2008

Bounded Rationality in Newsvendor Models

Xuanming Su
University of Pennsylvania

Recommended Citation

Su, X. (2008). Bounded Rationality in Newsvendor Models. *Manufacturing & Service Operations Management, 10* (4), 566-589. http://dx.doi.org/10.1287/msom.1070.0200

At the time of publication, author Xuanming Su was affiliated with the University of California. Currently (July 2016), he is a faculty member in the Operation, Information and Decisions Department of the Wharton School at the University of Pennsylvania.

Bounded Rationality in Newsvendor Models

Abstract

Many theoretical models adopt a normative approach and assume that decision makers are perfect optimizers. In contrast, this paper takes a descriptive approach and considers bounded rationality, in the sense that decision makers are prone to errors and biases. Our decision model builds on the quantal choice model: While the best decision need not always be made, better decisions are made more often. We apply this framework to the classic newsvendor model and characterize the ordering decisions made by a boundedly rational decision maker. We identify systematic biases and offer insight into when overordering and underordering may occur. We also investigate the impact of these biases on several other inventory settings that have traditionally been studied using the newsvendor model as a building block, such as supply chain contracting, the bullwhip effect, and inventory pooling. We find that incorporating decision noise and optimization error yields results that are consistent with some anomalies highlighted by recent experimental findings.

Keywords

bounded rationality, newsvendor, logit choice, random utility, quantal response, supply chain, bullwhip effect, inventory, pooling

Disciplines

Operations and Supply Chain Management | Other Business | Policy History, Theory, and Methods

Comments

At the time of publication, author Xuanming Su was affiliated with the University of California. Currently (July 2016), he is a faculty member in the Operation, Information and Decisions Department of the Wharton School at the University of Pennsylvania.

Bounded Rationality in Newsvendor Models

Xuanming Su

Haas School of Business, University of California, Berkeley, CA 94720, USA
xuanming@haas.berkeley.edu

Many theoretical models adopt a normative approach and assume that decision-makers are perfect optimizers. In contrast, this paper takes a descriptive approach and considers *bounded rationality*, in the sense that decision-makers are prone to errors and biases. Our decision model builds upon the *quantal choice* model: while the best decision need not always be made, better decisions are made more often. We apply this framework to the classic newsvendor model and characterize the ordering decisions made by a boundedly rational decision-maker. We identify systematic biases and offer insight into when over-ordering and under-ordering may occur. We also investigate the impact of these biases on several other inventory settings that have traditionally been studied using the newsvendor model as a building block, such as supply chain contracting, the bullwhip effect, and inventory pooling. We find that incorporating decision noise and optimization error yields results that are consistent with some anomalies highlighted by recent experimental findings.

1. Introduction

The newsvendor model is one of the main building blocks of inventory theory. The model derives its name from the canonical setting of a newsvendor facing random demand, who has to decide how many copies of newspapers to order: excess quantities that remain unsold have no value, but ordering too few copies means that customers have to be turned away and potential profits are lost. This model has a simple and elegant solution that offers insights into the optimal balance between the costs of supply-side investment and the costs of potential foregone profits. It has extensive applications, including inventory management, capacity planning, and pricing and revenue management.

The current theoretical literature is based primarily on the paradigm of perfect rationality. In existing models, the newsvendor is a perfect optimizer: without fail, he always chooses stocking levels that attain the maximum possible level of expected profits. This infallible newsvendor is an important workhorse of inventory theory and many results are based upon it. In contrast, recent experimental studies suggest that human decision-makers do not solve these inventory problems as theory predicts. One of the most robust experimental findings is that subjects in newsvendor experiments systematically deviate from the optimal critical fractile solution. Other anomalies have also been identified in more general supply chain settings. These empirical deviations challenge the validity of theoretical results in the literature. This is a significant issue because in practice, many of these newsvendor-type decisions are made by human decision-makers who suffer from similar psychological biases. In this regard, we feel that theoretical results based on perfect rationality need to be reconciled with actual human behavior.

The goal of this paper is to investigate the effect of bounded rationality in traditional newsvendor models. What happens when newsvendors make mistakes? Are the consequences of these errors consistent with experimental findings? To address these questions, we adopt a descriptive decision framework that allows for decision errors. As these errors become negligibly small, we obtain the standard normative setting of perfect rationality as a special case. We hope to generalize existing theoretical results while capturing a wider spectrum of empirically observed behavior.

Our decision model of bounded rationality is derived from classical *quantal choice* theory (see Luce, 1959). When faced with alternatives $i \in \mathcal{I}$ generating utility u_i, decision-makers do not always choose the utility-maximizing alternative $i^* \in \arg\max_i u_i$. Instead, all possible alternatives are candidates for selection, but more attractive alternatives (yielding higher utility) are chosen with larger probabilities. For analytical convenience, we focus on the logit choice model in which the probability of choosing alternative i is proportional to e^{u_i} (see McFadden, 1981, and Anderson, de Palma, and Thisse, 1992). When applied to newsvendor models, we interpret each possible order quantity x as a "candidate alternative" and the corresponding expected profit $\pi(x)$ as the "utility." The probabilistic choice setup implies that the newsvendor is subject to decision noise and may make suboptimal ordering decisions. Unlike conventional models, newsvendor order quantities are no longer deterministic (at the optimal quantity $x^* = \arg\max_x \pi(x)$); they are now random variables. Nevertheless, order quantities that lead to higher expected payoffs are chosen more often. On the terminology: we stress that "bounded rationality" refers to a wide range of behavioral phenomena (e.g., psychological biases, heuristics or rules of thumb, cognitive constraints), and we shall focus on noisy decision-making as one possible way of incorporating bounded rationality.

