

Inventory Decisions in the Presence of Strategic Customers: Theory and Behavioral Evidence

Yang Zhang

Department of Industrial Engineering and Logistics Management
The Hong Kong University of Science and Technology

ielmyzhang@ust.hk

Benny Mantin

Luxembourg Centre for Logistics and Supply Chain Management
University of Luxembourg

benny.mantin@uni.lu

Yaozhong Wu

National University of Singapore Business School

yaozhong.wu@nus.edu.sg

June 18, 2018

Abstract

We consider the inventory decision of a retailer facing strategic customers. We develop a behavioral theory that accounts for reference dependence, which makes predictions on how the presence of strategic customers leverages retailer behavior. Specifically, the reference-dependent retailer shall decrease her order quantity when there are more strategic customers in the population. As such, the conventional pull-to-center bias for newsvendor is generalized, since the presence of strategic consumers may induce the retailer to pull her order further below the center even when the production cost is low. Furthermore, increasing proportion of strategic customers reduces the retailer's ordering bias under low cost, yet amplifies it when the cost is high. Our subsequent experiments find pull-below-center effect, validate the theoretical predictions,

and establish the asymmetry of reference dependence with the estimated behavioral parameters. We also study extensions to our model and carry out robustness checks of our experimental results.

Key words: newsvendor problem, strategic customers, behavioral theory, reference dependence, pull-below-center bias

History: Received: November 2016; accepted: June 2018 by Elena Katok after two revisions.

1 Introduction

Since the seminal work of Schweitzer and Cachon (2000), human behavior in inventory decision making has received extensive attention, primarily within the newsvendor framework. The conventional newsvendor setting elegantly captures a core tradeoff between overstocking and understocking in inventory management, and its empirical tests reveal that in practice actual decision makers exhibit consistent biases, such as placing orders that deviate away from the optimal order towards the mean demand. Despite its elegance, the setting abstracts away from two important considerations widely observed in practice: markdowns and the consequential strategic customer behavior that emerges in such pricing environment. Markdown pricing, which often follows a predetermined pricing path, has been established as a mechanism to clear excess inventory (Fisher and Raman, 2010), while at the same time a significant proportion of customers strategically delay their purchase in expectation of future markdowns (ranging from 5%–20% according to Li et al. (2014) to 77% according to Osadchiy and Bendoly (2015)).¹ Hence, it comes at no surprise that an ever-increasing share of retail sales is transacted at marked down prices (see Phillips (2005) for the industry statistics).

The presence of strategic customers exposes the newsvendor to new challenges. In principle, the customers' strategic behavior hinges critically on the likelihood of the newsvendor having leftover inventory at the end of the main selling season, so that the excess inventory will be marked down. Hence, if the newsvendor carefully chooses her order quantity she can mitigate (some of) the impact of strategic customer behavior on profit. Namely, a low order quantity can discourage strategic customer waiting for the markdown period. In this paper, we focus on the behavioral aspects of inventory decisions making in the presence of strategic customers: How does the presence of strategic customers change the ordering patterns and biases of the newsvendor? To address this research question, we first develop a

¹The estimation of the proportions of strategic and myopic customers has been carried out in both empirical works (e.g., Nair 2007; Chevalier and Goolsbee 2009; Li et al. 2014, using video games, textbooks and airfares, respectively) and laboratory experiments (e.g., Mak et al. 2014; Osadchiy and Bendoly 2015).

modeling framework that captures the interaction between a newsvendor and customers in a two-period predetermined markdown setting while accounting for the newsvendors potential behavioral biases; we then test our model in a lab setting.

Our modeling framework enriches the newsvendor setting to introduce a second period where the leftover inventories are priced to a predetermined markdown. Predetermined price trajectories are common in practice (see, e.g., Talebian and van Ryzin 2014) and may apply when the retailer has limited pricing flexibility, such as in the case of the manufacturer's suggested retail price (MSRP). They are also commonly assumed in the revenue management literature (e.g., Gallego et al. 2008; Zhang and Cooper 2008; Mersereau and Zhang 2012; Altug and Aydinliyim 2016). Taking advantage of the possible discounted leftover inventory, some customers may strategically postpone their purchase. These are the strategic customers. Consistent with the literature, we assume the customer population consists of both myopic as well as strategic customers (Su 2007; Zhang and Cooper 2008; Gallego et al. 2008; Zhang and Cooper 2008; Mersereau and Zhang 2012).

Our behavioral model extends the basic setup from Cachon and Swinney (2009) by incorporating the reference dependent preference, which associates separate psychological losses with overstocking and understocking. Reference dependence has been established as a common bias in decision making (Kahneman and Tversky, 1979) and is demonstrated to be effective in explaining the behaviors in newsvendor problem (Ho et al., 2010), contracting (Ho and Zhang, 2008), brand choice (Hardie et al., 1993), and relation to dynamic pricing strategies (Popescu and Wu, 2007).

Our reference dependence behavioral theory predicts three important regularities of the retailer's orders. First, the retailer should decrease her inventory level as the proportion of strategic customers in the population increases, so as to induce strategic customers to purchase early at the full price rather than wait for a discount and face the risk of stockout. This result holds regardless of the retailer's degree of rationality (i.e., it holds when she is fully rational or when she is subject to reference dependence bias to any extent). Second, we generalize the pull-to-center bias. While in a high cost setting we show that the pull-to-center

effect may prevail as expected, in a low cost setting, the presence of strategic customers may induce a pull-below-center bias. We refer to these combined effects as the generalized pull-to-center bias. Third, the generalized pull-to-center effect varies in opposite ways with the proportion of strategic customers under different cost conditions. When the unit cost is low, the generalized pull-to-center bias decreases in the proportion of strategic customers in the population, whereas this bias increases in the proportion of strategic customers when the product unit cost is high. These predicted regularities are validated in subsequent laboratory experiments.

We conduct laboratory experiments to test our theory. As such, our paper contributes to the body of literature, initiated by Schweitzer and Cachon (2000), investigating the newsvendor decision making experimentally. This stream of research devotes efforts to characterizing and explaining the biases that human newsvendors exhibit when placing orders, as well as the ways in which these biases can be corrected, or explored for benefits. A central finding of the newsvendor behavior is the pull-to-center bias—subjects tend to shift away their order quantity from the optimum towards the average of market demand (see, e.g., Schweitzer and Cachon 2000, Su 2008, Bostian et al. 2008, Kremer et al. 2010, Ho et al. 2010, Ren and Croson 2013, Chen et al. 2013). Since the pull-to-center bias induces suboptimal retailer orders, several papers propose potential remedies to pull-to-center bias (e.g., Bolton and Katok 2008; Becker-Peth et al. 2013) or transformation of this bias into a competitive advantage (e.g., Ovchinnikov et al. 2015).²

Does a similar pull-to-center bias persist in the presence of strategic customers? As predicted by our model, the order quantity is pulled towards the mean demand if the ratio of psychological costs falls within some prescribed range; importantly, however, the order is pulled *beyond* the mean demand, if there are sufficiently many strategic customers in the market and if there is sufficient asymmetry in the psychological losses. In that sense, we

²Although our paper is primarily related to the research on a single newsvendor problem, we are also aware of the many experimental papers that reach out to competing newsvendors (Li et al. 2016; Ovchinnikov et al. 2015) and vertical coordination with newsvendor (Davis 2015; Davis et al. 2014; Katok and Pavlov 2013; Katok and Wu 2009), among other topics.

generalize the notion of pull-to-center behavior in the classic newsvendor environment to capture deviations from the well-established pull-to-center bias. Our experiments exhibit both pull-to-center bias and pull-below-center bias, thus validating the theoretical predictions. Also, we find that our generalized pull-to-center bias is stronger with lower unit cost, which parallels the regularity in newsvendor behavior (Schweitzer and Cachon 2000, Ho et al. 2010). This, too, can be explained by the unbalanced psychological costs. In this way, our behavioral model extends (the single-newsvendor version of) the setting studied in Ho et al.'s (2010) into the context of strategic customers.

Prior literature (e.g. Ho et al. 2010) has established the asymmetry of reference dependence (the marginal psychological cost due to wasted inventory greater than that for lost demand) in traditional newsvendor behavior. If the asymmetry is adequately severe, the newsvendor will order less than optimum in the low cost condition. Although this bends the order towards the mean demand, but as we observed in the vast literature of newsvendor experiments, the order seldom goes beyond the mean demand. In our study, the introduction of strategic customers gives another incentive for the retailer to lower her quantity, and consequently we observe pull-below-center effect at low cost. Therefore, it is plausible that the asymmetry of psychological costs alone is *not* sufficient in driving the retailer orders below the average demand, and pull-below-center effect can only take place with the presence of strategic customers.

Our paper is also akin to the ever-growing literature on revenue management that considers strategic customer behavior (see, e.g., Aviv and Pazgal 2008; Gallego et al. 2008; Su 2007; Liu and van Ryzin 2008; Mersereau and Zhang 2012; and a summarized review in Kremer et al. 2017). In particular, there are papers that examine behavioral issues in revenue management, which shares some similarities with the newsvendor problem in their settings (Kocabiyikoglu et al., 2015). Bearden et al. (2008) and Bendoly (2011) inspect biases and judgment errors in accepting and rejecting bids with a revenue management setup. Ozer and Zheng (2016), Song and Zhao (2017, 2016) and Kim and Dasu (2013) investigate the strategies of boundedly rational customers for a given price and inventory level, as well as

their implications for the operations of the selling firm. Hariss et al. (2016) accounts for consumers' price-based quality perceptions to reveal that retailers are better off announcing and committing to a markdown strategy. Cason and Mago (2013), Mak et al. (2014), and Kremer et al. (2017) focus on the sellers' dynamic pricing decisions, where the inventory is exogenously endowed. Differing from these papers, we focus on the inventory, rather than pricing, decisions of a seller facing strategic customers, and we relate our findings to the common pull-to-center bias documented in the literature of newsvendor experiments.

2 Retailer's reference dependence in the presence of strategic consumers

A traditional newsvendor is faced with the fundamental trade-off of overordering and underordering costs when setting the optimal ordering decisions. In multi-period environments when excess inventory is marked down over time and in the presence of strategic consumers, arriving at the optimal ordering decision is more challenging as strategic consumers may wait for the discount. The actual decisions may be subject to reference dependent preferences. Below, we first outline the basic model and then we incorporate reference dependence in the retailer's decision making.

