5 is arbitrary and we optimize the percentage discount d. A t = 3 corresponds to a 9-week delay, meaning the unit of time is 3 weeks, a constraint imposed by our estimation procedure (see Section 6). Nine weeks is half of the median price duration in a broad panel of consumer goods and services (Bils and Klenow 2004), a reasonable time to wait for a markdown. We set the (3-week) discounting rate parameter at $\rho = 0.13$, and the price-discount anchoring parameter at $d_0 = 0.5$. The discounting rate for the retailer is $\omega = 0.05$, reflecting the time value of money and opportunity costs. Consistent with the markdown management literature, we assume that retailers are more patient than consumers at $d = d_0$ (von der Fehr and Kuhn 1995). For dPTT, we set $\beta = 0.9$ and $\mu = 1.95$ (see Section 6) and extend the parameter region to $(\beta, \mu) \in [0, 1] \times [0, 6.5]$. The DEU model is the corner $(\beta, \mu) = (1, 0).$

Figure 1. (Color online) (a) Revenue, (b) Equilibrium threshold H, (c) Probability q_2 , (d) Buy-now λ_1 , and total quantity sold as a function of markdown d, under dPTT ($\beta = 0.9$, $\mu = 1.95$) and DEU models.



Notes. Dashed vertical lines mark optimal markdowns under DEU and dPTT. Parameters: Q = 0.625, p = 0.5, t = 3, $\lambda = 1$, $u \sim U[0, 1]$, $\omega = 0.05$, $d_0 = 0.5$, and $\rho = 0.13$.

(i) Retailers can offer larger markdowns under dPTT. Figure 1 illustrates the base-case instance of the markdown optimization problems. The revenue function under DEU is maximized at $d^{\text{DEU}} = 0.107$, and the dPTT revenue function is maximized at $d^{\text{dPTT}} = 0.186$, a 74% increase. Another helpful benchmark is the case of myopic (i.e., infinitely impatient) consumers, $\rho = \infty$, for whom $d^{myopic} = 0.5$. That is, dPTT sets the optimal markdown between the DEU and the myopic values.

One could argue that retailers have historically underestimated strategic consumer behavior. By setting markdowns close to 50%, they effectively trained consumers to wait strategically, which, in turn, led to the stream of research on markdown management with strategic consumers. In particular, multiple studies showed that disregarding strategic consumer behavior and mistakenly setting markdowns at d^{myopic} while consumers are strategic can be very costly (Aviv and Pazgal 2008, Ovchinnikov and Milner 2012). This is also easy to see from Figure 1(a) where the revenue at d = 0.5is smaller than at d = 0. However, the predictions from the DEU model are equally unsatisfying because the markdown of $d^{\text{DEU}} \approx 10\%$ is intuitively too small. The $d^{\text{dPTT}} \approx 20\%$ makes intuitive business sense and agrees with conventional wisdom that markdowns should be large enough to impact behavior, yet small enough to control strategic waiting.

(ii) *dPTT pricing generates higher revenue.* Consistent with Proposition 7, the dPTT revenue function has a positive derivative at d^{DEU} . Indeed, $R^{\text{dPTT}} = 0.270$ versus $R^{\text{dPTT}}(d^{\text{DEU}}) = 0.266$. Thus, accounting for the dPTT behavior in markdown optimization generates a 1.5% revenue gain (Figure 1(a)). In the figure, note that charging a single price (d = 0) generates a revenue of 0.25. Hence, the value of our model is better captured by the *incremental* revenue that the firm generates from dynamic pricing, i.e., 0.02 (dPTT) instead of 0.016 (DEU), a 25% increase. We refer to this incremental increase as *effectiveness*. Clearly, dPTT makes markdown management substantially more effective.

(iii) *dPTT behavior drives more sales at the tag price and in total.* Figure 1(b) shows that for $d \leq 0.54$, the number of

items sold at the tag price is greater under the dPTT model (that $H^{dPTT} \leq H^{DEU}$ implies that fewer customers wait). At optimal markdowns, the dPTT model sells 0.416 units at the tag price and 0.177 units at the markdown. The respective quantities for the DEU model are 0.388 and 0.166 units (Figure 1(d)). In both cases the split is close to 70% of units sold at the tag price, and 30% at the markdown, despite dPTT's markdown being almost twice as large. The behavioral anomalies enable retailers to increase markdowns, and open up the market, without sacrificing the tag price sales. Both d^{dPTT} and d^{DEU} correspond to the abundant supply regime, i.e., $q_2 = 1$, and every consumer with $u \ge p(1 - d)$ receives an item (Figure 1(c)). In total, dPTT sells 0.593 units, while DEU sells 0.553 units; this amounts to a gain of 7.1% (Figure 1(d)).

Note that dPTT qualitatively predicts quite different behavior as to how consumers react to discounts. For instance, for suboptimally large markdowns, $d \ge 0.25$, DEU predicts that the number of tag price sales will remain the same; see the dashed line in Figure 1(b). Because DEU treats the surplus and probability linearly, the increase in surplus is compensated by a decrease in probability so that the threshold *H* and, consequently, the tag-price sales λ_1^{DEU} are constant. This is not so under dPTT: Consumers become more patient in response to high *d*, decreasing λ_1^{dPTT} . As a result, the revenue losses from a suboptimally large markdown can be much larger than DEU predicts.

(iv) Impact of individual anomalies. To understand how the anomalies we study impact markdowns and revenue, we first show why these anomalies lead to larger markdowns at the optimum. The optimal solution balances the marginal benefit of increasing markdown with the associated marginal cost. The benefit comes from selling an additional unit of inventory, albeit at a lower price; it is a decreasing function of d, and is identical for DEU and dPTT. The cost comes from two sources, both of which increase in d: All markdown units sell at a lower price, and some consumers divert from tag-price purchases to markdowns. The former is also unaffected by the anomalies, but the latter is. Under dPTT

Figure 2. (Color online) Drivers of optimal markdown under dPTT. Panel (a) presents the base case (β , μ) = (0.9, 1.95), and combinations $(1, \mu)$, $(\beta, 0)$, and (1, 0) (or DEU). Panel (b) presents the same for $(\beta, \mu) = (0.4, 1)$.



Note. Parameters: Q = 0.625, p = 0.5, t = 3, $\lambda = 1$, $u \sim U[0, 1]$, $\omega = 0.05$, $d_0 = 0.5$ and $\rho = 0.13$.

more consumers buy at the tag price; the marginal cost is therefore smaller and hence it intersects the marginal benefit at a higher markdown value.