Next, we summarize our main results. Using the logit decision model, we offer a complete characterization of the boundedly rational newsvendor's order quantity (i.e. its distribution).

For the special case of uniform demand, we find that the order quantity follows a normal distribution (truncated at the appropriate cutoff points). For the general case, we offer a method to calculate the choice distribution and other quantities of interest, such as expected orders and expected profits. Then, we apply this decision framework to several inventory settings. Table 1 organizes our findings and compares them with the conventional predictions under perfect rationality. First, for basic newsvendor decisions, we show that random decision noise yields systematic biases from the optimal critical fractile solution, and we identify conditions that lead the boundedly rational newsvendor to over-order or under-order. Next, in supply contracting, while much of the literature proposes to achieve coordination by allocating a fixed fraction of total profits to the decision maker, we show that this is not feasible in our model of bounded rationality because the reduced stakes also diminishes the decision-maker's incentive to make "good" decisions. Next, our model shows that the bullwhip effect may arise when decision-makers do not trust their supply chain partners to order optimally and thus take actions to guard against and correct for others' biases. Finally, in inventory pooling, apart from the benefits that have been associated with variance reductions resulting from summing random variables, we also identify the (additional) behavioral benefits of pooling. As summarized in Table 1, there is a large number of experimental observations that seem to be at odds with conventional theory, but our analysis will show that many of them are consistent with bounded rationality explanations.

There are two main contributions in this paper. First, we apply the quantal choice framework to a fundamental operations model. The basic premise that "people need not optimize, but better decisions are made more often" is intuitively appealing. This quantal choice approach is sufficiently general to be applicable to a wide variety of settings that extend beyond newsvendor models, while remaining extremely parsimonious in that there is only a single parameter to be estimated. We stress that this represents an alternative theoretical paradigm in operations management. Instead of solving optimization problems (normative approach), we incorporate bounded rationality and characterize outcomes based on probabilistic choice models (descriptive approach). Second, our results (accounting for bounded rationality) are consistent with a wide range of experimental observations that seem to be at odds with conventional theory (assuming perfect rationality). Hence, we feel that decision noise and optimization error deserve consideration as a possible explanation for these empirical newsvendor anomalies. We hope that our results complement the behavioral and theoretical rationales that have been offered in the literature.

Context	Conventional results under *perfect rationality*	Some contradicting experimental / empirical / theoretical results	Can these inconsistencies be fully explained by *bounded rationality*?
Basic Newsvendor	• Critical fractile solution	• Systematic over-ordering and under-ordering (Schweitzer and Cachon, 2000) • Bias toward the mean (Bostian et. al, 2007) • Bias toward low-probability demand realizations • Learning effects (Bolton and Katok, 2005; Lurie and Swaminathan, 2005) • Previous-period effects (Benzion et. al, 2005)	• Yes • Yes • Yes • No • No
Supply Chain Contracting	• Double marginalization problem • Different types of coordinating contracts	• Double marginalization may be beneficial • Coordination is not achieved (Katok and Wu, 2007) • Stake sizes and the value of committing to standing orders (Bolton and Katok, 2005) • Differences between contractual forms (Katok and Wu, 2007) • Fairness concerns (Keser and Paleologo, 2004; Wu and Loch, 2007)	• Yes • Yes • Yes • No • No
Bullwhip Effect	• Bullwhip effect (variance of orders increases upstream) • Physical causes	• Behavioral causes of bullwhip effect (Sterman, 1989; Croson and Donohue, 2006) • Under-weighting of supply line (Sterman, 1989; Croson and Donohue, 2006) • Need to account for mistakes of supply chain partners (Croson et. al, 2005) • Demand uncertainty and supply uncertainty may both propagate upstream	• Yes • Yes • Yes • Yes
Inventory Pooling	• Pooling benefits • No benefit when demand is perfectly correlated	• Behavioral benefits of pooling • Cost savings when demand is perfectly correlated • Cost savings when demand is deterministic • Pooling of demand uncertainty and supply uncertainty	• Yes • Yes • Yes • Yes

Table 1: Summary of newsvendor-type results under perfect rationality and bounded rationality

The remainder of the paper is organized as follows. The literature is reviewed in Section 2. We describe our model of bounded rationality in Section 3 and apply it to the newsvendor problem in Section 4. In Section 5, we fit our model to experimental data and show that the data provides empirical support for our specification of bounded rationality. Next, we analyze the effect of bounded rationality in several different contexts, as shown in Table 1: Section 6 discusses over-ordering and under-ordering, Section 7 explains why "coordinating contracts" may fail to coordinate the system, Section 8 shows how bounded rationality generates the bullwhip effect, and Section 9 identifies the behavioral benefits of inventory pooling. Finally, we offer concluding remarks in Section 10. All proofs are provided in the Appendix.