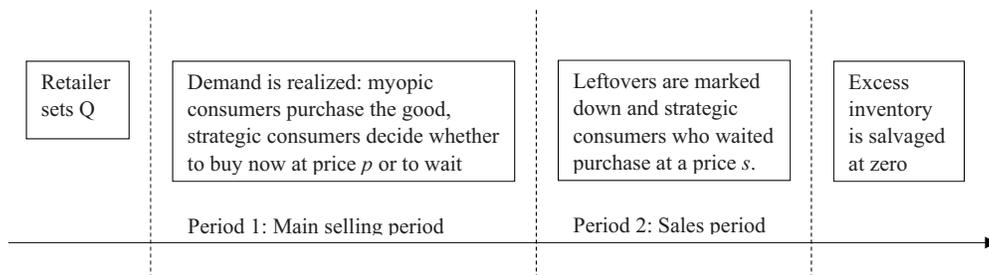
2.1 Model Overview

Our base model is a simplified version of Cachon and Swinney's (2009) model. The sequence of events in the model, which is depicted in Figure 1, is as follows. A retailer sells her products to myopic and strategic customers over a selling horizon consisting of two periods. In the first period—called the *selling* period—the retail price of the product is p , and any leftover units from the selling period are offered during the second period—the *sales* period—at a discounted price s , $s < p$. Both prices are predetermined and known to the customers.³ At

³In practice, predetermined retail prices apply when downstream members of the supply chain have limited pricing flexibility. For example, the practice of setting the manufacturer's suggested retail price (MSRP) or retail price maintenance (RPM), which are widely used in print, music, and automotive industries, limit retailers' pricing flexibility. Some recent studies also feature predetermined pricing paths, e.g., Altug and Aydinliyim (2016), Liu and van Ryzin (2008) (in their main model), among others.

the beginning of the selling horizon, before knowing the actual demand, the retailer places and receives her order of Q units at a unit ordering cost c . Then, during the main selling period, the total demand is realized (consisting of both consumer types), and purchases are made by myopic customers and strategic customers who decide to buy during the selling period. Any leftover inventory is carried over to the sales period;⁴ otherwise, the game ends. During the sales period, purchases are made by strategic customers who delayed their purchases. Any leftover inventory from the sales period is disposed at zero value and the game ends.

Figure 1: Sequence of events



The demand—i.e., the number of customers that show up, M —is a random variable that follows a probability distribution with cdf. $F(\cdot)$ and pdf. $f(\cdot)$ with a mean of μ , and its realization is denoted by m . We assume that each customer purchases at most one unit of the product. We impose the following assumption regarding the demand.

Assumption 1. *The demand distribution has an increasing generalized failure rate (IGFR).*

Let $\bar{\theta}(m) := \frac{mf(m)}{1-F(m)}$. A probability distribution function has an increasing generalized failure rate (IGFR) when $\bar{\theta}(m)$ weakly increases in m (Lariviere, 2006; Banciau and Mirchandani, 2013). The concept of IGFR is broader than that of the increasing failure rate (IFR), in which the failure rate $\frac{f(m)}{1-F(m)}$ increases in m , as the set of IGFR distributions entails all

⁴We do not consider the opportunity of inventory replenishment at the beginning of the sales period. Note that the retailer has little incentive to replenish, even if there is a stockout in the selling period, since strategic customers expecting a replenishment will be more inclined to wait until the sales season. In fact, strategic customers in our base model will possibly make a purchase in the sales period only if the firm orders more than the total customer demand, in which case no replenishment is needed for the sales period. See Section 2.2 for details.

IFR distributions (such as normal and uniform distributions), but it is not limited to those. In fact, many distributions with a decreasing failure rate also feature an IGFR. IGFR has been shown to be a useful property for many OM problems. We refer the reader to Banciau and Mirchandani (2013) and Lariviere (2006) for further discussion.

The market consists of two types of customers: *myopic* and *strategic*, where the fraction of strategic [resp., myopic] customers in the market is denoted by α [resp., $1 - \alpha$], $0 \leq \alpha \leq 1$, and is public information (which could be learned through market surveys, media, prior experience, etc.). Both customer types have a valuation v_M for the product during the selling period, $v_M \geq p$. Each of the myopic customers purchases in the main selling period, as his myopia excludes forward looking into opportunities of buying at a discounted rate. A strategic customer decides on the timing of his purchase by choosing between purchasing a product during the selling period at the retail price p and waiting to buy during the sales period at the discounted price, s . Customers' valuations for the product deteriorates over time. Specifically, customers' valuation v in the sales period exhibits heterogeneity, such that v follows a probability distribution on the interval $[\underline{v}, \bar{v}]$ with a cdf. $G(\cdot)$ and a pdf. $g(\cdot)$ such that $s < \underline{v} < \bar{v} < v_M$.⁵ We also assume that both customer types have an outside option yielding zero utility. The option of not buying a product at all is excluded from customers' choice sets. In the event that the total number of purchasing customers in either period is higher than the retailer's inventory, each customer has an equal probability of obtaining a unit of the available product.

The following aspects of our baseline model are simplified from Cachon and Swinney (2009). First, we assume that both the retailer's order quantity, Q , and the total customer arrivals, M , are public knowledge. Second, we impose an equilibrium selection rule (Assumption 2 in §2.2) when there are multiple equilibria of strategic customer behavior. These

⁵As illustrated in Cachon and Swinney (2009), customer heterogeneity in our model has two alternative specifications, both referring to δv as the discounted valuation in the sales period ($\delta \in (0, 1)$ represents the discount factor). One specification allows heterogeneity and randomness in δ and keeps v constant. The other assumes v is random, while the discount factor remains the same for all customers. These two alternative models produce similar insights, as does the current model.

simplifications facilitate the design of an experiment that focuses on the retailer’s ordering behavior, so that we can test our theory in the laboratory.⁶ In the Online Supplement that accompanies this manuscript, we show that the major insights of our behavioral model are robust to alternative setups, such as incomplete information on demand and inventory (Appendix 2.1), and alternative equilibrium of customer game (Appendix 2.2). Third, we assume an exogenous pricing path. This allows us to focus on the study of the retailer’s inventory decisions. As aforementioned, this assumption applies when the retailer has little pricing flexibility (e.g. MSRP). Fourth, in order to keep the parsimony of the experimental design we exclude the “bargain hunters”—the class of customers who only show up during the sales period. As a result, the complexity of customer population in our experiment is controlled by a single parameter α .

We solve the game described above via backwards induction. That is, we first address the decision problem faced by the strategic consumers at the beginning of the main selling period (§2.2); then we return to the study of the retailer’s ordering decision (§2.3).

2.2 The strategic customers’ problem

If a strategic customer chooses to purchase in the main selling period, then his expected utility is $\varphi_1(v_M - p)$, where $\varphi_1 \in [0, 1]$ is the probability of obtaining the product in that period. If the strategic customer chooses to wait and purchase in the sales period, then his expected utility is $\varphi_2(v - s)$, where $\varphi_2 \in [0, 1]$ is the probability of obtaining the product in the sales period. The strategic customers weigh these two expected utilities to decide when to purchase. The probabilities φ_1 and φ_2 depend on the total number of purchasing customers (strategic and myopic), as well as on the retailer’s inventory level. The following

⁶Under the assumption of complete information, the retailer and customers play a sequential game. This enables us to design the experiment such that the subjects (who play the role of the retailer) interact with computerized strategic customers that are automated to adopt the best response to the retailer’s ordering decision. This design provides the desired control for a focused study on retailer behavior. Such design is not feasible under incomplete information (where the customer does not know the order quantity or the number of customers) because, in that case, the game becomes *de facto* simultaneous between customers and the firm (c.f. Cachon and Swinney 2009). It is then impossible to separate the decisions and control one another. The equilibrium selection assumption is detailed and explained in Section 2.2.

Lemma shows that the equilibrium strategy of strategic customers can be characterized using a threshold strategy.

Lemma 1. *In the equilibrium of the customer game, there exists a threshold valuation v^* , such that all strategic customers whose valuation v in the sales period is lower [higher] than v^* purchase in the selling [sales] period. A strategic customer with value v^* is indifferent to purchasing in either period.*

All proofs are delegated to the appendix. The equilibrium threshold v^* in the lemma satisfies $\varphi_1^*(v_M - p) = \varphi_2^*(v^* - s)$, where φ_1^* and φ_2^* are equilibrium probabilities of obtaining a product in the selling and sales periods, respectively. In some cases, the threshold v^* may fall out of the bounds so that the equilibrium reduces to one in which all strategic customers purchase in the selling [sales] period, hereafter referred to as the *all-buy* [*all-wait*] equilibrium. Lemma 2 fully characterizes the equilibrium of the customer game in all scenarios.

Lemma 2. *Given an order quantity Q and a realized demand m , the equilibrium of the customer game takes the following forms:*

1. *If $Q \leq (1 - \alpha)m$, the unique equilibrium is an all-buy equilibrium;*
2. *If $Q > (1 - \alpha)m$, there exists a threshold such that a strategic customer whose valuation in the sales period is smaller [higher] than the threshold purchases in the selling [sales] period. Specifically,*
 - a. *if $(1 - \alpha)m < Q < m$, the equilibrium can be either an all-buy equilibrium or a threshold equilibrium with threshold $v_{(1-\alpha)m,m}^*$, where $v_{(1-\alpha)m,m}^*$ solves*

$$v_M - p = \frac{Q - (1 - \alpha)m - \alpha m G(v_{(1-\alpha)m,m}^*)}{\alpha m (1 - G(v_{(1-\alpha)m,m}^*))} (v_{(1-\alpha)m,m}^* - s);$$

- b. *if $Q \geq m$, the equilibrium is characterized by threshold $v_{m,\infty}^* := v_M - p + s$.*

For the interest of our core research question, in the ensuing analysis and experimental study, we focus on the all-buy equilibrium as the outcome of strategic customer behavior when multiple equilibria exist. This assumption is formally stated as follows.⁷

Assumption 2. *The all-buy equilibrium emerges in the customer game when $(1 - \alpha)m < Q < m$.*

Assumption 2 captures some customer behavior that exists in reality. For example, customers may “panic” when there is uncertainty in getting the goods (when $Q < m$), and they herd into purchases because they believe others will do the same.⁸ As an assumption on equilibrium selection, Assumption 2 removes the ambiguity in equilibrium prediction when $(1 - \alpha)m < Q < m$. The all-buy equilibrium is implemented in our experimental study, particularly for its ease of presentation, which reduces the subjects’ mental burden in understanding how their order quantities may shift demand.

2.3 The Retailer’s Problem

Our focus in this paper lies in the retailer’s problem. We are interested in the biases involved in the retailer’s ordering decision, and how such biases are leveraged by strategic consumers. Specifically, we consider the retailer’s reference dependence in making the ordering decision. Reference dependence is particularly relevant to our setting, as the retailer, upon placing her order, may be aiming at the natural target: the realized demand. Specifically, the retailer evaluates outcomes using the realized demand as her reference point, and incurs psychological losses if her order quantity differ from the realized demand.⁹ Another plausible behavioral module that could play a role in our setting is mental accounting, which could be present

⁷Relaxation of Assumption 2 is discussed in Section 3.2 and the full analysis is provided in the Online Supplement 2.2, where we allow the customers to play the (interior) threshold equilibrium when it exists. Therein we show that the major insights of the behavioral model persist.

⁸For example, Zara’s “fast fashion” strategy creates a sense of scarcity, which may make customers rush into early purchases.

⁹There are alternative choices of modeling the behavior, such as using profit as reference point (Uppari and Hasija 2016, Long and Nasiry 2015). However, we choose the present setup for its parsimony, and its intuitive interpretation of experimental results. Further extension of the behavioral model, such as embracing the full form of prospect theory, would be a priority for future research.

due to the retailer's interaction with two different customer types—myopic and strategic. However, the empirical analysis does not lend sufficient support to this potential behavioral bias (see Footnote 10).

Our model follows that of Ho et al. (2010), while integrating the multi-period nature of the selling environment and the presence of strategic customers. Let δ_o denote the psychological cost associated with each overordered unit and let δ_u denote the psychological cost associated with each underordered unit.¹⁰ Given these psychological costs, the retailer's total psychological cost of overordering is $\delta_o(Q - m)^+$. Similarly, the retailer's total underordering psychological cost is $\delta_u(m - Q)^+$.