Figure 2(a) illustrates this logic. The marginal benefit is linear, decreasing in d, and drops to 0 at d = 0.25 due to the inventory constraint. The marginal cost increases in dfor both models and is, in fact, linear for DEU; the two intersect at 10.7%. For dPTT with $(\beta, \mu) = (0.9, 1.95)$ the marginal cost is nonlinear and smaller, and it therefore crosses the benefit line at a higher value, 18.6%. Both are the optimal markdowns we saw earlier. The marginal cost under dPTT is affected by the anomalies. Setting $(\beta, \mu) =$ (1, 1.95) isolates subendurance (dashed line), and $(\beta, \mu) =$ (0.9, 0) isolates the decreasing sensitivity to psychological distance (dotted line). At the base-case parameters, the larger markdown is driven mostly by subendurance (see Figure 2(a)). This is not surprising since at $\beta = 0.9$ the curvature in the psychological distance function is small. However, if $(\beta, \mu) = (0.4, 1)$ (Figure 2(b)), the difference is mostly driven by the decreasing sensitivity to psychological distance. Significantly, in both cases, the effects are complementary.

(v) Value of considering anomalies jointly. We next investigate the joint effect of β and μ on the optimal markdown and revenue. Recall that $\beta < 1$ and $\mu > 0$ are considered deviations from DEU: Smaller β reflects a stronger common ratio and common difference effects, and larger μ reflects stronger subendurance. We find that stronger behavioral anomalies lead to greater markdowns (Figure 3(a)), and higher revenues (Figure 3(b)). Consistent with (iv), Figure 3(a) shows





Note. Parameters: Q = 0.625, p = 0.5, t = 3, $\lambda = 1$, $u \sim U[0, 1]$, $\omega = 0.05$, $d_0 = 0.5$, and $\rho = 0.13$.

11

that a deviation along one dimension (β or μ) is sufficient for a larger optimal markdown, and the deviation along the other reinforces the effect. The optimum markdown plateaus at the inventory clearing level.

Figure 3(b) shows the revenue opportunity offered by the behavioral anomalies. By traversing Figures 3(a) and 3(b) from the point $(\beta, \mu) = (0, 1)$ to, for example, $(\beta, \mu) = (0.5, 3)$, the optimal markdown increases from 10.7% to 22.5% (Figure 3(a)), resulting in a revenue opportunity of 6.89% compared to the DEU behavior (Figure 3(b)). The smaller DEU-optimal markdown leaves about half of revenue opportunity untapped. dPTT captures this untapped revenue, with a 3.04% gain relative over DEU pricing (Figure 3(c)), and increases the effectiveness of markdowns by 46.5% (Figure 3(d)).

Figure 3 also emphasizes the value of considering a model that integrates multiple anomalies pertinent to wait-or-buy decisions. While a deviation along β or μ is sufficient for an increase in the effectiveness of markdowns, such increased effectiveness does not necessarily require strong individual anomalies. For instance a 30% effectiveness gain could be obtained from a very strong deviation from rationality in sensitivity to psychological distance ($\beta \approx 0.4$) or a more moderate deviation in distance combined with moderate subendurance (e.g., $\beta \approx 0.8$, $\mu \approx 2$). That is, both the decreasing sensitivity to psychological distance and subendurance are important, and the dPTT model that simultaneously considers them allows the seller to capture the revenue opportunity offered by the behavioral anomalies, and thus increase the effectiveness of markdowns.

5.2. Strategic Capacity Rationing

Proposition 7 requires that Q be sufficiently large to ensure, at minimum, a boundary between the intermediate and abundant supply regimes. Intuitively, the seller wants to be in the region of intermediate supply, or its frontiers. Under the limited supply, the retailer can raise p and increase revenue; under the abundant supply, the retailer can reduce Q to such a boundary and not affect revenue. Hence, our assumption is not that restrictive.

However, a well known result in dynamic pricing is that the retailer can increase revenue by reducing inventory beyond such a boundary. This concept is called *strategic capacity rationing* (Liu and van Ryzin 2008). Doing so increases the shortage risk and induces some consumers to purchase at the tag price rather than wait. Thus, a necessary condition for the shortage risk to increase revenue is the increased number of tag price purchases, or, equivalently, the decrease in threshold H in response to the increased d. Proposition 8 shows when this can happen under dPTT.

PROPOSITION 8. For given λ , Q, p, t and $u \sim U[0, 1]$, let $d^* = 1 - (1 - Q/\lambda)/p$ be the inventory clearing markdown. Then $(\partial H/\partial d)|_{d=d^*} < 0$, provided

$$\frac{p - e^{-s(r(d^*)t)}(1 - Q/\lambda)}{1 - e^{-s(r(d^*)t)}} \leq 1, \quad and$$

$$\begin{split} &1 - s'(r(d^*)t) \bigg(1 - r'(d^*) \frac{1 - Q/\lambda}{p} \bigg) \\ &< \frac{s'(r(d^*)t)r'(d^*)(1 - e^{-s(r(d^*)t)}((1 - Q/\lambda)/p))}{1 - e^{-s(r(d^*)t)}}. \end{split}$$

The first condition is technical. The second is satisfied when $s'(\sigma)r'(d^*)$ is large. For example, this can occur if the patterns of the common ratio, common difference, and subendurance effects are strong, or if *t* is small. If the condition is not satisfied, the seller sets *d* to be at the frontier between intermediate and abundant supply.

Note that for the case of DEU, the second condition reduces to the equality $H'_d(d^*) = 0$, i.e., inducing scarcity risk is never optimal under DEU. This replicates the result of Liu and van Ryzin (2008) who showed that rationing cannot be optimal if consumers have a linear value function, u - p, as we assume. Their result requires a concave function of u - p, which, as discussed in Section 2, is at odds with experimental evidence. By introducing probability distortions à la prospect theory, dPTT justifies strategic capacity rationing from a more plausible modeling angle.

In our baseline numerical example, the conditions of Proposition 8 do not hold. If, however, the selling season t is sufficiently short, or if consumers are patient (small ρ), then these two conditions can be satisfied and the seller may find it optimal to set a markdown that induces the intermediate supply regime with $q_2 < 1$. This is illustrated in Figure 4(a), where Q is set slightly above the limited supply level, and the markdown d = 0.018 clears the entire inventory. At d = 0.018, approximately 96% of consumers choose to wait and tag-price sales fall to $\lambda_1 = 0.0161$, negatively impacting revenue. However, increasing d beyond 0.018 creates the scarcity risk: The optimal d = 0.037 results in $q_2 =$ 0.9, which increases tag price sales to $\lambda_1 = 0.417$ so that only 18% wait. In other words, scarcity risk induces early purchases. In this example Q = 0.509; for any larger Q, the retailer can increase revenue by reducing supply to this quantity. Strategic capacity rationing can have large implications for profit if the inventory is costly Liu and van Ryzin (2008).