2. Literature Review

Traditional theory associate rationality with the ability to optimize perfectly: rational agents will settle for nothing less than the best. In contrast, the concept of bounded rationality recognizes the inherent imperfections in human decision-making. The seminal work of Simon (1955) proposes "satisficing" as a more accurate way to model decision-making behavior: rather than optimizing perfectly, agents search over the choice domain until they find something satisfactory. Another broad approach towards bounded rationality is to study heuristics or rules of thumb; see Geigerenzer and Selten (2001). When the "bounds" on rationality render optimization infeasible, agents may instead adopt simple heuristics to make complex decisions. Several well-studied examples include the representativeness heuristic, the availability heuristic, and the anchoring and adjustment heuristic, which are described in Tversky and Kahnemann (1974). Yet another approach is to explicitly model agents' cognitive limitations and the computational complexity of decision tasks. Rubinstein (1998) surveys this work and discusses it in the context of economic models of decisions and games. For a review of the evolution and development of bounded rationality, readers are referred to Simon (1982) and Conlisk (1996).

This paper models bounded rationality by incorporating stochastic elements into the decision process. Instead of choosing the utility-maximizing alternative all the time, decision-makers adopt a probabilistic choice rule such that more attractive alternatives are chosen more often. In particular, we focus on the logit choice rule. Our approach is related to three separate streams of literature. First, there is a rich academic tradition on stochastic choice rules with the consistency property that *better options are chosen more often.* This approach can be traced back to Thurstone (1927) and Luce (1959), who set up the mathematical framework and develop invariance properties. Blume (1993) motivates this stochastic approach by showing that the choice distributions are analogous to Gibbs states, which have proven to be a useful tool in studying Ising models in statistical mechanics even though their stationary distributions are not completely known. McKelvey and Palfrey (1995) develop a framework that admits generalizations of the “better-options-are-chosen-more-often” structure and applies it to game-theoretic settings. Chen, Friedman, Thisse (1997) consider individuals with latent utility functions and study the stochastic choice probabilities that emerge. Next, our model of bounded rationality is mathematically equivalent to the *random utility* approach in discrete choice models (see Anderson, de Palma, and Thisse, 1992). In these models, individuals’ utilities over different alternatives have idiosyncratic taste shocks reflecting unobserved heterogeneity. The characterization of stochastic choice models as random utility models was first established by Block and Marschak (1960). The logit choice framework was originally developed by Luce (1959) and McFadden (1981). The third stream of related models involves *evolutionary adjustment* in decision processes. See, for example, Young (1993), Binmore and Samuelson (1997), Hofbauer and Sandholm (2002) and Anderson, Goeree, Holt (2004). In these papers, a main interest is in characterizing the steady-state distribution of decisions over the long run. In particular, Anderson, Goeree, Holt (2004) develop a model in which agents adjust their decisions toward higher payoffs, subject to normal error, and show that the long run steady state distribution of this Gaussian process agrees with the logit choice rule. In our view, the three streams of work reviewed above provide different justifications to our model of bounded rationality. For concreteness, we shall focus on the first interpretation; in other words, when we refer to “bounded rationality,” we mean that individuals need not always pick the best option, but they choose better options more often.

There is a recent stream of work on behavioral operations management; see reviews by Bendoly, Donohue, and Schultz (2006), Gino and Pisano (2006), and Loch and Wu (2007). This emerging literature points out fundamental inconsistencies between empirical observations and theoretical predictions in a variety of operations settings, and underscore the need to reconcile these findings. In a similar spirit, our current work applies quantal choice models to newsvendor-type scenarios, and shows that bounded rationality provides a potential explanation for some of these inconsistencies. Here, our goal is theory-building: we seek to extend the decision-making foundations of existing theory in order to better match empirical observations in the laboratory.

We shall organize this review into the four areas listed in Table 1. First, for the basic newsvendor model, it is well-known that the optimal solution is characterized by the critical fractile; see Porteus (2002). However, Schweitzer and Cachon (2000) present experimental evidence of decision biases in this basic model. With a uniform demand distribution, they find that subjects tend to order too much of "low-profit" products and too few of "high-profit" products. These results are consistent with two behavioral explanations: subjects may have a preference to reduce ex-post inventory error, or they may suffer from the anchoring and insufficient adjustment heuristic (so their orders are biased towards the mean). There are subsequent studies that conduct experiments to investigate the effect of feedback and learning in the newsvendor problem. Bolton and Katok (2005) show that requiring newsvendors to commit to standing orders focuses their attention on long-term profits and results in better decision-making. Benzion, Cohen, Peled, and Shavit (2005) identify a significant previous-period effect, but it weakens over time as subjects learn. Lurie and Swaminathan (2005) show that more frequent feedback may sometimes degrade performance. Bostian, Holt, and Smith (2007) find that the bias of order quantities towards the mean can be explained by an adaptive learning model. In this paper, we demonstrate that bounded rationality, modeled via stochastic choice rules (i.e. decision noise), can generate some of these laboratory observations.