Under Assumption 2, the expected utility of the retailer involving reference dependence can be thus written as

$$\begin{aligned}
 U(Q) &= \int_0^Q [p(1 - \alpha + \alpha G(v_{m,\infty}^*))m + s(1 - G(v_{m,\infty}^*))\alpha m - \delta_o(Q - m)] dF(m) \\
 &+ \int_Q^\infty [pQ - \delta_u(m - Q)] dF(m) - cQ.
 \end{aligned} \tag{1}$$

Note that when all customers are myopic ($\alpha = 0$), our model in this section reduces to the single-store newsvendor version of Ho et al.'s (2010) model (see Section 2 therein).

3 Model Analysis

In this section, we provide the theoretical results of our behavioral model. We characterize the optimal order quantity of the retailer under reference dependence and study its properties. We also show the implication of our model on the traditional pull-to-center effect—the decision bias known to prevail in the standard newsvendor setting. In particular, we show how the ordering pattern and biases of the retailer are leveraged by the presence of strategic

¹⁰Since the retailer faces customers from separate classes (myopic and strategic), we also consider a different model in which the retailer manages separate accounts for myopic and strategic customers. Mathematically, it means the psychological costs associated with myopic and strategic customers are differentiated. Let δ_o^m (δ_o^s) be the psychological cost associated with each overordered unit with respect to myopic (strategic) customers; and δ_u^m (δ_u^s) the psychological cost associated with each underordered unit with respect to myopic (strategic) customers. Such a model, however, does not receive sufficient support from the laboratory experiments.

customers in the market. We start with the analysis of the model from Section 2 and then we proceed with analyses of extended models.

3.1 Analysis of the Main Model

We start with characterization of the optimal behavioral order quantity.

Lemma 3. *The optimal behavioral order quantity that maximizes the retailer's reference-dependent utility, denoted by Q_r^* , solves the following equation:*

$$(p + \delta_u)(1 - F(Q_r^*)) - \delta_o F(Q_r^*) - \alpha Q_r^* f(Q_r^*)(p - s)(1 - G(v_{m,\infty}^*)) - c = 0. \quad (2)$$

The optimality condition (2) is established from the first-order condition of the retailer's objective function, of which the unimodality follows from the IGFR property (Assumption 1). For simplicity, we refer to Q_r^* as the *behavioral* order quantity or *reference dependent* order quantity. When the demand follows uniform distribution $U[\underline{M}, \bar{M}]$, the explicit solution of Q_r^* is given by

$$Q_r^* = \frac{(p + \delta_u - c)\bar{M} + (\delta_o + c)\underline{M}}{\delta_o + \delta_u + p + \alpha(p - s)(1 - G(v_{m,\infty}^*))}, \quad (3)$$

which will be used in our subsequent experiments. Notice that the optimal solution that maximizes the retailer's *monetary* profit, denoted by Q^* , can be obtained from (2) by setting $\delta_u = \delta_o = 0$. That is, Q^* solves

$$p - c - pF(Q^*) - (p - s)(1 - G(v_{m,\infty}^*))\alpha Q^* f(Q^*) = 0.$$

We also refer to Q^* as the *rational* order quantity, or optimal order quantity when there is no confusion. In particular, if we set $\alpha = 0$, Q^* then reduces to the solution to the standard newsvendor problem, i.e., $F(Q^*)|_{\alpha=0} = \frac{p-c}{p}$.

Proposition 1. *The behavioral order quantity, Q_r^* , decreases in the unit cost, c , and in the fraction of strategic customers, α .*

Proposition 1 demonstrates the comparative statics of the behavioral order quantity in regard to the key parameters of our model. Particularly, as there are more strategic customers in the population, the retailer's order shall decrease. By doing so, the retailer shows a lack of inventory so that strategic customers herd into purchases at full price. This property also implies that the retailer facing strategic customers orders less than does a conventional newsvendor, while the latter corresponds to $\alpha = 0$ in our case. Thus, our model nests as a special case the single newsvendor version of the reference dependence model presented in Ho et al. (2010). Proposition 1 also implies that the retailer's rational order quantity Q^* decreases in the production cost (c) and in the fraction of strategic customers (α). It also follows from Proposition 1, by an application of envelope theorem, that both the maximal profit (attained with Q^*) and the maximal utility (attained with Q_r^*) of the retailer decline in the proportion of strategic consumers (α).

To proceed, we develop the theoretical framework to address the potential pull-to-center effect. In the spirit of Proposition 1 in Ho et al. (2010), we have the following result.

Proposition 2. *A behavioral retailer orders less [more] than the rational optimum, if and only if*

$$\frac{\delta_u}{\delta_o} < [>] \frac{F(Q^*)}{1 - F(Q^*)}. \quad (4)$$

Note that a low [high] unit cost leads to a larger [smaller] rational order quantity Q^* (by Proposition 1), which increases [reduces] the value of RHS in (4). Thus, the behavioral retailer tends to order below optimum when the cost is low, and above the optimum when the cost is high. In other words, the behavioral retailer may exhibit the well-established pull-to-center bias, yet in our two-period setting with strategic customers. Furthermore, notice that since the RHS of (4) decreases in α (as Q^* decreases in α by Proposition 1), there exists some threshold of α below [above] which the retail orders less [more] than the rational order (although both the behavioral order and the rational order decrease in α). This argument is further explored under different cost conditions in Proposition 5, and validated experimentally in Section 4.2.

Define a threshold unit cost, $\hat{c} := p(1 - F(\mu)) - \alpha\mu f(\mu)(p - s)(1 - G(v_{m,\infty}^*))$, such that the rational order quantity equals the average demand when $c = \hat{c}$. We say that the unit cost is *low* [*high*] if it is less [more] than \hat{c} , and the resulting rational order quantity is greater [less] than the average demand μ . The following proposition establishes the comparison between reference dependent order quantity and the mean demand.

Proposition 3. *A behavioral retailer orders less [more] than the average demand, if and only if*

$$\delta_u(1 - F(\mu)) - \delta_o F(\mu) < [>]c - \hat{c}. \quad (5)$$

While the traditional pull-to-center effect suggests that the order quantity falls between optimum and mean demand, Proposition 3 implies that it may exceed these bounds for two reasons. First, observe that in (5), only if the individual's average psychological cost is adequately high [low] relative to that of the underage, it will be possible for her to order less [more] than μ in a low [high] cost condition. In particular, Proposition 3 implies a *pull-below-center* effect in the low cost setting, since the psychological cost for overstocking tends to significantly outweigh that for understocking (given the saliency of leftover inventory—see Ho et al. (2010)) Second, noting that \hat{c} declines in α , an increase in the proportion of strategic customers results with an increase in the RHS of (5), which implies that the reference dependent order is more likely to fall short of the average demand, lending further support to the notion of pull-below-center effect.¹¹ Since we seldom observe pull-below-center effect in standard newsvendor experiments, it is plausible that the asymmetry of psychological costs alone is *not* sufficient in driving the retailer orders below the average demand. It is the introduction of strategic customers that further pulls the order downward below the mean demand. In other words, the emergence of pull-below-center effect in our experiment, but not in standard newsvendor experiments, should be attributed to the presence of strategic customers.

¹¹This argument can also be reached by simply observing that Q_r^* decreases in α (by Proposition 1).

We refer to the combined effects of pull-to-center and pull-below-center as the *generalized pull-to-center* bias. Propositions 2 and 3 provide a necessary and sufficient characterization (including the pattern of deviations from rationality and from the mean demand) of this generalized effect. Next, we shall investigate factors that influence the extent of the generalized pull-to-center bias.

Let $Q_{c\alpha}^*$ and $Q_{rc\alpha}^*$ denote the rational and behavioral order quantities, respectively, for given c and α . Consider two different levels of cost, \bar{c} and \underline{c} , such that $\bar{c} > \underline{c}$. A pair of costs $\{\bar{c}, \underline{c}\}$ is *symmetric* if $F(Q_{\bar{c}\alpha}^*) = 1 - F(Q_{\underline{c}\alpha}^*)$. Under a symmetric distribution F , symmetric cost conditions imply that the rational orders $Q_{\underline{c}\alpha}^*$ and $Q_{\bar{c}\alpha}^*$ are symmetrically positioned relative to the mean demand μ . This setting is often employed in the design of newsvendor experiments, where these conditions can be explicitly written in the form of critical ratios (see, e.g., Ho et al. 2010). Similar to Ho et al. (2010), let $\Omega_{\underline{c}\alpha} := Q_{\underline{c}\alpha}^* - Q_{rc\underline{c}\alpha}^*$ and $\Omega_{\bar{c}\alpha} := Q_{rc\bar{c}\alpha}^* - Q_{\bar{c}\alpha}^*$ to measure the degree of the generalized pull-to-center effect under cost \underline{c} and \bar{c} , respectively.¹² The ensuing Propositions 4 and 5 establish the effects of c and α on the generalized pull-to-center bias.¹³ The following result generalizes Proposition 2 of Ho et al. (2010) into an environment with strategic customers and IGFR demand distributions.

Proposition 4. *The generalized pull-to-center effect is stronger in the lower cost condition relative to that of the higher cost condition, i.e., $\Omega_{\underline{c}\alpha} > \Omega_{\bar{c}\alpha}$ if*

$$\frac{\delta_u}{\delta_o} < \frac{1}{2} \left(\frac{F(Q_{\bar{c}\alpha}^*)}{1 - F(Q_{\bar{c}\alpha}^*)} + \frac{F(Q_{\underline{c}\alpha}^*)}{1 - F(Q_{\underline{c}\alpha}^*)} \right). \quad (6)$$

Furthermore, if the cost conditions are symmetric, $\Omega_{\underline{c}\alpha} > \Omega_{\bar{c}\alpha}$ holds if

$$\frac{\delta_u}{\delta_o} < 1. \quad (7)$$

The generalized pull-to-center effect conjectures that the retailer orders less [more] than optimum under low [high] cost. By Proposition 2, this means the ratio $\frac{\delta_u}{\delta_o}$ falls within the

¹²Since we do not impose the mean demand as a limit for retailer order quantity in generalized pull-to-center scenario, our measure of the retailer's bias ($\Omega_{c\alpha}$) is defined upon the deviation from the rationality.

¹³To prove Propositions 4 and 5 we need to introduce a technical assumption relating to the reference-dependent utility (see Assumption 1.1 in the Appendix).

interval $\left(\frac{F(Q_{\underline{c}\alpha}^*)}{1-F(Q_{\underline{c}\alpha}^*)}, \frac{F(Q_{\bar{c}\alpha}^*)}{1-F(Q_{\bar{c}\alpha}^*)}\right)$. Moreover, inequality (6) implies that the order quantity bends towards the lower end of the interval, i.e. $\frac{F(Q_{\underline{c}\alpha}^*)}{1-F(Q_{\underline{c}\alpha}^*)}$. In this case, a larger departure from optimum is incurred under \underline{c} than under \bar{c} .¹⁴ Consequently, a generalized pull-to-center effect being stronger with lower production cost, can be interpreted as an outcome of psychological cost for excess inventory adequately exceeding that for unsatisfied demand.¹⁵ When the demand follows uniform distribution between \underline{M} and \bar{M} , we can obtain the necessary and sufficient condition for $\Omega_{\underline{c}\alpha} > \Omega_{\bar{c}\alpha}$ as

$$\frac{2p\delta_o - (\bar{c} + \underline{c})(\delta_u + \delta_o)}{\delta_u\bar{M} + \delta_o\underline{M}} > \frac{2\alpha(p-s)(1-G(v_{m\infty}^*))}{\bar{M} - \underline{M}}. \quad (8)$$

The fundamental effect of α on the generalized pull-to-center bias can be observed from (6). Since the RHS of (6) decreases in α , the condition (6) is less likely to be satisfied when α increases. In other words, with more strategic customers, it is less likely that $\Omega_{\underline{c}\alpha} > \Omega_{\bar{c}\alpha}$. More specifically for uniform distribution, the condition in (8) is less likely to be met as α becomes larger. The impact of α on the generalized pull-to-center effect is formalized in the following proposition.