5.3. Markdown Timing

In some contexts retailers can also control the timing of the markdown. When we set *t* as a decision variable, we find that for a broad range of parameters, the dPTT model suggests that a retailer offers the markdown sooner. Let t^{dPTT} and t^{DEU} be the revenue-maximizing delays under the dPTT and DEU models, respectively. The following result can be established.

PROPOSITION 9. For given λ , p, d, and $u \sim U[0, 1]$, let $Q > \lambda(1 - p(1 - d))$. Then $t^{dPTT} < t^{DEU}$, provided $e^{s(r(d)t)} - e^{\rho t} > (s''(r(d)t)(e^{s(r(d)t)} - 1))/(s'(r(d)t))^2$ for all $t < t^{DEU}$.

Note that the RHS of the condition is negative. Therefore, the condition is guaranteed to be satisfied if $s(r(d)t) > \rho t$. By the properties of the function *s*, the condition holds when *t* is sufficiently small. Numerically, under the baseline



Note. Parameters: Q = 0.509, p = 0.5, t = 0.01, $\lambda = 1$, $u \sim U[0, 1]$, $\omega = 0.05$, $d_0 = 0.5$, and $\rho = 0.13$. (b) Revenue under dPTT and DEU as a function of delay and optimal markdown. Parameters: Q = 0.625, p = 0.5, $\lambda = 1$, $u \sim U[0, 1]$, $\omega = 0.05$, $d_0 = 0.5$, and $\rho = 0.13$.

scenario, and d = 0.2, we observe that the optimal delay under dPTT is almost twice as short ($t^{dPTT} = 5.7$ versus $t^{DEU} = 10.4$). We also observe that by setting shorter delay dPTT extracts 1% additional revenue (0.273 versus 0.270).

If a retailer can jointly optimize d and t, in the baseline scenario, we find that dPTT and DEU set markdowns at the same inventory clearing level $d^{dPTT} = d^{DEU} = 0.25$, however, the dPTT sets the markdown sooner $t^{dPTT} = 7.2$ versus $t^{\text{DEU}} = 11.9$ with the incremental revenue gain of 0.7% (0.274 versus 0.272) (Figure 4(b)). As Q increases both models further increase and delay markdowns, with the increasing spread between t^{dPTT} and t^{DEU} , and slightly increasing the revenue gain. Absent supply constraints, the dPTT timing achieves 1.0% higher revenue compared to the DEU optimal timing, with a delay almost twice as short (7.7 versus 14.6), and requiring fewer units of inventory (0.63 versus 0.67). Thus, incorporating the dPTT behavior into the markdown magnitude and timing optimization allows retailers to extract more revenue over a shorter period of time, and with a smaller inventory. This, again, can be even more profitable if the inventory is costly.

5.4. Robustness Analysis

To conclude the numerical illustrations, we considered 10 scenarios in which we varied parameters p, t, Q, F, d_0 , and ρ , keeping the rest at their baseline. Table 2 presents the results, and the general observation is evident: The key qualitative managerial insights obtained above are not affected by the model parameters. Indeed, the condition of Proposition 7 is satisfied for a wide range of parameter combinations. Quantitatively, the revenue gain can exceed 3.5%, a fairly remarkable increase, considering that the optimization is performed with respect to d alone. In other studies, similar revenue gains require joint optimization of pricing and capacity (Özer and Zheng 2016). Moreover, the effectiveness gain in markdown management is, in many cases, 25% or more.

The benefit of dPTT is even greater when markdowns are jointly optimized with tag prices. The bottom row in Table 2 presents the results of such joint optimization for the basecase scenario: The revenue gain increases to 1.87% and the effectiveness gain reaches nearly 30%. Across all the scenarios (not depicted for conciseness), the average revenue

Table 2.	Impact of model	parameters on mar	kdowns, revenues,	, and revenue an	d markdown	effectiveness	gains from c	IPTT
----------	-----------------	-------------------	-------------------	------------------	------------	---------------	--------------	-------------

Scenario	Parameters	d^{dPTT} (%)	d^{DEU} (%)	$R^{ m dPTT}$	$R^{\mathrm{dPTT}}(d^{\mathrm{DEU}})$	Rev. gain (%)	Eff. gain (%)
Base case		18.6	10.7	0.2700	0.2660	1.50	25.0
Expensive product	p = 0.75, Q = 0.375	17.3	10.7	0.2318	0.2234	3.76	23.4
Cheap product	p = 0.25, Q = 0.825	18.6	10.7	0.1925	0.1915	0.52	25.0
Short selling season	t = 1	10.50	3.8	0.262	0.2568	2.02	76.5
Long selling season	t = 5	22.8	16.5	0.2728	0.2708	0.74	9.6
Concentrated valuations	$u \sim \text{Beta}(4, 4), Q = 0.75$	18.0	12.1	0.2921	0.2871	1.74	13.5
Disperse valuations	$u \sim \text{Beta}(0.4, 0.4)$	18.7	10.5	0.2609	0.2586	0.89	26.7
Low ref. markdown	$d_0 = 0.25$	13.2	10.7	0.2626	0.2621	0.19	4.1
High ref. markdown	$d_0 = 0.75$	24.5	10.7	0.2781	0.2683	3.65	53.6
Patient consumers	$\rho = 0.06$	10.3	1.5	0.2589	0.2523	2.62	287
Impatient consumers	$\rho = 0.2$	23.8	18.1	0.2771	0.2753	0.65	7.1
Base case, joint (p, d)	$p^{\text{dPTT}} = 0.5434,$	18.57	10.67	0.2717	0.2667	1.87	29.9
optimization	$p^{\text{DEU}} = 0.5156$						

Note. The parameters are set to the baseline values, except those explicitly specified.

gain increases to 1.86% versus 1.66% when only *d* is optimized, and the effectiveness of markdowns increases from 50% to 55%.

To summarize, by correctly accounting for the behavioral anomalies in consumer wait-or-buy decisions, the retailer can set higher, more meaningful markdowns *and* generate more revenue. However, the magnitude of revenue gains clearly depends on the dPTT model parameters. Hence, we next calibrate the dPTT model and quantify its practical impact.

6. Experimental Calibration and Validation of the dPTT Model

Here our goal is to estimate the dPTT model parameters. For the baseline time discount parameter we used $\rho = 13\%$, which is the average discount rate of 18% for monetary gains between €50 and €100 from Baucells et al. (2009) adjusted for the fact that a unit of time in their study was a month and in ours it was 3 weeks ($18 \times 3/4 \approx 13$). Interestingly, fitting the DEU model to our wait-or-buy data also results in $\rho = 13\%$. For the reference discount parameter, we used $d_0 = 50\%$, which we obtained through an online survey with N = 32 student participants from Canada. We asked them: "Think about an end-of-season sale (markdown) at a retail store, such as Boxing Day, for example. What is the percentage price discount that first comes to mind?" The average response was 51%; both the mode and median were 50%. For sensitivity to discount μ and sensitivity to psychological distance β , we conducted an experiment described below.