Second, we position our work with respect to the theoretical literature on supply chain coordination, which is well-developed. When decisions in a supply chain are made by individual parties with misaligned interests, the double marginalization problem arises and system optimal profits can not be attained. To rectify this situation, various supply contracts have been proposed. These contracts coordinate the system by aligning individual incentives with system objectives. Examples of such contracts include buy-back contracts (e.g., Pasternack, 1985), quantity flexibility contracts (e.g., Tsay, 1999), markdown money (e.g., Tsay, 2001), sales rebates (e.g., Taylor, 2002), and revenue sharing contracts (e.g., Cachon and Lariviere, 2005). For a review of the supply chain contracting literature, readers are referred to Cachon (2003). Recently, Katok and Wu (2006) investigate the performance of these coordinating contracts in the laboratory. They test the performance of two mechanisms: buy-back contracts and revenue sharing contracts. They observe that, in contrast to theoretical predictions, coordination is not achieved. In a similar spirit, we also find that coordination may not be feasible when the decision-maker is boundedly rational.

Third, we discuss the theoretical and experimental literature related to the bullwhip effect. The bullwhip effect refers to the tendency for the variance of orders to increase upstream along the supply chain. This is demonstrated in the important paper by Lee, Padmanabhan, and Whang (1997), who also identify four separate causes of bullwhip effect: demand signal processing, inventory rationing, order batching, and price fluctuations. Theoretical studies suggest that in the absence of the physical causes listed above, the bullwhip effect will not arise. On the experimental side, the first study demonstrating the bullwhip effect is by Sterman (1989). There are two important contributions in this seminal paper: first, the study introduces an empirical framework for predicting subjects' choices by assuming that they follow a decision rule based on the anchor-and-adjust heuristic, and second, it identifies underweighting of the supply line as another cause of the bullwhip effect. That is, because subjects do not fully account for quantities in the supply line, they may over-order and generate instability that triggers off the bullwhip effect. In a subsequent experiment with a commonly known demand distribution, Croson and Donohue (2006) controls for all four physical causes but find that the bullwhip effect still persists. Finally, even in an experiment in which the demand is constant and publicly known, Croson, Donohue, Katok, and Sterman (2006) find evidence of the bullwhip effect. They suggest that it arises because subjects may place excessive orders to address the perceived risk that others will not behave optimally, and call this coordination risk. These experimental studies show that beyond its physical properties, the bullwhip effect is also very much a behavioral phenomenon. In this paper, we apply our framework of boundedly rational decision-making in such settings. To a large extent, we find that the theoretical predictions of our model agree with the experimental findings reviewed above.

Fourth, we discuss the theoretical literature related to inventory pooling. The classic study by Eppen (1979) shows that in a multi-location inventory setting, consolidating stocks at a centralized location (instead of holding separate inventories at individual locations) leads to a reduction in total costs; further, the magnitude of these pooling benefits depends on the correlation of demands. Although inventory pooling leads to lower costs, Gerchak and Mossman (1992) show that it does not necessarily lead to lower inventory levels. We extend this work by showing that beyond the physical benefits, there are also behavioral benefits to inventory pooling.

From a meta-modeling perspective, the existing body of work on random supply processes provides a physical analogue to our behavioral notion of bounded rationality. This literature, which traces back to Karlin (1958), studies the impact of physical phenomena such as yield uncertainty in the production process, supply unreliability and disruptions (e.g., breakdowns, natural disasters or labor strikes), and inventory record inaccuracy. For a sample of different modeling approaches, readers are referred to the review by Lee and Yano (1995), as well as more recent papers by Chen, Yao, and Zheng (2001), Tomlin and Wang (2005), Kok and Shang (2006), and Dada, Petruzzi and Schwarz (2006). Ultimately, the modeling root in most of this literature, which also forms the core of our notion of bounded rationality, is that the supply X is a random variable. The causes may be behavioral or physical in nature, but the consequence is the same: supply is uncertain. Under physical constraints, the agent is capable of making an optimal decision but may still experience a suboptimal outcome because the agent is not perfectly capable of implementing the decision. In contrast, under behavioral constraints, the agent faces a suboptimal outcome because he/she is not able to make the optimal decision in the first place (even though the decision, once made, can be implemented accordingly). Although the motivation and practical contexts behind these two perspectives are completely different, they can be examined from a similar modeling angle.

There are two noteworthy differences between supply uncertainty and bounded rationality as modeled in this paper. First, we capture boundedly rational decision-making by postulating that better choices are made more often. This implies that the resulting decision noise is intricately related to the underlying decision problem. In contrast, most practical interpretations of supply uncertainty models do not require a close relationship between supply disruptions and the underlying economic context; for instance, less severe disruptions (analogous to better decisions) do not necessarily occur more frequently. Second, in many models of supply uncertainty, the actual quantity being supplied is usually less than the intended quantity. On the other hand, in our model of bounded rationality, the chosen quantity may either be larger or smaller than the optimal quantity; the only consistency condition we impose is that better decisions are made more often. Nevertheless, from a managerial perspective, the insights that can be gleaned from studying random supply processes may have analogous interpretations for behavioral settings of bounded rationality. In our analysis below, we shall draw such parallels wherever possible.