Proposition 5. *Under the low cost condition, the generalized pull-to-center effect decreases with the proportion of strategic customers, α . For the high cost condition, we show a necessary condition always holds true for the generalized pull-to-center effect to increase in α .*

Proposition 5 predicts opposite effects of strategic customers on the retailer's ordering bias under different cost conditions. Although both the reference dependent order and the rational order decrease in α , the gap between the two may grow or shrink depending on the level of cost. Because $Q_{\underline{c}\alpha}^*$ decreases faster than $Q_{\bar{c}\alpha}^*$, their gap diminishes in α under low cost, and expands when the cost is high. Stated alternatively, the presence of strategic

¹⁴Under symmetric production costs, we have $\frac{F(Q_{\underline{c}\alpha}^*)}{1-F(Q_{\underline{c}\alpha}^*)} \times \frac{F(Q_{\bar{c}\alpha}^*)}{1-F(Q_{\bar{c}\alpha}^*)} = 1$; and thus we use the inequality $\frac{1}{2} \left(\frac{F(Q_{\underline{c}\alpha}^*)}{1-F(Q_{\underline{c}\alpha}^*)} + \frac{F(Q_{\bar{c}\alpha}^*)}{1-F(Q_{\bar{c}\alpha}^*)} \right) \geq \sqrt{\frac{F(Q_{\underline{c}\alpha}^*)}{1-F(Q_{\underline{c}\alpha}^*)} \times \frac{F(Q_{\bar{c}\alpha}^*)}{1-F(Q_{\bar{c}\alpha}^*)}} = 1$ to simplify the sufficient condition into $\frac{\delta_u}{\delta_o} < 1$.

¹⁵Noticeably, Proposition 4 is compatible with the existing knowledge that the pull-to-center effect is stronger with higher product profitability (see Schweitzer and Cachon 2000 and Ho et al. 2010, among others), and generalizes it in the context of strategic customers under IGFR demand.

customers *offsets* the retailer's ordering bias upon a low unit cost, but *aggravates* it when the cost is high. In the low [high] cost case, we show a sufficient [necessary] condition holds for the respective result, given as (D.31) [(D.34)] in the proof. As we will later see, the conjectures of Proposition 5 receive support from the experiment.

3.2 Extended Models

Our model adopts two important assumptions that facilitate a simplified exposition of the analysis and corresponding experimental tests. Specifically, we assume (i) complete information in the sense that customers know the retailer's available inventory as well as the total demand in the market prior to making their wait-or-buy decisions, and that (ii) customers coordinate on the all-buy equilibrium when it is available (Assumption 2). Both assumptions can be relaxed, while some important insights of the model are preserved. While herein we offer an overview of the results, the complete analysis of the extended models is delegated to the Online Supplement.

In the absence of information regarding inventory and demand, the setting can be conceptually perceived as a simultaneous game between customers and the retailer. Solving the resulting game in the same fashion as in Cachon and Swinney (2009), we find that most of our results are robust to incomplete information. If the customers play the threshold equilibrium instead of the all-buy equilibrium (see Lemma 2), we show that the key messages of the model sustain, under some parametric restriction and uniform sales period valuation. Those assumptions are made for analytical necessities, noting that the uniform valuation assumption is widely used in the revenue management literature (including Cachon and Swinney 2009; Desai 2007; Lazear 1986, among others). Along with the main results on retailer behavior, we also provide a characterization of the customer equilibrium and its properties.

For both model relaxations (i.e., (i) incomplete information and (ii) altered customer equilibrium), we prove that Proposition 2 and 3 hold. Analogous to Proposition 4, the generalized pull-to-center effect is stronger in the lower cost condition relative to that of the higher cost condition, if $\frac{\delta_u}{\delta_o} < \frac{F(Q_{\bar{c}\alpha}^*) + F(Q_{c\alpha}^*)}{(1 - F(Q_{\bar{c}\alpha}^*)) + (1 - F(Q_{c\alpha}^*))}$ (which reduces to $\frac{\delta_u}{\delta_o} < 1$ if the

cost conditions are symmetric). All in all, we demonstrate the robustness of the patterns in which reference dependence affects the retailer’s ordering behavior in the presence of strategic customers.

4 Validation Experiments

4.1 Experimental Design

In this section we design a simple experiment to test our behavioral model. In order to maintain focus and frugality, we base our experiment on the main model in Section 3.1, which drafts on complete information and simplified customer behavior (Assumption 2). We employ a factorial design where the fraction of strategic customers (α) and the unit cost (c) are the two main factors under examination. The design factor α has 3 levels, $\alpha \in \{0, 0.5, 1\}$, and c has 2 levels, $c \in \{3, 6\}$, resulting in 6 treatments in total. The choice of high cost, $c = 6$, and low cost, $c = 3$, enables us to test the pull-to-center effect from both below and above the mean demand. Note that the treatments with $\alpha = 0$ (i.e., no strategic customers) coincide with the conventional newsvendor setting. The experimental design is summarized in Table 1. The selling and sales period prices are fixed at $p = 12$ and $s = 6$, respectively, for all treatments.¹⁶ Demand follows a discrete uniform distribution between 1 and 300, i.e. $M \sim U[1, 300]$. Note that a uniform distribution exhibits IGFR, thereby satisfying Assumption 1. A treatment is labeled by a combination of a letter indicating whether it is a low-cost, L, or high-cost, H, treatment, and a number indicating the value of α —i.e., 0, 0.5, or 1. The experimental design is summarized in the table below, with the rational order quantity of each treatment provided in parentheses.

In each round of experiments, each subject acts as a retailer, makes a single decision (order quantity), and plays against automated customers who behave rationally according to Lemma 2 and Assumption 2. Automating the customers allows us to focus on the retailer’s

¹⁶An initial markdown of 50% is consistent with practice. For example, in Sports Unlimited this was the preferred markdown so as “to really attract customer attention and move merchandise” (Talebian and van Ryzin, 2014).

Table 1: Experimental design

		Treatment (rational order quantity)		
		Fraction of strategic customers, α		
		0	0.5	1
Cost, c	3	L0 (225.25)	L0.5 (180.2)	L1 (150.17)
	6	H0 (150.5)	H0.5 (120.4)	H1 (100.33)

$p = 12, s = 6, M \sim U[1, 300]$, between-subjects design

behavior, and is common in the literature (see, e.g., Bolton and Katok (2008); Kremer et al. (2017)). We choose not to navigate the subjects through the detailed problem faced by strategic customers, and instead directly inform them about the outcome of customer purchases as a function of the subject's order quantity and the realized demand. We also set up the experiment such that the customer equilibrium is reduced to all-wait when $Q \geq m$ (e.g., by assuming a common v that exceeds $v_{m,\infty}^*$) and all-buy when $Q < m$. As a result, one can obtain the rational order for each treatment (listed in Table 1) by setting $\delta_o = \delta_u = 0$ and $G(v_{m,\infty}^*) = 0$ in (3), i.e. $Q^* = \frac{(p-c)\bar{M}+cM}{p+\alpha(p-s)}$.

To further ensure the subjects understand the game, they were asked to complete a quiz that covers critical aspects of the game. The experiment began only after the quiz questions were answered correctly. Before the actual experiment, the subjects were given three practice rounds to help them become familiarized with the game. To facilitate the subjects' decision making, we provided them with a decision support tool that calculates the profit for any combination of an order quantity and a demand realization set by the subject. Subjects played 50 rounds of the game, where each round was an independent repetition of the game with no inventory carried over between rounds. Sample instruction can be found in Appendix C.

The experiment was conducted in a behavioral lab in a public university. In total, we recruited 215 subjects who were undergraduate students. Each treatment has more than 30 subjects and no subject participated in more than one treatment. The sample sizes are

reported in Table 2. Subjects were paid a fixed show-up fee and an additional performance-based payment that is proportional to the total profit over the 50 rounds of the experiment. We used different conversion rates from profit to payment in different treatments in order to maintain similar average amount of cash payments. On average, subjects earned a total amount equivalent to \$15.

4.2 Preliminary Results

The summary statistics are provided in Table 2, where we also report the statistical tests of the observed orders against the optimal order quantities predicted by the theory. In the reported t -tests, the unit of analysis is the average order quantity of each subject over the course of 50 rounds. The results indicate that subjects fail to place optimal orders in all treatments except for treatment H0.¹⁷

Table 2: Summary statistics

Treatment	Mean	Std. dev.	Optimum	p -value*	# subjects	Mean - optimum
H0	146.70	35.17	150.5	0.528	35	3.8
H0.5	128.85	20.71	120.4	0.028	32	8.45
H1	114.43	28.44	100.33	0.006	35	14.1
L0	168.56	38.93	225.25	0.000	36	56.69
L0.5	138.43	22.37	180.2	0.000	37	41.77
L1	136.14	29.61	150.17	0.005	40	14.03

* Two-tailed t -test; optimal order quantity as the null hypothesis.

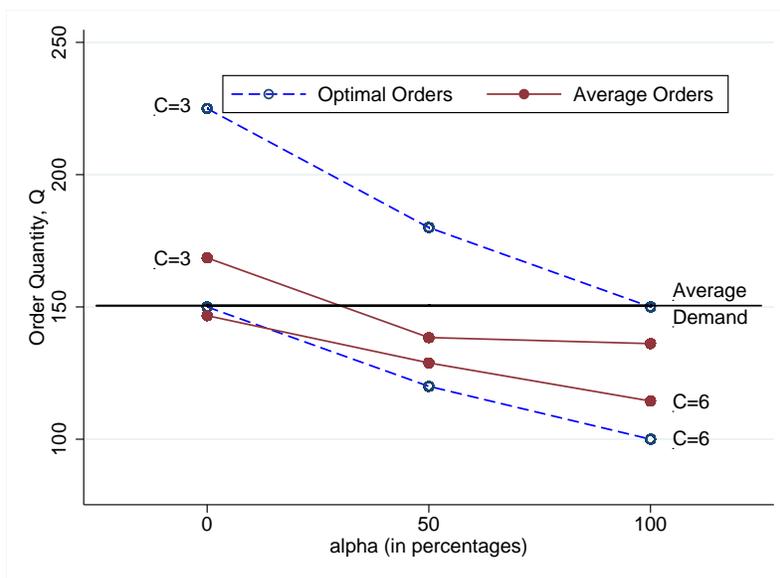
4.2.1 Comparative statics

Figure 2 below depicts the rational orders (dashed lines), the actual orders placed by subjects (solid lines) in each treatment, and the mean demand (solid horizontal line). Clearly, the subjects order less as the unit cost increases and that they order less as the proportion of

¹⁷A plausible explanation for this exception is that at H0 the optimal order quantity is equal to the mean demand—a relatively easy benchmark for subjects to converge on. However in L1, subjects' orders tend to pull below the average demand, even when ordering the average demand is nearly optimal. This is explained later in Section 4.2.2.

strategic customers increases, as predicted by Proposition 1. A two-way analysis of variance (ANOVA) test confirms the visual message ($p < 0.001$ for the main effects of both c and α based on average orders by subjects). Although Figure 2 implies that the change in order quantity with respect to one factor may depend on the level of the other factor, this interaction effect is not statistically significant in the ANOVA test.