6.1. Design of the Experiment

To estimate (β, μ) parameters, it suffices to collect consumer wait-or- buy choice data for various discount levels *d*, and probabilities of product availability *q*. The dPTT model allows us to hold the benefit of consumption *u*, the tag price *p*, and the time delay *t* constant.

The choice data can be collected in two ways. An intuitive approach is to collect the binary choice data, i.e., present subjects with combinations of (d, q) and ask if they would buy or wait. The challenge of this approach is that to have good coverage of input parameter space, each subject would have to answer dozens of nearly identical questions, causing subjects fatigue and disengagement, which are known to significantly lower the quality of elicitation. An alternative approach is to use choice lists, e.g., present subjects with blocks of questions each containing a list of binary wait-orbuy questions for different values of q where d is held constant in each block, and varied across blocks. This approach has been shown to increase the quality of elicitation in situations where there is an implicit indifference point (this will be shown for our data). See the well known paper by Holt and Laury (2002) for an example of using choice lists.

In the experiment, we set the benefit of consumption u = 250 and the tag price p = 200. The time delay is set at three weeks, which we normalized as t = 1. The "buy

later" option was varied over a wide range of probabilities of product availability, q = 10%, 20%, ... 90%, and discounts d = 5%, 15%, 25%, 50%, 75%. The selected discounts are based on evidence from Pilehvar et al. (2016, Table 1) who documented that inventory liquidators in practice recover 27 ± 5 cents on a dollar of the tag price; hence, discounts larger than 75% seem ineffective. Thus, each subject saw five choice lists, one for each d, with nine q values in each list. We ensured that all 45 combinations were assigned to at least one subject. The five lists appeared in random order (see Figure 8(a) in the appendix for a screenshot). As a reliability test, and after answering the choice lists, subjects answered five binary choice questions, one from each choice list, drawn and ordered randomly as well (see Figure 8(b)).

To make the elicitation incentive-compatible we determined the payments using a random incentive scheme (Savage 1954). Specifically, we implemented the Prince method, Johnson et al. (2014). Under this method, the scenario for potential compensation is randomly assigned to a subject *before* the experiment, but is unknown until the experiment concludes. The scenario is provided to the subjects in a tangible/physical form (usually in a sealed envelope), and subjects' answers are framed as instructions to the experimenters about how to implement the real choice situation in the envelope. Johnson et al. (2014) show that this improves the quality of elicitation.

6.2. Subjects and Procedure

Subjects, 64 business school students, were recruited through an online system to participate in a decision-making experiment that promised earnings of a minimum of \$5 and a maximum of \$200. Upon arrival, subjects picked physical sealed envelopes and experimental instructions (see the appendix). The instructions gave the following description of the decision situation:

Suppose that you went to a retail store and saw a product that you know you can resell for \$250 at any time. The product was priced at \$200 (two hundred dollars), so you picked the product from the shelf and were about to purchase. However, then you started thinking that in three weeks from today this product may be marked down. Thus the question was: Should you buy the product now or wait for the markdown?

From reading the instructions, the subjects learned that the sealed envelopes contained two numbers, i.e., markdown percentage and probability of product availability. They further learned that two of them will be selected at the end of the experiment. The envelope of a selected subject will be privately opened and the scenario in the envelope will be played as per the choices she will make during the experiment. That is, if in that scenario (s)he chose to "buy now" then (s)he will receive the "buy now" payment, \$50, immediately, plus the \$5 participation fee paid to all subjects. If (s)he chose to "wait for the markdown" then (s)he would come again to see the experimenter in three weeks to learn about the product's availability (determined by whether a U[1,100] random integer is less than or equal to q) which will, in turn, determine her payoff: 0 if the product is unavailable, or $250 - 200 \times (1 - d) \in [\$60, \$200]$ otherwise. The three-week delay was selected due to constraints associated with the subjects' availability.

Per instructions, two subjects were randomly selected, ID49 (male) with d = 25% and q = 50%, and ID11 (female) with d = 25% and q = 80%. ID49's response to the q = 50%question in the d = 25% choice list was to "buy now," thus, he was awarded \$55 and left. ID11's response to the q =80% question in the d = 25% choice list was to "wait for the markdown," thus, she left the experiment with \$5 payment, but in three weeks came again to see the experimenter. The random number drawn was $15 \leq 80 \equiv$ (her q), which meant that the product was available. Thus she "bought it" for \$200 * (1 - 0.25) = \$150 and immediately resold it to the experimenter for a surplus of \$100. That concluded our experiment.

6.3. Structure of the Data and Initial Checks

The experimental data consists of a series of wait-or-buy decisions that the subjects made for different (d, q) combinations. However, within each choice list, such data are not independent, particularly if the choice lists reveal indifference points. An indifference point would imply that within each choice list, as q's increase, subjects switch from selecting "buy" for low q's to selecting "wait" for high q's, and they do so only once (i.e., they do not switch back to "buy" at even higher q's). Then, somewhere between the highest "buy" q and the lowest "wait" q is the probability at which, for that specific d, the subject is indifferent between buying and waiting

Our data shows overwhelming support for the existence of indifference points: 62/64 subjects exhibited such an indifference-point pattern in all choice-lists, one subject switched more than once in a single choice list, and one subject (ID56) exhibited behavior that is largely inconsistent with the notion of the indifference point. We believe that (s)he misunderstood the task; this is further supported by the consistency analysis that follows. Thus we removed all subject ID56 data, and for the rest defined indifference points as the mid-value between the highest "buy" q and the lowest "wait" q (Holt and Laury 2002). Whenever a choice list had all "buys," the indifference points were defined as 95%, and conversely, in a choice list with all "waits," the indifference point was defined at 5%. In both cases, these are the midpoints between the corresponding subject's answer and the boundary of the [0, 1] interval for probabilities. This led to 315 indifference points.

The binary choice data were used to verify within-subject consistency by comparing the decision made in a binarychoice question to the corresponding decision made in the respective choice list. This analysis also reveals formidable consistency: 67% of our subjects were consistent in all decisions; 23% were consistent in all-but-one choice. Only one subject, ID56, was inconsistent in more than three, which supported our decision to exclude his/her data from the analysis. The remainder is a set of highly consistent wait-or-buy data on 63 usable subjects, which we next use to estimate the dPTT model parameters (β , μ).