3. A Model of Bounded Rationality

The standard approach in most normative analysis assumes perfect rationality on the part of the decision-maker. Specifically, when faced with a choice among different alternatives

$i \in \mathcal{I}$, the perfectly rational decision-maker always chooses the most preferred option(s) $i^* \in \arg\max_i u_i$. In contrast, to capture bounded rationality, we apply the multinomial logit choice model and assume that the decision-maker chooses alternative $i \in \mathcal{I}$ with probability

$$\psi_i = \frac{e^{u_i/\beta}}{\sum_{i \in \mathcal{I}} e^{u_i/\beta}}. \tag{1}$$

Similarly, the logit choice probabilities over a continuous domain $\mathcal{Y}$ are given by the density

$$\psi(y) = \frac{e^{u(y)/\beta}}{\int_{y \in \mathcal{Y}} e^{u(y)/\beta}} \tag{2}$$

with distribution $\Psi(y) \equiv \int_{-\infty}^{y} \psi(v)dv$. In other words, the agent's choice is a random variable $Y \in \mathcal{Y}$. As noted by Anderson, de Palma and Thisse (1992) (page 4), this probabilistic approach provides a way to model bounded rationality. With this logit structure, *better alternatives are chosen more often.* Although the best option is no longer chosen with probability one, it nonetheless will be the mode of the choice distribution. The logit model is sometimes called the log-linear model because the log odds of choosing one alternative over another is proportional to the payoff difference between the two alternatives.

The parameter β can be interpreted as the extent of cognitive and computational limitations suffered by the decision-maker. To understand this, observe that as $\beta \to \infty$, the choice distribution in (1) approaches the uniform distribution over $\mathcal{I}$ in the limit. In this extreme case, the decision-maker lacks the ability to make any informed choices and instead randomizes over the alternatives with equal probabilities. On the other hand, as $\beta \to 0$, the choice distribution in (1) becomes entirely concentrated on the utility-maximizing alternative (assuming it is unique), and this coincides with the choice of a perfectly rational decision-maker; when there are multiple alternatives attaining the maximum utility, the choice distribution approaches the uniform distribution over these utility-maximizing alternatives, which is also consistent with perfect rationality. Therefore, we shall interpret the magnitude of β as the extent of bounded rationality.

The multinomial logit model described above has been frequently used in different contexts. Under one interpretation, the noise terms ϵ_i reflect heterogeneity that is unobserved by the modeler; despite the presence of noise, the decision-maker is perfectly rational, but he is just taking some unobserved factors into account. Alternatively, as in our model, the noise terms are the explicit result of bounded rationality. Both interpretations are reasonable, and they lead to the same probabilistic choice outcomes. In this paper, we shall adopt the bounded rationality interpretation as this facilitates comparing our results with recent experimental findings, in which all other external attributes have been controlled for.

Before proceeding, we put forth two invariance properties. First, consider affine transformations of the decision domain, i.e., instead of choosing $y \in \mathcal{Y}$, suppose that the decision-maker chooses $\tilde{y} \in \tilde{\mathcal{Y}}$, where $\tilde{y} \equiv ay + b$ and $\tilde{\mathcal{Y}} \equiv a\mathcal{Y} + b$ for some constants a and b. Such transformations of the decision domain does not affect utility, so the utility function over $\tilde{\mathcal{Y}}$ is given by $\tilde{u}(\tilde{y}) = u(y)$. This can be interpreted as purely a change in the way choices are named or labeled. In particular, the multiplicative factor a changes the units of the choices (e.g. from kilograms to tonnes) and the additive term b changes the location of "zero."

Lemma 1. *The choice distribution is invariant to affine transformations of the decision domain. Specifically, let $\Psi(y)$ be the choice distribution over $\mathcal{Y}$ and let $\tilde{\Psi}(\tilde{y})$ be the choice distribution over $\tilde{\mathcal{Y}}$. Then, $\Psi(y) = \tilde{\Psi}(\tilde{y})$.*

There is a similar result for additive transformations in the utility function, i.e., the decision-maker faces utility function $\tilde{u}(y) = u(y) + a$ instead of $u(y)$.

Lemma 2. *The choice distribution is invariant to additive transformations in the utility function. Specifically, let $\Psi(y)$ be the choice distribution with respect to $u(y)$ and let $\tilde{\Psi}(y)$ be the choice distribution with respect to $\tilde{u}(y)$. Then, $\Psi(y) = \tilde{\Psi}(y)$.*

However, the choice distribution is affected by multiplicative transformations in the utility function. As we shall see in the subsequent analysis, this effect has important implications. Most significantly, it suggests that stake sizes have an impact on decision outcomes.

4. The Newsvendor Model under Bounded Rationality

We now apply the logit choice framework to the newsvendor problem. Recall that the canonical setting involves a newsvendor who has to determine how many copies of newspapers to order. Each copy costs c but can be sold at price p, where $p > c$. The random demand D has density f and distribution F; we shall write $\bar{F} \equiv 1 - F$. Demand that is not fulfilled is lost, and leftover copies have zero value. (Although it is straightforward to incorporate a salvage value, we prefer to suppress it for notational clarity; by virtue of Lemma 2, the analysis can easily be modified to include a salvage value s by replacing p and c by $p - s$ and $c - s$.) Given this setup, the newsvendor's expected profit when ordering x copies is

$$\pi(x) = pE\min(D, x) - cx, \tag{3}$$

which is uniquely maximized at $x^* = F^{-1}(\xi)$; here, $\xi \equiv 1 - (c/p)$ is the critical fractile and $1 - \xi \equiv c/p$ is the optimal stockout probability.