Figure 2: Decision biases in our experiment: pull-to-center and pull-below-center



4.2.2 Generalized pull-to-center effect

Recall that the generalized pull-to-center bias predicted by our model entails both the traditional pull-to-center effect and the novel pull-below-center phenomenon. Consistent with the vast literature on behavioral newsvendor, in the absence of strategic consumers, we find only evidence for the traditional pull-to-center effect. By Figure 2, the actual average order quantity in treatment L0 clearly lies between the mean demand (150.5) and the rational order (225.25); whereas in treatment H0, the average order sticks very close to mean demand, as the latter coincides with the optimal order in this case (see the t -test in Table 2). When strategic consumers are introduced, the average orders in all four treatments fall below the mean demand, 150.5. In the high cost treatments (H0.5 and H1), this does not go against the pull-to-center conjecture, because the quantities still exceed their respective optimum

(120.4 for H0.5 and 100.33 for H1). In the low cost treatments where the optimal orders are above or close to the mean demand (180.2 for L0.5 and 150.17 for L1), however, the observed patterns break the conventional pull-to-center framework, because the actual orders do not fall between the optimum and the average demand. One-sided t -tests based on the average order quantities by subject show that the order quantity is significantly lower than the mean demand in both treatments ($p = 0.001$ for L0.5 and $p = 0.002$ for L1). This confirms the pull-below-center bias predicted by our model; and it is clearly the introduction of strategic customers that gives an additional downward pull of the order quantity, giving rise to the pull-below-center effect.¹⁸ While the discussion of pull-below-center effect in this section is held at the aggregate level, Appendix B shows that similar patterns also prevail at the individual level.

Our game differs from the traditional newsvendor problem in the introduction of strategic customers. Along this line, we also discover interesting insights regarding how the strategicness of the customer population influences the retailer's decision bias. Our Proposition 5 predicts that, under high [low] production cost, the generalized pull-to-center bias will be strengthened [alleviated] as the proportion of strategic customers in the population rises. This prediction is validated by the empirical results in Table 2, where the departure of the actual order quantity from the optimum increases in α when the cost is high ($3.8 \nearrow 8.45 \nearrow 14.1$), and declines in α in the low cost condition ($56.69 \searrow 41.77 \searrow 14.03$). Also observe that, while both the rational and behavioral order quantities decrease with the ratio of strategic consumers, the former falls faster than the latter (Figure 2). In other words, the actual orders are *less sensitive* than the rational orders to the increase of strategic customers in the population.

Another salient observation lies in the differences in the level of generalized pull-to-center bias across treatments. When the cost is low, the subjects' order quantities drift further away

¹⁸In addition to the experiment we report here, we have played the low-cost treatments of our experiment in four different operations management and inventory management classes. While a different incentive is adopted (bonus point for best performer), subjects in those treatments consistently exhibit the pull-below-center bias.

from the theoretical optimum. The last column of Table 2 reveals that the deviations from the optimal quantity are larger in the low cost treatments than in the high cost treatments. This validates the prediction of Proposition 4.

5 Model Estimation and Discussion

In this section, we estimate the subject's psychological cost parameters for overstocking and understocking using our experimental data. Since the demand in our experiment is uniformly distributed between $[M, \bar{M}]$ and strategic customers all choose to purchase in the selling season when $Q < m$, and they wait for the sales season when $Q \geq m$, we can express the retailer's expected utility as follows:

$$\begin{aligned}
 U(Q) &= \int_M^Q [p(1 - \alpha)m + s\alpha m - \delta_o(Q - m)] \cdot \frac{1}{\bar{M} - M} dm \\
 &+ \int_Q^{\bar{M}} [pQ - \delta_u(m - Q)] \cdot \frac{1}{\bar{M} - M} dm - cQ.
 \end{aligned}$$

Thus, the reference-dependent order quantity maximizing $U(Q)$ takes the explicit form as follows.

$$Q_r^* = \frac{(p + \delta_u - c)\bar{M} + (\delta_o + c)M}{\delta_o + \delta_u + p + \alpha(p - s)}. \quad (9)$$

Remark 1 below characterizes the necessary and sufficient condition for pull-below-center effect specific to uniform demand distribution, which will be useful for interpreting our experimental estimations later. The proof of Remark 1 is implied from Proposition 3 and hence omitted.

Remark 1. *Under a uniform demand distribution $U[M, \bar{M}]$, the necessary and sufficient condition for $Q_r^* < \mu$ is $\delta_o - \delta_u > p - 2c - (p - s)\alpha \frac{2\mu}{M - \bar{M}}$.*

For structural estimation of the psychological costs, we assume that the order quantities are normally distributed with a mean Q_r^* and a treatment-specific standard deviation σ_i , where i is the treatment indicator from our experiment, $i \in \{L0, L0.5, L1, H0, H0.5,$

H1}. Using the maximum likelihood estimation method, we estimate the full model with all parameters, as well as two nested models. One of them is a rational model where the subjects make rational decisions, the other is a restricted model where the psychological costs associated with overage and underage are assumed to be identical, i.e. $\delta_o = \delta_u$.

Table 3: Estimation results

Parameters	Rational model	Restricted model	Full model
	$\delta_o = \delta_u = 0$	$\delta_o = \delta_u$	
δ_u	-	11.64	10.53
	-	(3.03)	(2.65)
δ_o	-	11.64	13.96
	-	(3.03)	(3.21)
σ_{L0}	79.21	55.80	55.83
	(4.83)	(4.64)	(4.89)
$\sigma_{L0.5}$	58.99	47.93	42.92
	(3.23)	(3.02)	(2.54)
σ_{L1}	48.95	49.00	46.95
	(3.28)	(3.28)	(3.71)
σ_{H0}	51.88	51.88	52.76
	(3.19)	(3.19)	(3.40)
$\sigma_{H0.5}$	41.61	41.92	40.84
	(2.57)	(2.71)	(2.56)
σ_{H1}	46.55	46.58	44.45
	(2.57)	(2.70)	(2.30)
LL	-58045.1	-57045.7	-56661.8
Wald test and p-value	$\chi^2(2) = 21.93$	$\chi^2(1) = 20.12$	
(test against the full model)	0.000	0.000	

Notes: All estimates are significant at the 0.01 level.

Standard errors are clustered by subject and reported in parentheses.

The estimation results are reported in Table 3. Since the observations are not independent at the subject level, we resort to the Wald test (Wooldridge, 2010) in order to compare the rational and restricted models with respect to the full model. The tests show that

the full model performs better than the former two. Using the estimated parameters, we present the predictions of various models in Table 4. Generally speaking, the predictions of behavioral models (restricted and full) exhibit less sensitivity in α than does the rational model (consistent with the trend of the data — see Section 4.2.2). As expected, the full model yields predictions that are closer to the actual observations than does the restricted model. More importantly, compared to the restricted model, the full model predicts a further downward shift of the order quantities which is consistent with the concept of pull-below-the-center. Also, the tests in Table 3 suggest that subjects incur remarkably greater psychological costs associated with overstocking than with understocking. This can stem from the saliency of leftover inventory relative to unmet demand (Ho et al., 2010), and in that sense, generalizes the classical finding of Ho et al. (2010) into an environment with strategic consumers.

Table 4: The predictions of models using the estimated parameters

		Fraction of strategic customers, α		
Cost		0	0.5	1
Rational	3	L0 (225.25)	L0.5 (180.2)	L1 (150.17)
	6	H0 (150.5)	H0.5 (120.4)	H1 (100.33)
Restricted	3	L0 (175.93)	L0.5 (162.14)	L1 (150.35)
	6	H0 (150.5)	H0.5 (138.71)	H1 (128.61)
Full	3	L0 (161.03)	L0.5 (148.8)	L1 (138.29)
	6	H0 (136.45)	H0.5 (126.08)	H1 (117.18)

$$p = 12, s = 6, M \sim U[1, 300]$$

Proposition 2 characterizes the conditions under which the behavioral order is higher or lower than the rational order quantity. Using the rational orders reported in Table 1 and the psychological parameters estimated in Table 3, it is easy to confirm that the relationships between the parameters, summarized in (4), reflect consistently the aggregate patterns (i.e.

whether the actual order falls below or above the optimum) observed for all treatments except H0. This is not surprising since the optimal order quantity in treatment H0 equals to the mean demand, thus leaving no space for the pull-to-center effect to stretch out.

Next, we turn to cases where the orders are pulled below the mean demand. The necessary and sufficient condition of the pull-below-center effect under uniform demand distribution is given in Remark 1, and it is shown to hold for all treatments except for L0 (which is not surprising as subjects do not pull below the center when the cost is low in the standard newsvendor setting, i.e. $\alpha = 0$). Hence, our behavioral model featuring reference dependence has successfully explained the rationale behind the pull-below-center bias in the treatments with strategic customers.

Last, we address the observed difference in the generalized pull-to-center effect over different cost treatments. In the context of a standard newsvendor game, this could stem from the fact that human decision makers are more sensitive to the “visible” waste generated by the disposed products, relative to the opportunity cost from lost demand (Ho et al. 2010). In our game, the leftover inventory is also responsible for driving strategic customers into the sales season, incurring new loss of profit in addition to the conventional cost of overage. Therefore, the issue where the mental cost for overordering is larger than that of underordering in newsvendor environment aggravates with the presence of strategic customers. Indeed, our estimation shows both $\delta_o > \delta_u$ (see Table 3) and that (6) holds with the estimated values for all α . As a result of Proposition 4, subjects’ orders deviate more from the theoretical optimum when the unit cost is low.

6 Robustness and extensions

In this section we present additional robustness and extensions carried out and which are elaborated in the appendix.

6.1 Learning, heuristics, and individual-level behavior

Our behavioral model abstracts away from the effects of learning. As detailed in Appendix A, the study of subject learning does not qualitatively change the insights shed by our existing results. Although subjects exhibit some mild learning in the low-cost treatments, the analysis of subject behavior using the latter half of experimental data does not show significantly different insights than using the whole set of data. We then model subjects' decision dynamics via the mean anchor heuristic (Schweitzer and Cachon, 2000). Our analysis indicates that accounting for learning effects does bring the anchor closer to the mean; and the pull-below-center bias persists in the presence of strategic customers.

Another facet of analysis we left out of the main text is individual-level behavior. In Appendix B we show that, while subjects exhibit considerable heterogeneity, their individual behavior echoes the aggregate patterns of generalized pull-to-center effect found in each of the treatments.

6.2 Subjects as customers

In order to focus on the retailer behavior, we automated the purchase decisions made by the strategic customers in our main experiments. In this section, we check the robustness of our findings when retailers interact with real human customers. The experimental design and the associated model are provided in the online supplement 2.4. In brief, our experiment simplifies the continuum space of customers and controls the social preference of interacting subjects, while preserving the core tradeoff in the customer's problem between buy and wait.