6.4. Parameter Estimation

Given the structure of our data, we designed the estimation procedure to minimize the deviations between the observed and implied indifference points. Note that we purposefully did not use a somewhat more intuitive maximum likelihood estimation for the binary choice data because the existence of the indifference points for most subjects implies that many of the binary choice data points are not independent. Indeed, if for d = 25% one's indifference point is q = 65%, then the only two independent binary choice data points are a "buy" for d = 25%, q = 60% and a "wait" for d = 25%, q = 70%; all data points with q < 60% will have a "buy" decision, and those with q > 70% will have a "wait."

The choice-list data consists of a set of indifference pairs in the form of $(\tilde{d}_{ii}, \tilde{q}_{ii})$, where \tilde{d} is the observed discount value and \tilde{q} is the imputed indifference probability value as explained above; *j* is the index for the subject, and *i* is the index for the observation. In the context of our model, the indifference between buying and waiting given a (d, q) pair and model parameters β , μ implies

$$u - p = (u - p(1 - d)) \exp\left[-\left(\ln\frac{1}{q} + \rho e^{\mu(d_0 - d)}t\right)^{\beta}\right].$$
 (11)

Since the expression on the left is independent of q and the expression on the right is a monotone function of q, for any d there exists the implied q(d) for which Equation (11) holds. Solving for q(d) we obtain that:

$$q(d) \equiv q(d)|_{\beta,\mu}$$
$$= \exp\left\{-\left(-\ln\left(\frac{u-p}{u-p(1-d)}\right)\right)^{1/\beta} + \rho e^{\mu(d_0-d)}t\right\}.$$
 (12)

Figure 5 presents the observed indifference probabilities for different discounts; the size of the bubble is proportional to the number of subjects with the same (d_{ii}, \tilde{q}_{ii}) combination. The figure also presents the implied q(d) values for the optimal fit. Finally, it highlights an important observation that motivates how we fit our model to this data: The data points (and, hence, the resultant estimation errors) are censored. Indeed, since the indifference probability cannot be larger than 100% or smaller than zero, for small discounts, the estimated errors will be censored on the left, and for large discounts, on the right. As argued by Powell (1984), in a situation with censored observations and errors, the least absolute deviation (LAD) criterion is more appropriate than the standard least squares (LS) estimation. Thus we measure the goodness of fit between the implied and observed



Figure 5. (Color online) Observed, \tilde{q}_{ij} and implied $q(\tilde{d}_{ij})$ indifference probabilities.

probabilities for a given discount and model parameters by the absolute difference: $LAD_{ij}(\tilde{d}_{ij})|_{\beta,\mu} = |(\tilde{q}_{ij} - q(\tilde{d}_{ij}))|_{\beta,\mu}|$. We select the parameters of the pooled model to minimize the total LAD over all subjects and lists, i.e., by solving: $\min_{\beta,\mu} [\sum_{i,j} LAD_{ij}(\tilde{d}_{ij})|_{\beta,\mu}]$. Individual models can be similarly defined for each subject *j* by taking a sum over *i* only. Note that, similar to how LS regression is interpreted as a conditional mean, the LAD regression is interpreted as a conditional median (Powell 1984). The estimation was performed in VBA using the multistart generalized reduced gradient (GRG) engine in the Premium Solver Platform.

Figure 6 presents the estimated β and μ parameters. The \times s on Figure 6 represent the estimated individual β_j , μ_j pairs using the LAD method based on the choice-list data. Note that, although our model implies $\beta \leq 1$, we relaxed this constraint in fitting the model. The majority of subjects indeed had $\beta \leq 1$; seven subjects had estimated $\beta > 1$, and for another seven , the estimation problem is unbounded: Those subjects exhibit reverse subendurance and our model therefore best fits their choices at $\beta \rightarrow \infty$. The median of individual estimates, $\beta = 0.9232$, $\mu = 1.9694$, (depicted by the square on the figure) is unaffected by these few subjects.

The triangle in Figure 6 represents the pooled LAD estimate, i.e., $\beta = 0.9029$, $\mu = 1.9474$. The error bars

Figure 6. (Color online) Estimates of the parametric dPTT model, β and μ .



correspond to two standard deviations from the estimated parameter. Given that our model is highly nonlinear, and hence its error structure is unclear, we calculated standard errors by bootstrapping (the jackknife approach was used, Efron 1979; Boos and Stefanski 2013, Chapter 10). The estimated errors are 0.015 for β and 0.162 for μ . It is thus evident that the subjects significantly deviated from the fully rational estimates ($\beta = 1, \mu = 0$), but the pooled estimate and individual median are not statistically different. Both suggest that $\beta \approx 0.9, \ \mu \approx 1.95$ which, not coincidentally, were used in numerical illustrations of Section 5.1. There the dPTT sets the optimal markdown at 18.6% (versus 10.7% of DEU), increasing revenue by 1.5%, and the effectiveness of markdowns by 25%.

6.5. Robustness to Errors in Parameter Estimation

As is true for most models, an important question is whether our qualitative prediction is robust to errors in parameter estimation. For instance, if the "true" (β , μ) are different from our estimate , then how well would our suggested markdown perform? Specifically, would it still be better to use the dPTT-optimal markdown (with misestimated parameters), or should one use the DEU optimum instead, which does not have those parameters?

Figure 7(a) shows the revenue loss due to errors in parameter estimates. Here the seller applies d = 0.186 (as is optimal for $\beta = 0.9$, $\mu = 1.95$) but consumers behave as if they have different (β , μ). Note that for a large range of "true" parameters, the loss is negligible. For example, at $\beta = 0.7$, $\mu = 1.6$ the revenue loss is 0.02%, i.e., nearly 100× smaller than the gain over DEU. Figures 7(b) and (c) further explore the gains in revenue and effectiveness of markdown pricing over DEU. Significantly, the gains are positive for a very wide range of parameters; in fact, no error in estimating β could lead to DEU outperforming dPTT.

We note that sensitivity results of a similar magnitude were also presented in Özer and Zheng (2016), who also found that the improvement from their policy is approximately ten times larger than the loss from the possible misestimation of parameters. As their model also has two additional parameters, it seems that building more complex models that capture behavioral effects may indeed be superior to using DEU, which has fewer parameters, but disregards such effects.

6.6. Out-of-Sample Validation Experiment

To provide an external validation of the effectiveness of the dPTT model, we conduct an out-of-sample experiment. We use the scaled parameters from the earlier numerical illustration, $u \in U[0, 100]$, p = 50, t = 3, $\rho = 0.13$, $\omega = 0.05$. Recall that for these parameters, the optimal discount under DEU is predicted at 10.7% while under dPTT it is at 18.6%. Therefore, in the validation experiment, we compare the control group that sees a markdown price of \$45 (d = 10%) and the treatment group that sees a markdown price of \$40





Note. Parameters: $\beta = 0.9$, $\mu = 1.95$, Q = 0.625, p = 0.5, t = 3, $\lambda = 1$, $u \sim U[0, 1]$, $\omega = 0.05$, $d_0 = 0.5$ and $\rho = 0.13$.