Let us introduce some terminology. We shall refer to the profit-maximizing ordering quantity x^* as the *optimal* solution, which is chosen whenever the newsvendor is perfectly rational. However, under bounded rationality, the newsvendor's ordering quantity is subject to noise and becomes a random variable. We shall refer to this as the *behavioral* solution and denote it using $x^\flat$ (for the realization) and $X^\flat$ (for the random variable).

Given the problem data, it is straightforward to use (2) to write down the behavioral solution of the newsvendor problem. We assume that the decision domain $S \subseteq \mathbb{R}$ is the smallest interval containing the support of f. In other words, the boundedly rational newsvendor may order any quantity between the smallest possible and largest possible demand realizations. Then, the probability density function of the behavioral solution is

$$\psi(x) = \frac{e^{\pi(x)/\beta}}{\int_S e^{\pi(v)/\beta} dv} = \frac{e^{(pE\min(D,x)-cx)/\beta}}{\int_S e^{(pE\min(D,v)-cv)/\beta} dv}. \tag{4}$$

We stress that the decision domain S plays an important role in this logit choice setup. This will become evident in subsequent analysis, where we study the behavioral solution $X^\flat$ through its density ψ.

4.1 Uniform Demand

Suppose that the demand D is uniformly distributed between a and b, with $b > a \geq 0$. Then, the newsvendor's profit function in (3) can be simplified into a quadratic function

$$\pi(x) = Ax^2 + Bx + C, \tag{5}$$

with coefficients $A = -\frac{p}{2(b-a)}$, $B = \left(\frac{pb}{b-a} - c\right)$, and $C = -\frac{pa^2}{2(b-a)}$. This quadratic structure is essential and gives us the following result.

Proposition 1. *Let $D \sim U[a, b]$. Then, the behavioral solution to the newsvendor problem follows a truncated normal distribution over $[a, b]$, with mean μ and variance σ^2 given by*

$$\mu = b - \frac{c}{p}(b-a), \tag{6}$$

$$\sigma^2 = \beta\frac{b-a}{p}. \tag{7}$$

Corollary 1. *Let $D \sim U[a,b]$. Then, the expected behavioral solution is*

$$EX^\flat = \mu - \sigma \cdot \frac{\phi(\frac{b-\mu}{\sigma}) - \phi(\frac{a-\mu}{\sigma})}{\Phi(\frac{b-\mu}{\sigma}) - \Phi(\frac{a-\mu}{\sigma})}, \tag{8}$$

where $\phi(\cdot)$ and $\Phi(\cdot)$ denote the standard normal density and distribution functions.

Observe that the parameter μ of the behavioral solution coincides with the optimal solution x^*. This is natural because the optimal solution uniquely maximizes expected payoffs and thus should be the unique mode of the behavioral solution, which occurs at μ for a truncated normal distribution. Next, observe that the variance of the behavioral solution σ^2 is proportional to β, as expected, since bounded rationality (larger β) increases noise in the decision-making process. It is also intuitive that the variance increases with the range $b-a$ of the demand distribution (a wider range generates a more complex decision task) but decreases with p (higher prices increases the stakes and results in better decisions).

The main message from this result is: under bounded rationality, uniform demand yields normally distributed choices (with appropriate truncations). From an experimental standpoint, there is no dearth of laboratory data for the uniform demand case. Many experiments are run using uniform demand because this is easier to understand for subjects. Therefore, our theory provides testable implications that can immediately be put to the test. In particular, given experimental data on subjects' ordering decisions, we may fit the truncated normal distribution to the data to obtain estimates of the parameters. Since the model of perfect rationality (with $\beta = 0$) is a special case of our model, we may test this hypothesis to detect the presence of bounded rationality. Significant evidence for $\beta > 0$ would support the presence of bounded rationality. This procedure is reported in detail in Section 5.

4.2 Triangular Demand

Next, we consider the special case of triangular demand distributions. Experimentally, apart from uniformly distributed demand, the triangular distribution is another special case that can be easily understood by the subjects. Besides, compared to the uniform distribution, the triangular distribution offers a more accurate representation of demand in practical settings.

Here, we consider triangular demand distributions with range [0,100]. This is without loss of generality via a straightforward translation. We use $h \in [0,100]$ to denote the peak of the triangular demand density. Then, the probability density function for demand is

$$f(x) = \begin{cases} \frac{x}{50h}, & x \le h, \\ \frac{100-x}{50(100-h)}, & x > h. \end{cases} \tag{9}$$

Under this demand density, it is straightforward to derive the newsvendor profit function $\pi(x)$ using (3). Observe that while the profit function is quadratic in the case of uniform demand, it is now cubic in the case of triangular demand. As such, the behavioral solution $X^\flat$ can not be characterized using standard probability distributions. Therefore, we shall proceed to study its properties numerically.