The game is divided into two phases, with 25 periods in each phase. In each phase, each subject playing the retailer is matched in each round with a single subject who acts on behalf of all strategic customers and has valuation v for the product in the sales period;¹⁹

¹⁹To maintain the frugality of experimental design and keep consistency with the main experiment, strategic customers have homogeneous valuations in this experiment. Therefore, all strategic customers will behave the same in equilibrium. That allows us to set up the experiment so that their decision makings are all delegated to a single human subject (called the focal subject). The focal subject chooses between buy now and wait—which is the core trade-off that we seek to capture in the experiment. In making this decision, the focal customer has to compete with the automated myopic customers as well as all other strategic customers who replicate his/her decision.

and we rematch them every round so that no partners meet more than once in each phase (to avoid reputation building). The game is identical in the two phases, except that subjects exchange their roles across phases (that is, subjects who play retailers [customers] in phase I play customers [retailers] in phase II). Since each subject experiences both roles (retailer and customer) for an equal number of rounds, the payoff received by each subject is ex-ante similar to his or her counterparts and issues of fairness are therefore unlikely to arise. A similar role-flip approach is adopted in Özer et al. (2017). In this design, we still automate the role of the myopic customers.²⁰

For consistency with the main experiment, we maintain the parameters for the retailer's problem (e.g. $p = 12, s = 6$). For the customer game, we choose $v = 12, v_M = 16, c = 3, \alpha = 0.5$ and 0.35 . The resulting treatments are respectively coded as L0.5h and L0.35h (where h is abbreviated for human-customer). The choices of α reflect the consistency with the main experiment ($\alpha = 0.5$) as well as a relatively low ratio ($\alpha = 0.35$) to accommodate scenarios where strategic behavior is mild for customers.

The experiment was conducted in a behavioral lab in a public university. In total, we recruited 100 subjects who were undergraduate students, 50 for each of the treatments. Subjects, upon playing the role of the strategic customers, exhibit a highly rational behavior. Specifically, in treatment L0.5h [L0.35h] the average frequency of making optimal buy-or-wait decision is 93.28% [96%], with a standard deviation of 7.89% [5.71%]. Given the high rationality of customer decision and our focus on retailer behavior, we shall base our further analysis assuming customers act optimally, which gives the optimal order quantity for the retailer as $Q^* = \frac{c(-\bar{M} + \underline{M}) + \bar{M}p}{\beta(2 - (1 - \alpha)\beta)(p - s) + s}$, where $\beta = \frac{v - s - \alpha(v_M - p)}{(1 - \alpha)(v - s)}$, and the reference-dependent order

$$Q^h = \frac{c(-\bar{M} + \underline{M}) + \underline{M}\delta_o + \bar{M}(p + \delta_u)}{\delta_o + \beta(2 - (1 - \alpha)\beta)(p - s) + s + \delta_u}. \quad (10)$$

The details of derivation is delegated to Online Supplement 2.4. Based on that, we can show the following comparative statics for the retailer's order.

²⁰Given the trivial decision of myopic customers and our focus on strategic customers, we automate the role of myopic customers.

Remark 2. *The reference-dependent order quantity of the retailer facing human customers, Q^h , decreases in the fraction of strategic customers in the population, α .*

Now we report the findings regarding the retailer behavior in our robustness check experiment. The summary statistics with respect to the ordering decisions are reported in Table 5. In the reported t -tests, the unit of analysis is the average order quantity of each subject when acting as retailer.

Table 5: Summary statistics

Treatment	Mean	Std. dev.	Optimum	p -value*	# subjects	Mean - optimum
L0.35h	155.27	18.40	183.53	0.000	50	28.26
L0.5h	146.55	24.89	162.18	0.000	50	15.63

* Two-tailed t -test; optimal order quantity as the null hypothesis.

The ordering behavior of retailer subjects (Table 5) confirms Remark 2 — The retailer cuts down his or her order quantity when there are more strategic customers. Although no direct prediction is available in theory, we also observe that the magnitude of retailer’s bias (measured by the deviation from optimum) is reduced with increasing strategic customers under the low production cost, which is aligned with the finding in the main experiment.²¹

7 Conclusions

In this paper, we study inventory decision making involving strategic customers from a behavioral perspective. Our model based on reference dependence makes predictions on how the retailer behavior is leveraged by the strategicness of customer population. Specifically, the reference-dependent retailer shall decrease her inventory level as the proportion of

²¹We find a pull-below-center effect in L0.5h, although it is not statistically significant (p -value = 0.1336 from one-sided t -test based on the average order quantities by subject as retailer). This finding is also confirmed by checking a necessary and sufficient condition for pull-below-center effect that applies to the robustness check experiment (see Remark 3 in Appendix 2.4). In treatment L0.35h, the pull-below-center effect is not observed at the aggregate level, as the treatment mean is already greater than the mean demand. Since L0.5h and L0.35h differ only by α , the observed difference confirms our earlier finding that a sufficiently high portion of strategic customers (α) is necessary to the emergence of pull-below-center behavior (p. 15-16, Section 3.1).

strategic customers rises in the population (in order to discourage strategic customers from waiting for markdowns). As such, the conventional pull-to-center bias is generalized with the presence of strategic consumers, since the retailer may now pull her order below the center when the cost is low. Furthermore, the generalized pull-to-center effect varies in opposite ways with the proportion of strategic customers under different cost conditions. When the unit cost is low [high], the generalized pull-to-center bias shrinks [grows] as the fraction of strategic customers increases in the population. Our subsequent experiment validates these predictions, calibrates the (asymmetric) reference dependence of the decision makers by the estimate of corresponding psychological costs. As for robustness checks, we conduct a series of treatments where strategic customers are also played by human subjects, and find qualitative support to our theory.

Our research contributes to the literature on behavioral research related to inventory management by accounting for the important dimension of strategic customer behavior. As a first attempt, we consider a predetermined pricing path in the problem setting. Future research may allow pricing to serve as a decision variable in addition to inventory. In the current study, we concentrate on structurally modeling the ordering behavior, and find relatively little impact of subject learning. A future study could be devoted to the investigation of time-dependent strategies or heuristics used in inventory decision setting against strategic customers, and how these are affected by operational environments such as demand dynamics.

Acknowledgments

The authors thank the Departmental Editor, Elena Katok, the Senior Editor and the two anonymous reviewers for valuable recommendations that have significantly improved the contribution and presentation of the paper. We thank seminar participants at the Chinese Academy of Sciences, Tsinghua University, University of Maryland, University of Luxembourg, POMS 2015, INFORMS 2015, MSOM 2015, BOM 2016, INFORMS 2016, INFORMS

2017, and the Retail Operations Conference 2018, for helpful comments and discussions. This work was supported by National Natural Science Foundation of China [grant number 71501108] and Beijing Natural Science Foundation [grant number 9164030]. Yang Zhang is grateful for the research assistance provided by Xiaomin Du, Yang Liu, and Xiaojin Fu. Yaozhong Wu gratefully acknowledges financial support by NUS HSS Seed Fund R-314-000-099-646. The experimental software was developed by Taizan Chan.

References

- Altug, M. S. and Aydinliyim, T. (2016). Counteracting strategic purchase deferrals: The impact of online retailers return policy decisions. *Manufacturing & Service Operations Management*, forthcoming.
- Aviv, Y. and Pazgal, A. (2008). Optimal pricing of seasonal products in the presence of forward-looking consumers. *Manufacturing & Service Operations Management*, 10(3):339–359.
- Banciu, M. and Mirchandani, P. (2013). Technical note—New results concerning probability distributions with increasing generalized failure rates. *Operations Research*, 61(4):925–931.
- Bearden, J. N., Murphy, R. O., and Rapoport, A. (2008). Decision biases in revenue management: Some behavioral evidence. *Manufacturing & Service Operations Management*, 10(4):625–636.
- Becker-Peth, M., Katok, E., and Thonemann, U. W. (2013). Designing buyback contracts for irrational but predictable newsvendors. *Management Science*, 59(8):1800–1816.
- Bendoly, E. (2011). Linking task conditions to physiology and judgment errors in rm systems. *Production and Operations Management*, 20(6):860–876.
- Bolton, G. and Katok, E. (2008). Learning-by-doing in the newsvendor problem: A laboratory investigation of the role of the experience and feedback. *Manufacturing & Service Operations Management*, 10(3):519–538.

- Bostian, A. A., Holt, C. A., and Smith, A. M. (2008). Newsvendor pull-to-center effect: Adaptive learning in a laboratory experiment. *Manufacturing & Service Operations Management*, 10(4):590–608.
- Cachon, G. P. and Swinney, R. (2009). Purchasing, pricing, and quick response in the presence of strategic consumers. *Management Science*, 55(3):497–511.
- Cason, T. N. and Mago, S. D. (2013). A laboratory study of duopoly price competition with patient buyers. *Economic Inquiry*, 51(2):1123–1141.
- Chen, L., Kök, A. G., and Tong, J. D. (2013). The effect of payment schemes on inventory decisions: The role of mental accounting. *Management Science*, 59(2):436–451.
- Chevalier, J. and Goolsbee, A. (2009). Are durable goods consumers forward-looking? evidence from college textbooks. *The Quarterly Journal of Economics*, 124(4):1853–1884.
- Davis, A. M. (2015). An experimental investigation of pull contracts in supply chains. *Production and Operations Management*, 24(2):325–340.
- Davis, A. M., Katok, E., and Santamaría, N. (2014). Push, pull, or both? a behavioral study of how the allocation of inventory risk affects channel efficiency. *Management Science*, 60(11):2666–2683.
- Desai, P. S. (2007). The role of production lead time and demand uncertainty in meeting durable goods. *Management Science*, 53(1):150–158.
- Fisher, M. and Raman, A. (2010). *The New Science of Retailing: How Analytics are Transforming the Supply Chain and Improving Performance*. Harvard Business Review Press, Boston, MA.
- Gallego, G., Phillips, R., and Şahin, Ö. (2008). Strategic management of distressed inventory. *Production and Operations Management*, 17(4):402–415.

- Hardie, B. G., Johnson, E. J., and Fader, P. S. (1993). Modeling loss aversion and reference dependence effects on brand choice. *Marketing science*, 12(4):378–394.
- Hariss, R., Perakis, G., Sinchaisri, W., and Zheng, Y. (2016). Markdown pricing with quality perception and consumer optimism: From experiment to theory. Working Paper.
- Ho, T.-H., Lim, N., and Cui, T. H. (2010). Reference dependence in multilocation newsvendor models: A structural analysis. *Management Science*, 56(11):1891–1910.
- Ho, T.-H. and Zhang, J. (2008). Designing pricing contracts for boundedly rational customers: Does the framing of the fixed fee matter? *Management Science*, 54(4):686700.
- Kahneman, D. and Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2):263–291.
- Katok, E. and Pavlov, V. (2013). Fairness in supply chain contracts: A laboratory study. *Journal of Operations Management*, 31(3):129–137.
- Katok, E. and Wu, D. Y. (2009). Contracting in supply chains: A laboratory investigation. *Management Science*, 55(12):1953–1968.
- Kim, S. and Dasu, S. (2013). An experimental study of posted prices. *Available at SSRN 2195171*.
- Kocabiyikoglu, A., Gogus, C. I., and Gonul, M. S. (2015). Revenue management vs. newsvendor decisions: Does behavioral response mirror normative equivalence? *Production and Operations Management*, 24(5):750–761.
- Kremer, M., Mantin, B., and Ovchinnikov, A. (2017). Dynamic pricing in the presence of myopic and strategic consumers: theory and experiment. *Production and Operations Management*, 26(1):116–133.
- Kremer, M., Minner, S., and Van Wassenhove, L. N. (2010). Do random errors explain newsvendor behavior? *Manufacturing & Service Operations Management*, 12(4):673–681.