(d = 20%). The DEU model predicts that $H^{\text{DEU}}(d = 0.1) = 0.6048$, i.e., consumers with $u = 51, \ldots, 60$ should wait, and those with $u = 61, \ldots, 100$ should buy. Similarly, $H^{\text{dPTT}}(d = 0.2) = 0.5938$, which results in identical predictions. Since experimental subjects would naturally exhibit some noise, we varied $u = 51, \ldots, 70$ to have an equal number of observations on either side of the predicted thresholds.

Each subject was provided a pair $(u, d), d = \{10\%, 20\%\}$ $u = \{51, 52, \dots, 70\}$ and was asked whether (s)he would purchase an item now at p = \$50 or wait for two months (recall that t = 3 is 9 weeks~2 months) and purchase at the corresponding markdown. The benefit of consumption was induced in the same way as in the elicitation experiment, through the ability to sell the item at u to the experimenter. The product availability in both cases was 100%, as that is the prediction at both optimal discounts. Before running the experiment, we determined via a simulation that a sample size of approximately N = 600 would be needed. Since this exceeds what is feasible in a lab, the experiment was implemented on Amazon Mechanical Turk. Subjects were exposed to randomly drawn (u, d) pairs and upon submitting their choice were given a unique code that they entered to get paid. The average hourly rate was US\$4.53.

We analyze the collected data as follows: For a given d, we call a complete set of subjects with $u = \{51, 52, ..., 70\}$ a market. We assign subjects to markets on a first-come, firstserved basis; that is, market 1 (for a given d) consists of the first arriving subject with u = 51, the first arriving subject with u = 52, ... for that d; market 2 consists of the second arriving subject with u = 51, the second arriving subject with u = 52, ... for that d, and so on. Due to randomness in generating (u, d) pairs, the data contains n = 14 markets for both d's. We padded the data with u = 71, ..., 100 subjects who buy now, and u = 45, ..., 50 for d = 10% (u = 40, ..., 50 for d = 20%) who nonstrategically wait. We then counted the subjects who wait (or buy) in each market and computed the revenue.

The observed revenue (standard deviation) at d = 10% is \$2,601.44 (19.67), and the observed revenue at d = 20%

is \$2,648.63 (30.16). As predicted, increasing the markdown from the DEU to the dPTT optimum *increased* revenue. The absolute difference between the two is \$47.19 > 0 with p < 0.01 (two-sample, double-sided, unequal variances *t*-test). Additionally, because the markets were constructed in an identical, first-come, first-served basis, a paired comparison is valid; here the, revenue with d = 20% was larger in 13 out of 14 markets. The relative improvement of $47/2,601 \approx 1.8\%$ is very similar to the theoretically anticipated improvement of 1.5%. Furthermore, since the revenue of \$2,500 can be obtained without dynamic pricing at all, the incremental revenue from markdowns increased from \$101 to \$148, or by nearly 50\%, clearly a sizable gain in effectiveness of markdown pricing.

To summarize, the validation experiment confirms the main finding of our paper: due to the behavioral anomalies that affect how people decide to wait or buy, retailers should offer higher markdowns than the standard theory suggests, and by doing so, obtain higher revenue.

7. Conclusion

The importance of markdown management to modern retailers is hardly a question: Nearly 1/3 of unit sales and 1/5 of dollar sales are generated at markdowns (Smith and Achabal 1998, Agrawal and Smith 2009). Furthermore, with retailers' net profit margins being approximately 3%, each percent of extra markdown revenue translates into major profit increases. The proliferation of markdowns, however, fueled strategic waiting: effectively, every time a consumer enters the store (s)he mentally "solves" a wait-or-buy problem: should (s)he buy the item now or wait for a possible markdown? Being aware of the behavioral regularities surrounding this decision, and incorporating them into markdown management offers substantial revenue opportunity for retailers.

In this paper, we study this fundamental wait-or-buy problem from a unique, behavioral, perspective. The core idea of our study is that the wait-or-buy decision reflects a multidimensional trade-off between the delay in getting an item, the likelihood of getting it, and the magnitude of the price discount. Multiple studies in decision analysis, psychology, and behavioral economics, showed that all these trade-offs are prone to behavioral regularities by which decision makers deviate from the discounted expected utility model used in the current literature.

We present behavioral preference conditions (axioms) that support a modification of the discounted expected utility model. Our preference conditions capture three behavioral anomalies widely documented in laboratory experiments, i.e., the common ratio effect in risk perception, the common difference effect in time perception (a.k.a., hyperbolic discounting), and the magnitude effect in time discounting (a.k.a., subendurance). Key in our formulation is the concept of psychological distance. The result is a parsimonious modification of the discounted expected utility. We solve the consumer's wait-or-buy problem and embed it into the firm's markdown optimization problem. We calibrate the model parameters using experimental data, validate it out-ofsample, and show that accounting for the behavioral anomalies results in substantially larger markdowns that the current literature suggests and leads to noticeable revenue gains.

Supplemental Material

Supplemental material to this paper is available at https://doi.org/ 10.1287/opre.2016.1547.

Acknowledgments

The authors thank Casey Lichtendahl (University of Virginia), Laurence Ashworth (Queen's University), Dan Adelman (University of Chicago), the associate editor, and three anonymous referees for their constructive feedback. M. Baucells received support from the Spanish Ministerio de Economía y Competitividad [Grant PSI2013-41909-P].

Appendix. Experimental Instructions

You are about to participate in an experiment in the economics of decision making. There are no right or wrong answers; just express your preferences. By doing so you can earn a substantial amount of money that will be paid to you as is explained below. If you have a question at any time, please raise your hand and the experimenter will answer it; do not talk with one another for the duration of the experiment.

Overview of the Experiment

The context of the experiment is the following:

Wait or Buy? Suppose that you went to a retail store and saw a product that you know you can resell for \$250 at any time. The product was priced at \$200 (two hundred dollars), so you picked the product from the shelf and were about to purchase. However, then you started thinking that in three weeks from today this product may be marked down. Thus the question was: Should you buy the product now or wait for the markdown?

As you entered the room you selected a sealed envelope. Do not open the envelope until the end of the experiment. The envelope contains two numbers: (1) the markdown percentage and the (2) likelihood that the product will be available when you visit the store again in three weeks. Either number could vary between 5% and 90%. As yet, you do not know which two numbers are inside your envelope.