To generate Figure 1, we consider several different triangular demand distributions, with peak densities at $h = 0, 20, 40, 60, 80, 100$; the range is maintained as $[0, 100]$. Each of these demand distributions is represented by one of the six charts in Figure 1. We set price $p = 1$. For each demand density, we plot the expected behavioral orders $EX^\flat$ against the profit margin (defined as a percentage of price); we do this for $\beta = 1, 5, 10, 20$. To compute the expected behavioral solution $EX^\flat$, we calculate the behavioral density using (4).

Now, we make some observations using Figure 1. Notice that in all six plots, along the x-axis, there is some profit margin level (call it PML) where the curves corresponding to different values of β approximately intersect. For profit margins below PML, the expected order quantities $EX^\flat$ tend to increase as the bounded rationality parameter β increases; in contrast, the reverse is true for profit margins above PML. This suggests that with triangular demand distributions, bounded rationality leads to an increase in order quantities under low margin conditions, but it leads to a decrease in order quantities under high margin conditions. This is reminiscent of results with the flavor of "regression to the mean" except that here, bounded rationality is pushing order quantities towards the mid-point m of the range of possible demand realizations (namely, $m = 50$) instead of the mean, which is $(100 + h)/3$. Another observation is that the "threshold" profit margin level PML, which distinguishes low-margin conditions from high-margin conditions, decreases as the peak density h increases. In our numerical example, as h increases from 0 to 100, we see that PML decreases from 0.75 to 0.25 (approximately).

4.3 General Demand

When the newsvendor faces a general demand distribution, the behavioral solution can not be expressed in terms of explicit distributions such as the (truncated) normal. Nevertheless, we shall show that it is possible to characterize the expected order quantities and expected profits in the general case.

The key observation is that the behavioral solution $X^\flat$ belongs to an exponential family of probability distributions, parameterized by the price p and cost c.

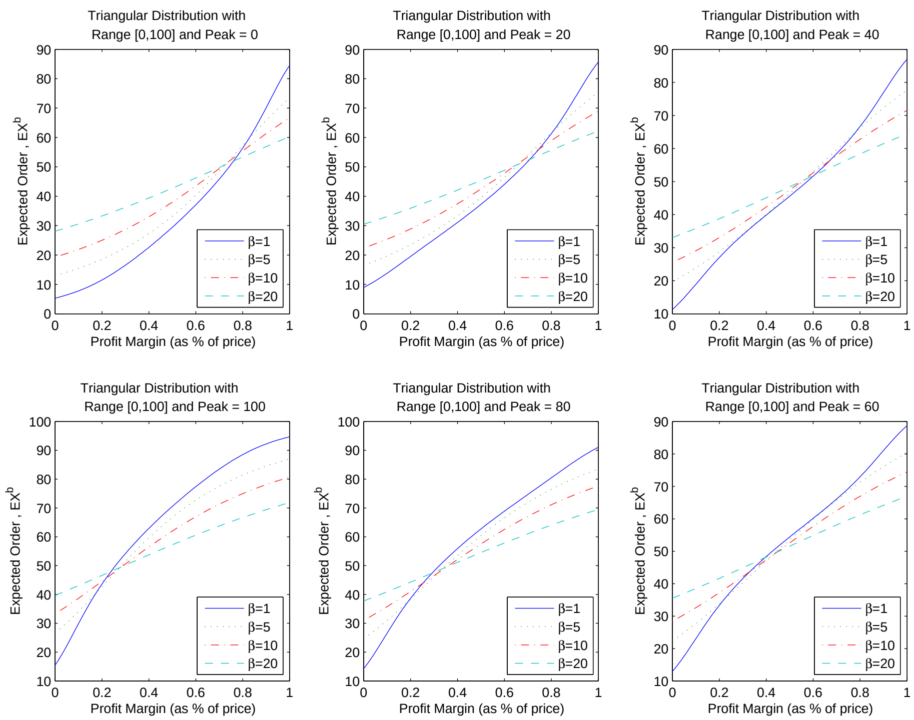

Figure 1: Expected order quantities for triangular demand distributions with range $[0, 100]$ and peaks at $h = 0, 20, 40, 60, 80, 100$. The price is fixed at \$1. In each case, we plot expected orders $EX^\flat$ against profit margins (as a fraction of price) for $\beta = 1, 5, 10, 20$.

Definition. *A family $\{H_\eta\}$ of probability distributions is said to form an s-dimensional exponential family if these distributions have densities of the form*

$$h(x; \eta) = l(x) \exp\left\{ \sum_{i=1}^{s} \eta_i T_i(x) - A(\eta) \right\}. \tag{10}$$

The parameters η are referred to as the natural parameters.

Observation. *The family of choice distributions $\{\Psi_{p,c}\}$, where $\Psi_{p,c}$ has density*

$$\psi_{p,c}(x) = \frac{e^{(pE\min(D,x)-cx)/\beta}}{\int_S e^{(pE\min(D,v)-cv)/\beta} dv}, \tag{11}$$

forms a 2-dimensional exponential family, with natural parameters $\eta_1 = p/\beta, \eta_2 = c/\beta, T_1(x) = E\min(D, x), T_2(x) = -x, l(x) \equiv 1,$ and

$$A(\eta_1, \eta_2) = \ln\left(\int_S e^{(\eta_1 E\min(D,v) - \eta_2 v)} dv \right). \tag{12}$$

We proceed to state a well-known fact about exponential families.