- Lariviere, M. A. (2006). A note on probability distributions with increasing generalized failure rates. *Operations Research*, 54(3):602–604.
- Lazear, E. P. (1986). Retail pricing and clearance sales. *The American Economic Review*, 76(1):14–32.
- Li, J., Granados, N., and Netessine, S. (2014). Are consumers strategic? structural estimation from the air-travel industry. *Management Science*, 60(9):2114–2137.
- Li, M., Petruzzi, N. C., and Zhang, J. (2016). Overconfident competing newsvendors. *Management Science*, 63(8):2637–2646.
- Liu, Q. and van Ryzin, G. (2008). Strategic capacity rationing to induce early purchases. *Management Science*, 54:1115–1131.
- Long, X. and Nasiry, J. (2015). Prospect theory explains newsvendor: The role of reference points. *Management Science*, 61(12):3009–3012.
- Mak, V., Rapoport, A., Gisches, E. J., and Han, J. (2014). Purchasing scarce products under dynamic pricing: An experimental investigation. *Manufacturing & Service Operations Management*, 16(3):425–438.
- Mersereau, A. J. and Zhang, D. (2012). Markdown pricing with unknown fraction of strategic customers. *Manufacturing & Service Operations Management*, 14(3):355–370.
- Nair, H. (2007). Intertemporal price discrimination with forward-looking consumers: Application to the us market for console video-games. *Quantitative Marketing and Economics*, 5(3):239–292.
- Osadchiy, N. and Bendoly, E. (2015). Are consumers really strategic? implications from an experimental study. Working paper, Emory University, Atlanta, GA.
- Ovchinnikov, A., Moritz, B., and Quiroga, B. F. (2015). How to compete against a behavioral newsvendor. *Production and Operations Management*, 24(11):1783–1793.

- Özer, Ö., Subramanian, U., and Wang, Y. (2017). Information sharing, advice provision, or delegation: What leads to higher trust and trustworthiness? *Management Science*.
- Ozer, O. and Zheng, K. (2016). Markdown or everyday low price? the role of behavioral motives. *Management Science*, 62(2):326–346.
- Phillips, R. (2005). *Pricing and Revenue Optimization*. Stanford Business Books. Stanford University Press.
- Popescu, I. and Wu, Y. (2007). Dynamic pricing strategies with reference effects. *Operations Research*, 55(3):413–429.
- Ren, Y. and Croson, R. (2013). Overconfidence in newsvendor orders: An experimental study. *Management Science*, 59(11):2502–2517.
- Schweitzer, M. and Cachon, G. (2000). Decision bias in the newsvendor problem with a known demand distribution: Experimental evidence. *Management Science*, 46(3):404–420.
- Song, Y. and Zhao, X. (2016). Strategic customer behavior facing possible stockout: An experimental study. *International Journal of Production Economics*, 180:57–67.
- Song, Y. and Zhao, X. (2017). A newsvendor problem with boundedly rational strategic customers. *International Journal of Production Research*, 55(1):228–243.
- Su, X. (2007). Intertemporal pricing with strategic customer behavior. *Management Science*, 53(5):726–741.
- Su, X. (2008). Bounded rationality in newsvendor models. *Manufacturing & Service Operations Management*, 10(4):566–589.
- Talebian, M. and van Ryzin, G. (2014). Case article markdown management at sports unlimited. *INFORMS Transactions on Education*, 14(2):96–99.

Uppari, B. S. and Hasija, S. (2016). On the consistency between prospect theory and the newsvendor pull-to-center effect. Working Paper, INSEAD.

Wooldridge, J. M. (2010). *Econometric analysis of cross section and panel data*. MIT press.

Zhang, D. and Cooper, W. L. (2008). Managing clearance sales in the presence of strategic customers. *Production and Operations Management*, 17(4):416–431.

Appendices

Appendix A Learning and Heuristics

In formalizing our behavioral model in Section 2, we have abstracted away from the effects of learning. In this appendix, we demonstrate that learning does not qualitatively change the insights derived in the foregoing sections. Figure 3 plots the average actual orders placed by subjects in each round of each treatment, with the corresponding 95% confidence interval for the orders shown as the vertical bar, and the respective optimal order quantity represented by the dotted line. Figure 3 suggests some learning trends that are especially visible in the low-cost treatments. However, the summary statistics of the second half of the experiment (i.e., rounds 26–50), shown in Table 6, yield little different insights compared to those based on data of the whole experiment (see Table 2 in the main text): Subjects' orders in both Tables 2 and 6 are significantly different from optimum except in Treatment H0 (with weak significance in L1 in Table 6).

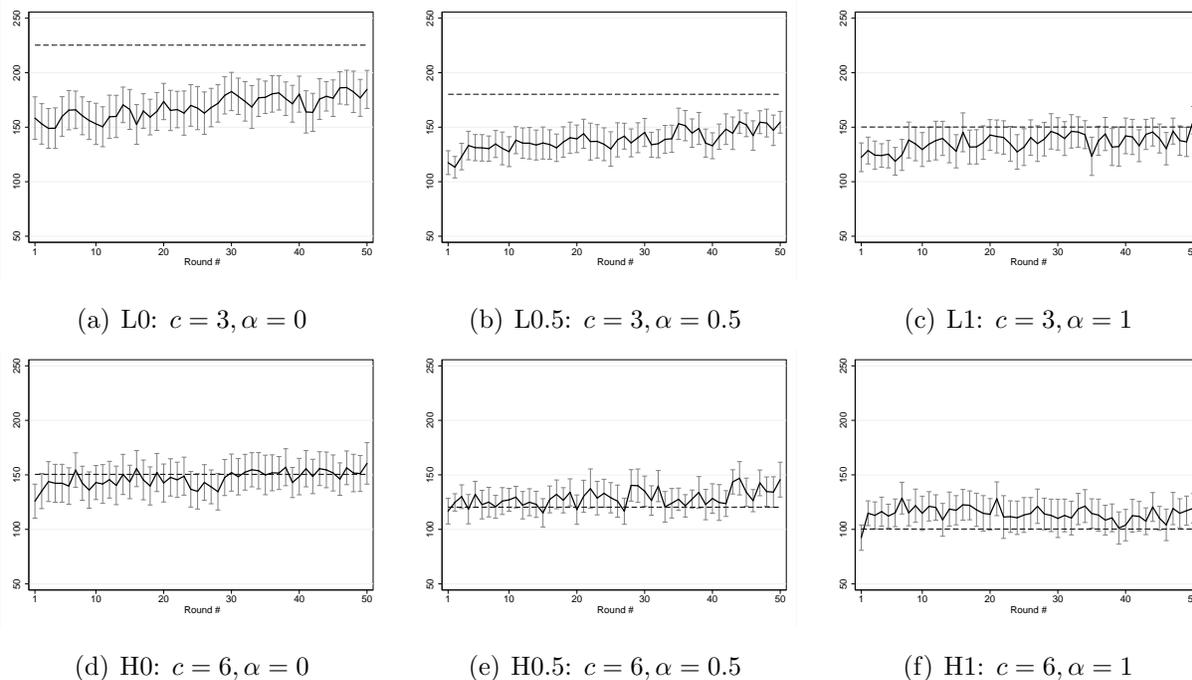


Figure 3: Orders by treatment (— Observed ---- Rational)

Table 6: Summary statistics of the second half of the experiment (rounds 26-50)

Treatment	Mean	Std. dev.	Optimum	p -value*	# subjects
H0	149.78	36.64	150.5	0.973	35
H0.5	131.39	22.44	120.4	0.007	32
H1	113.02	31.89	100.33	0.021	35
L0	175.84	38.69	225.25	0.000	36
L0.5	143.77	22.85	180.2	0.000	37
L1	139.66	33.33	150.17	0.057	40

* Two-tailed t -test; optimal order quantity as the null hypothesis.

We proceed by discussing the role of learning in conjunction with the ordering heuristics. We modify the mean anchor heuristic from Schweitzer and Cachon (2000), such that the anchoring point for the order quantity (which is not necessarily the mean demand in our case) can be estimated from regression analysis. Specifically, let E denote the anchor used by the subjects, Q denote their actual order quantity, and let Q^* denote the rational order quantity. Further, we introduce a coefficient λ , serving to measure the degree to which the subject adjusts her order quantity from the anchor toward the optimum.

$$\lambda = \frac{Q - E}{Q^* - E}, \quad (\text{A.1})$$

where $\lambda \in [0, 1]$, with $\lambda = 0$ implying the subject orders $Q = E$, whereas $\lambda = 1$ implies the subject orders $Q = Q^*$. Rewriting (A.1) in the following regression form highlights the role of λ in shifting weight between the optimum and the anchor:

$$Q = \lambda Q^* + (1 - \lambda)E. \quad (\text{A.2})$$

To capture the effect of strategic customers in leveraging the ordering behavior, we also include α in our regression. Specifically, we run a linear regression on the entire dataset,

while accommodating individual heterogeneity by employing a random effect on subjects²²:

$$Q_{it} = a_0 + a_1\alpha + a_2Q^* + u_i + \epsilon_{it}. \quad (\text{A.3})$$

The resulting estimates $\{\hat{a}_0, \hat{a}_1, \hat{a}_2\}$ provide a meaningful inference of the underlying heuristic, as (A.3) and (A.2) imply $\hat{\lambda} = \hat{a}_2$ and $\hat{E} = \frac{\hat{a}_0 + \hat{a}_1\alpha}{1 - \hat{a}_2}$.

The estimation results of this regression are provided in column (1) of Table 7. This estimation reveals that subjects allocate only about 30% of the weight for the rational order quantity, while the remainder is shifted toward the anchor which is estimated as $142.5 - 19.3\alpha$. This anchor is less than the mean demand, 150.5, and its (downward) deviation from the mean demand increases at a rate of 19.3 with the proportion of strategic customers.

To account for the potential chronological effects, we enrich the regression by adding a term $t \in \{1, 2, \dots, 50\}$, representing the round number:

$$Q_{it} = a_0 + a_1\alpha + a_2Q^* + a_3t + u_i + \epsilon_{it}. \quad (\text{A.4})$$

The modified estimation of the anchor is then $\hat{E} = \frac{\hat{a}_0 + \hat{a}_1\alpha + \hat{a}_3t}{1 - \hat{a}_2}$.²³ The resulting estimation is provided in column (2) of Table 7. This estimation reveals that the anchor E moves over time toward the mean demand—the value it would take under the traditional pull-to-center effect. In fact, when $t = 50$ and $\alpha = 0$, the value of the anchor is very close to the mean (153.1), which supports the idea of pull-to-center in traditional newsvendor environment. The involvement of strategic customers leads to a reduction up to 19.3 units (when all customers are strategic), resulting in an anchor of 133.8 at $t = 50$. Even when only half of the customers are strategic, the anchor at $t = 50$ is 143.5. This reveals that, although learning can bring the anchor closer to the mean, the presence of strategic customers still

²²An alternative way to incorporate individual heterogeneity is to allow the regression parameters a_0, a_1, a_2 , and a_3 to differ by subject. This is, however, infeasible with the between-subjects design in our case, as each subject faces a fixed Q^* and α .