With the help of the experimental interface (see the screenshot and the link on the next page) you will give the experimenter instructions as to whether you would like to buy the product now or wait three weeks until the markdown for each possible combination of the markdown percentage and the likelihood of availability that could be inside your envelope. You will first be asked five questions, each containing a single markdown percentage and nine different likelihoods of availability. Your answers to those questions will determine your pay. We will then ask you six additional questions each containing a single markdown percentage and a single likelihood of availability just to double-check your answers.

How You Will Be Paid

At a minimum you will be paid \$5 just for participation in this experiment. At the maximum you can earn over \$200. We will determine how much you will earn as follows: At the end of the experiment today we will randomly select two participants. These selected participants will stay with the experimenter to complete the process; everyone else will collect their \$5 and leave.

The experimenter ("we") and each selected participant ("you") will in private complete the following procedure. First, we will open your envelope and find out your markdown percentage and the likelihood of availability. Second, we will look up the Waitor-Buy choice you made during the experiment for that specific markdown percentage and likelihood of availability.

• If your choice was "Buy now" then you will buy the product for \$200 and immediately resell it to the experimenter for \$250, keeping the remaining \$250 - 200 = \$50.

• If your choice was "Wait for the markdown" then in three weeks (i.e., on April 23, 2014) you will have to stop by Prof. [Removed by the authors for confidentiality] office to learn about the product availability.

—Product availability will be determined by drawing a random number between 1 and 100 (all numbers being equally likely) and comparing it to the likelihood of availability from your envelope. If the random number is less or equal to the likelihood from the envelope, the product will be determined to be "available" or otherwise "unavailable."

—If the product is available, then its price will be adjusted according to the markdown percentage from your envelope, you will pay the adjusted price, immediately resell the product to the experimenter for \$250 and keep the rest of the money. For example, if the markdown percentage is 50% then the adjusted price will be $$200 \times 50\% = 100 , and after re-selling the product you will keep the remaining \$250 - 100 = \$150.

—If the product is unavailable, then you will have nothing to resell and thus there will be no additional money.

All money will be paid to you in cash. All decisions and earnings are confidential.

Screenshot of the Experimental Interface

[Removed by the authors to save space; it is identical to Figure 8(a) except that the actual markdown percentage was replaced with XX to ensure that all values receive equal amount of subjects attention.]

Link: [Removed by the authors for confidentiality.]

Figure 8. (Color online) Screenshots of the choice list (a) and binary choice (b) questions.

Recall that if you buy the product, you can re-sell it for \$250.

The product is currently priced at \$200. The store has announced that in three weeks the product will be marked down by **25%**. The problem, however, is that it may be sold-out by then. From your experience, the likelihood that the product will be available could vary between 10% and 90%.

Considering different possible likelihoods of the product being available, would you buy now or wait for the markdown?

	Buy now	Wait for the markdown
If there is a 10% chance that the product will be available at the markdown	•	•
If there is a 20% chance that the product will be available at the markdown	•	•
If there is a 30% chance that the product will be available at the markdown	•	•
If there is a 40% chance that the product will be available at the markdown	•	0
If there is a 50% chance that the product will be available at the markdown	0	•
If there is a 60% chance that the product will be available at the markdown	0	0
If there is a 70% chance that the product will be available at the markdown	•	•
If there is a 80% chance that the product will be available at the markdown	0	0
If there is a 90% chance that the product will be available at the markdown	۲	0
	(a)	

Thank you for making your choices. Now lets check some of your responses.

Next

Next

Recall that if you buy the product, you can re-sell it for \$250.

The product is currently priced at 200. The store has announced that the product will be marked down by 25% and there is a 50% chance that it will be available in three weeks.

Would you buy now or wait for the markdown?



Wait for the markdown

0

(b)

Endnotes

1. That risk and time distance are substitutes yields two nontrivial predictions, i.e., that the common ratio effect can be reproduced by adding a common delay to both options; and that the common difference effect can be reproduced using a common probability reduction for both options. Intuitively, both manipulations increase the distance $\sigma = \ln(1/q) + r(d)t$, and because $s(\sigma)$ is concave,

they induce loss of sensitivity with respect to probability or time and make the payoff dimension more salient. In particular, adding a common delay of three months to choice 1 in Table 1 would produce a reversal. Indeed, Baucells and Heukamp (2010, Table 1) show that only 43% of subjects prefer (9€, for sure, 3 months) to (12€, with 80%, 3 months). Similarly, adding a common reduction of probability to choice 3 in Table 1 would produce a reversal. Indeed, Keren and Roelofsma (1995, Table 1) show that only 39% of subjects prefer (100 fl, with 50%, now) to (110 fl, with 50%, 4 weeks).

2. The fluid model is a limiting case of multiple probabilistic demand when the demand rate and capacity grow proportionally large (Maglaras and Meissner 2006).

References

- Agrawal N, Smith SA (2009) *Retail Supply Chain Management* (Springer, New York).
- Allais M (1953) Le comportement de l'homme rationnel devant le risque: Critique des postulats et axioms de l'école américaine. *Econometrica* 21(4):503–546.
- Aviv Y, Pazgal A (2008) Optimal pricing of seasonal products in the presence of forward-looking consumers. *Manufacturing and Service Oper. Management* 10(3):339–359.
- Baucells M, Heukamp F (2010) Common ratio using delay. *Theory and Decision* 68(1-2):149–158.
- Baucells M, Heukamp FH (2012) Probability and time trade-off. Management Sci. 58(4):831–842.
- Baucells M, Heukamp F, Villasís A (2009) Trading-off probability and time: Experimental evidence. IESE Business School Working paper.
- Baucells M, Weber M, Welfens F (2011) Reference point formation and updating. *Management Sci.* 57(3):506–519.
- Besanko D, Winston W (1990) Optimal price skimming by a monopolist facing rational consumers. *Management Sci.* 36(5):555–567.
- Bils M, Klenow P (2004) Some evidence on the importance of sticky prices. J. Political Econom. 112(5):947–985.
- Booij A, van Praag B, van de Kuilen G (2010) A parametric analysis of prospect theory's functionals for the general population. *Theory and Decision* 68(1-2):115–148.
- Boos DD, Stefanski LA (2013) Essential Statistical Inference: Theory and Methods (Springer, New York).
- Correa J, Montoya R, Thraves C (2015) Contingent preannounced pricing policies with strategic consumers. Oper. Res. 64(1):251–272.
- Damodaran A (2015) Operating and net margins by sector (US). Accessed September 2015, http://pages.stern.nyu.edu/~adamodar/New_Home _Page/datafile/margin.html.
- Darke PR, Freedman JL (1993) Deciding whether to seek a bargain: Effects of both amount and percentage off. J. Applied Psych. 78(6):960–965.
- DellaVigna S, Malmendier U (2004) Contract design and self-control: Theory and evidence. Quart. J. Econom. 119(2):353–402.
- Ebert JE, Prelec D (2007) The fragility of time: Time-insensitivity and valuation of the near and far future. *Management Sci.* 53(9): 1423–1438.
- Efron B (1979) Booststrap methods: Another look at the jackknife. Ann. Statist. 7(1):1–26.
- Holt C, Laury S (2002) Risk aversion and incentive effects. Amer. Econom. Revi. 92(5):1644–1655.
- Huang T, Liu Q (2014) Is capacity rationing optimal when customers use anecdotal reasoning? Working paper.
- Johnson C, Baillon A, Bleichrodt H, Li Z, van Dolder D, Wakker PP (2014) The prior incentive system. Erasmus University Working paper.
- Kalish S (1983) Monopolist pricing with dynamic demand and production cost. *Marketing Sci.* 2(2):135–159.
- Keeney R, Raiffa H (1976) Decisions with Multiple Objectives: Preference and Value Tradeoffs (John Wiley & Sons, New York).
- Keren G, Roelofsma P (1995) Immediacy and certainty in intertemporal choice. Organ. Behav. Human Decision Processes 63(3):287–297.
- Kim S, Dasu S (2014) An experimental study of posted prices. Working paper.
- Laibson D (1997) Golden eggs and hyperbolic discounting. Quart. J. Econom. 112(2):443–477.
- Li J, Granados N, Netessine S (2014) Are consumers strategic? Structural estimation from the air-travel industry. *Management Sci.* 60(9): 2114–2137.