²³In our learning analysis, subjects shift the anchor over time rather than shift the weight on the anchor (as \hat{E} is a function of t , but not so is $\hat{\lambda}$). By doing so, we are able to examine how the generalized pull-to-center bias evolves, rather than how the behavior evolves given a generalized pull-to-center bias.

profoundly lowers the anchor, which gives rise to the pull-below-center bias described in the main text of this paper.^{24,25}

Table 7: Estimation results of the observed order quantities

Coefficient	(1)	(2)
a_2	0.293** (0.0660)	0.293** (0.0660)
a_1	-13.66* (6.422)	-13.66* (6.422)
a_3		0.316** (0.0246)
a_0	100.5** (12.74)	92.45** (12.76)
N	10750	10750
r_w^2	1.29e-32	0.0155
r_o^2	0.101	0.110
r_b^2	0.221	0.221
χ^2	60.01	225.3
p	9.33e-14	1.41e-48
E	$142.5 - 19.3\alpha$	$130.8 - 19.3\alpha + 0.45t$

Notes: Standard deviations in parentheses;

** $p < .01$, * $p < .05$.

Lastly, we note that the σ values from Table 3 can be perceived as approximations to the degree of rationality, because they represent the noise level in the decisions. We tested pairwise differences among them. We found that there is no statistical difference in σ between high and low cost treatments for a given α . Under the Low cost condition, σ_{L0} is statistically larger than $\sigma_{L0.5}$, and there is no statistical difference between σ_{L0} and σ_{L1} and between $\sigma_{L0.5}$ and σ_{L1} . Under the high cost condition, we found that σ_{H0} is statistically larger than both $\sigma_{H0.5}$ and σ_{H1} , and there is no statistical difference between $\sigma_{H0.5}$ and σ_{H1} . Therefore, it

²⁴In the spirit of Bostian et al. (2008), we could have developed a more sophisticated anchor heuristic that evolves over time such that, e.g., $\lambda_{it} = (1 + \Delta_i)\lambda_{i,t-1}$, where Δ_i captures the idiosyncratic adjustment of the weight. However, as we have demonstrated in this section, learning does not play a major role in altering the pull-below-center bias stemming from the presence of strategic customers, and hence any further exploration of learning mechanisms are out of the scope of this paper.

²⁵As subjects shift the anchor, they do not simply make less mistakes as the game proceeds—as seen from Figure 3, their orders move slightly upwards in most treatments (perhaps except for H1), but do not necessarily approach the optimal levels.

seems that the conditions with no strategic customers have a higher level of noise in the ordering decisions. Since we do not incorporate a rationality parameter explicitly in the model, we abstract away from discussing this aspect. Nevertheless, by visual observation of Figure 3, which plots confidence intervals by rounds, it looks like panels (a) and (d) of this figure have larger intervals, which is consistent with the statistics mentioned above.

Appendix B Individual-level behavior

In the main text the pull-below-center bias was identified on the aggregate level. In this section we will show that the same behavior also persists at the individual level. Figure 4 presents the average order placed by each subject, along with the corresponding 95% confidence interval for the orders shown as the vertical bar, and the respective optimal order quantity represented by the dotted line. Despite the subject heterogeneity observed in every treatment, more subjects in L0.5 generally order less than the mean demand than those showing the traditional pull-to-center bias (specifically, out of the 37 subjects, 25 exhibit PBC whereas 12 exhibit P2C based on their average orders – see Figure 4); similarly in treatment L1, there are apparently more participants who pull below center (specifically, 28 out of the 40 subjects based on average orders) than those who order above, and this asymmetry clearly exceeds the level that can be accounted by random noise. Thus the individual-level behavior echoes the aggregate-level findings of pull-below-center effect in treatments L0.5 and L1. In places where pull-to-center effect stretches out on aggregate, there is always a portion of subjects (albeit not dominant in the population) whose orders are pulled below center (see L0, H0, H0.5, H1 in Figure 4). In Proposition 3 (main text), we establish psychological conditions for both ordering more and less than center, thereby accommodating the individual heterogeneity in these treatments.

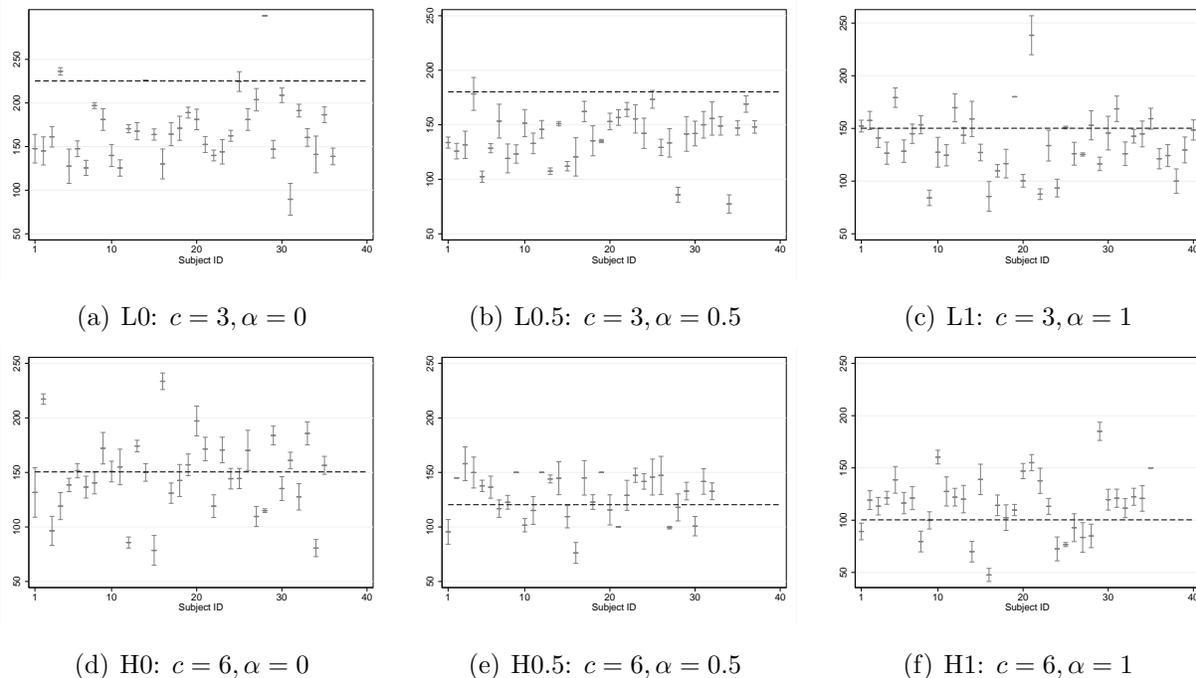


Figure 4: Orders by Subject (— Observed --- Optimal)

Appendix C Sample Instructions ($c = 3$)

Instructions

Welcome to this experiment on economic decision-making. If you follow the instruction carefully and make good decisions, you may earn a considerable amount of money that will be paid to you in cash at the end of the experiment. If you have any questions, please raise your hand and the experimenter will answer your questions in private. We ask you not to talk or makes noises, and that you do not look at the decisions of others during the experiment.

Description of the Decision Task

In this experiment, you will play the role of a retailer selling a seasonal product. There are 50 selling seasons, each of which consists of two selling periods: early selling period and late selling period. The unit retail price during the early selling period is 12 points and the unit retail price during the late period is 6 points. At the beginning of each season, you will need to decide how many units of the seasonal product to order from your supplier at a cost of 3 points for each unit. Notice that you decide on your order quantity before knowing the number of customers that will visit your retail store. Your order quantity has to be an integer value.

Once you have placed your order, your customers arrive, each of whom desires to purchase one unit of the product. The number of customers could be any integer value between 1 and 300 with equally probability. That is, the values 1, 2, 3, ..., 299, and 300 can equally

likely be the number of customers visiting your store during this selling season. Whether a customer will obtain a unit of the product depends on the total number of customers and the number of units you have ordered.

The customers that visit your retail store are classified into two types: simple customer and sophisticated customer.

Simple customers always purchase the product in the early selling period.

Sophisticated customers realize that they may obtain the product at a lower price during the late selling period. Hence, they decide whether to purchase in the early selling period or in the late selling period according to the following rule:

- If the total number of customers (i.e. the number of simple customers + that of sophisticated customers) is larger than the number of units you ordered, then all the sophisticated customers will try to purchase in the early selling period (at a price of 12 points per unit).
- If, on the other hand, the total number of customers is smaller than the number of units your ordered, then all the sophisticated customers will purchase in the late selling period (at a price of 6 points per unit).

Once the game starts, you will know the mixture of simple vs. sophisticated customers that you face. It could be either:

- All customers are simple,
- 50% of your customers are simple and the other 50% are sophisticated, or
- All customers are sophisticated.

All unsold products at the end of each selling season are disposed. No inventory is carried from one season over to the next. Your profit in each season depends on your ordering decision and the total number of customers, and is calculated as follows.

- If the number of customers is larger than your order quantity, then all the customers will purchase in the early selling period, but you cannot sell more than the available number of units you have ordered. In this case, your profit for the season is: Profit = (early period price) * (order quantity) - (unit cost)*(order quantity). Since the early period price is 12 points and the unit cost is 3 points, your profit is Profit = 12 * (order quantity) - 3 *(order quantity).
- If the number of customers is smaller than your order quantity, then all simple customers purchase in the early selling period and all sophisticated customers purchase in the late selling period. In this case you will have to dispose the unsold products, and your profit for the season is Profit = (early period price) *(Number of simple customers) +(late period price) *(Number of sophisticated customers) - (unit cost)*(order quantity). Since the early period price is 12 points, the late period price is 6 points, and the unit cost is 3 points, your profit is Profit = 12 * (Number of simple customers) + 6 *(Number of Sophisticated customers) - 3 *(order quantity).

Examples:

Assume that 50% of your customers are simple and the other 50% are sophisticated. Assume you have placed an order of 100 units and the number of customers who visit your retail store is 120. Since the total number of customers is large than the order quantity ($120 > 100$), all the customers purchase in the early period. As a consequence, you sell all 100 units you ordered. Your profit for the season in this example is $12 * 100 - 9 * 100 = 300$ points.

Now consider another example. Assume you have placed an order of 150 units and the number of customers that visit your retail store is 90. Since the total number of customers is smaller than the order quantity ($90 < 150$), all the simple customers purchase in the early period and all the sophisticated customers purchase in the late period. Since 45 customer are simple and the other 45 customers are sophisticated, your profit for the season in this example is $12 * 45 + 6 * 45 - 3 * 150 = 810 - 450 = 360$ points.

Note that the above two examples are for illustrative purposes only. You may or may not see the same numbers in the experiment.

At the end of each season, you will be informed of your decision outcomes and profit. You will have to make ordering decision for 50 seasons. For each season, a new number of customers will be drawn between 1 and 300. The number of customers in each season is independent of the number of customers realized in any other season. The unit cost of product, the unit price of early selling period and the unit price of late selling period are the same in all seasons.

The payment you will receive consists of two parts: 1) a show up fee of \$5 cash, and 2) a performance reward determined from your total profit accumulated over 50 seasons in the experiment, at a conversion rate as 1 point = \$0.001 in real cash. That is, for every 1000 points you have earned in the experiment, you will receive a \$1 performance reward in real cash. All payment will be in cash, and paid to you upon the completion of the entire experimental session.