RIGHTSLINK4)

- Liu Q, van Ryzin G (2008) Strategic capacity rationing to induce early purchases. *Management Sci.* 54(6):1115–1131.
- Maglaras C, Meissner J (2006) Dynamic pricing strategies for multiproduct revenue management problems. *Manufacturing Service Oper*. *Management* 8(2):136–148.
- Mak V, Rapoport A, Gisches EJ, Han J (2014) Purchasing scarce products under dynamic pricing: An experimental investigation. *Manufacturing* and Service Oper. Management 16(3):425–438.
- Nasiry J, Popescu I (2011) Dynamic pricing with loss-averse consumers and peak-end anchoring. Oper Res. 59(6):1361–1368.
- O'Donoghue T, Rabin M (1999) Doing it now or later. Amer. Econom. Rev. 89(1):103-124.
- Ortmeyer G, Quelch JA, Salmon W (1991) Restoring credibility to retail pricing. *Sloan Management Rev.* 33(1):55–66.
- Osadchiy N, Bendoly E (2010) Are consumers really strategic? Experimental study. Working paper.
- Osadchiy N, Vulcano G (2010) Selling with binding reservations in the presence of strategic consumers. *Management Sci.* 56(12):2173–2190.
- Ovchinnikov A, Milner JM (2012) Revenue management with end-ofperiod discounts in the presence of customer learning. *Production Oper. Management* 21(1):69–84.
- Özer Ö, Zheng Y (2012) Behavioral issues in pricing management. Özer Ö, Phillips R, eds. Oxford Handbook of Pricing Management (Oxford University Press, Oxford, UK).
- Özer Ö, Zheng Y (2015) Markdown or everyday low price? The role of behavioral motives. *Management Sci.* 62(2):326–346
- Pilehvar A, Elmaghraby W, Gopal A (2016) Market information and bidder heterogeneity in secondary market online B2B auctions. *Management Sci.*, ePub ahead of print June 3, 2016, https://doi.org/10.1287/ mnsc.2016.2444.
- Popescu I, Wu Y (2007) Dynamic pricing strategies with reference effects. Oper. Res. 55(3):413–429.
- Powell JL (1984) Least absolute deviations estimation for the censored regresson model. J. Econometrics 25(3):303–325.
- Prelec D (1998) The probability weighting function. *Econometrica* 66(3): 497–527.
- Savage LJ (1954) The Foundations of Statistics (John Wiley & Sons, New York).
- Smith SA, Achabal DD (1998) Clearance pricing and inventory policies for retail chains. *Management Sci.* 44(3):285–300.
- Starmer C, Sugden R (1993) Testing for juxtaposition and event-splitting effects. J. Risk and Uncertainty 6(3):235–254.
- Su X (2009) A model of consumer inertia with applications to dynamic pricing. Production Oper. Management 18(4):365–380.
- Su X (2010) Intertemporal pricing and consumer stockpiling. *Oper. Res.* 58(4):1133–1147.
- Swinney R (2011) Selling to strategic consumers when product value is uncertain: The value of matching supply and demand. *Management Sci.* 57(10):1737–1751.
- Tereyagoglu N, Fader PS, Veeraraghavan SK (2014) Multi-atribute loss aversoin and reference dependence: Evidence from the performing arts industry. Working paper.
- Thaler RH (1985) Mental accounting and consumer choice. Marketing Sci. 4(3):199–214.
- Tversky A, Kahneman D (1981) The framing of decisions and the psychology of choice. Science 211(4481):453–458.
- Tversky A, Kahneman D (1992) Advances in prospect theory: Cumulative representation of uncertainty. J. Risk and Uncertainty 5(4):297–323.
- von der Fehr N, Kuhn K (1995) Coase versus Pac-Man: Who eats whom in the durable-goods monopoly? J. Political Econom. 103(4):785–812.
- Wakker PP (2010) Prospect Theory: For Risk and Ambiguity (Cambridge University Press, New York).
- Zhang D, Cooper W (2008) Managing clearance sales in the presence of strategic consumers. *Production Oper. Management* 17(4):416–431.

Manel Baucells is an associate professor at the University of Virginia Darden School of Business. His research focuses on incorporating psychological realism into consumer behavior models by considering factors such as anticipation, reference point comparison, mental accounting, psychological distance, and satiation. Manel and Rakesh Sarin (UCLA) have published *Engineering Happiness* (UC Press), which has received the 2014 Best Publication Award by the Decision Analysis society. **Nikolay Osadchiy** is an assistant professor at the Goizueta Business School, Emory University. His research interests are revenue management, behavioral operations, and supply chain management, wherein he studies how supply networks affect risk and operational performance.

Anton Ovchinnikov is a Distinguished Faculty Professor of Operations Management and Management Science and a Scotiabank Scholar of Customer Analytics at the Smith School of Business, Queen's University, Canada. His research interests include, on the theoretical side, behavioral operations, revenue management, and environmental sustainability. On the applied side, he studies datadriven applications in business, government, and nonprofit sectors. Anton has published in leading academic journals and was recognized as the winner/finalist of the INFORMS Case Competition, JFIG Best Paper award, RMP Section Practice Prize, and the POMS Paul Kleindorfer award.