Fact. *Let X be from an exponential family with density (10). Then,*

$$E(T_i(X)) = \frac{\partial}{\partial \eta_i} A(\eta). \tag{13}$$

Based on this fact, it is then straightforward to write down the following result.

Proposition 2. *Let $X^\flat$ be the behavioral solution to the newsvendor problem with price p and cost c, so the natural parameters for the choice distribution are $\eta_1 = p/\beta$ and $\eta_2 = c/\beta$. Then, we have*

$$EX^\flat = -\frac{\partial A}{\partial \eta_2}(\eta_1, \eta_2), \tag{14}$$

$$E\pi(X^\flat) = p\frac{\partial A}{\partial \eta_1}(\eta_1, \eta_2) + c\frac{\partial A}{\partial \eta_2}(\eta_1, \eta_2). \tag{15}$$

This result highlights the important role played by the function $A(\eta_1, \eta_2)$ in our model of bounded rationality based on logit probabilities. Through this function, we can calculate moments of interest, such as expected order quantities and expected profits. This approach is useful for empirical and numerical studies in two ways. First, it allows us to replace integration (expectation) with differentiation, which is easier to compute. Instead of taking an expectation by integrating the density function (11), we can simply differentiate $A(\eta_1, \eta_2)$ and approximate its value using two data points. Second, it offers a way to estimate the value of β using aggregate data via the method of moments. For example, given laboratory data on average order quantities, an estimate of β would be the value at which the partial derivative $-\partial A/\partial \eta_2$, evaluated at $\eta_1 = p/\beta, \eta_2 = c/\beta$, matches the observation.

5. Empirical Evidence for Bounded Rationality

The goal of this section is to provide empirical evidence for our model of bounded rationality, using a dataset of newsvendor-type decisions made by individual subjects. In particular, we specify a statistical model for newsvendor decisions, and we fit our model to the data to obtain maximum-likelihood estimates of the bounded rationality parameter β. Finally, we show that our fitted model explains the data significantly better than the alternative that does not take bounded rationality into account.

First, we describe the data-set. This data-set consists of a series of newsvendor ordering decisions made by human subjects. Each subject participated either in the low-profit or high-profit condition, and made a sequence of 100 ordering decisions for the same parameter values. For the high-profit condition, demand is uniform between 1 and 100, price is 12, and cost is 3, so the optimal ordering quantity is 75. For the low-profit condition, demand is uniform between 51 and 150, price is 12, and cost is 9, so the optimal ordering quantity is again 75. There are 20 subjects participating in the low-profit condition and 18 subjects for the high-profit condition, so the data consists of 3800 quantity decisions altogether. Readers are referred to Bolton and Katok (2005) for more details on the experimental procedures used in collecting this data.

Our statistical model is

$$Y_k = X_k^\flat + \epsilon_k, \tag{16}$$

where Y_k is the observed order quantity for decision k, $X_k^\flat$ follows the same distribution as the behavioral solution $X^\flat$ obtained in Section 4, and ϵ_k are i.i.d. error terms. Since the demand is uniformly distributed, we know from Section 4.1 that the behavioral solution is truncated normal with mean at the optimal quantity $x^* = 75$ and standard deviation $\tau \equiv \sqrt{\beta \frac{b-a}{p}}$. We assume that

$$X_k^\flat \sim N^\dagger(x^*, \tau^2), \tag{17}$$

$$\epsilon_k \sim N(0, \sigma^2), \tag{18}$$

where $N^\dagger$ denotes the truncated normal distribution. There are two parameters τ and σ to be estimated. Given the data, the likelihood function is

$$L(\tau, \sigma | \mathbf{Y}) = \prod_{k=1}^{n} \int_L^U \phi\left(\frac{Y_k - m}{\sigma}\right) d\Psi(m), \tag{19}$$

where L and U denote the lowest and highest possible demand realizations, $\phi(\cdot)$ denotes the standard normal probability density function, and $\Psi(\cdot)$ denotes the probability distribution function of the behavioral solution $X^\flat$ as given in (17). We stress that the perfect rationality model, under which $\tau = 0$ and thus $X^\flat \equiv x^*$, is a special case of our model. In particular, to investigate whether the data suggests the presence of bounded rationality, we may test whether $\tau = 0$.

Our estimation strategy follows a data sub-sampling approach analogous to the bootstrap. We generate bootstrap samples from the dataset as follows. Let y_{ij} denote the j-th quantity decision made by subject i, where $i \in \{1, \ldots, I\}$ and $j \in \{1, \ldots, J\}$. Here $J = 100$ is the total number of decisions made by each subject and I is the number of subjects ($I = 18$ in the high-profit condition and $I = 20$ in the low-profit condition). Then, to generate each bootstrap sample $\{z_1, \ldots, z_I\}$, we randomly sample each z_i uniformly from $\{y_{ij} : 1 \leq j \leq J\}$. Each bootstrap sample is indicative of the quantity decisions made by our subject population. For each bootstrap sample (of size I), we obtain the parameter estimates $\hat{\tau}, \hat{\sigma}$ that maximize the likelihood function (19) above. We used a total of $B = 10,000$ bootstrap replicates. In other words, we obtain B estimates (one from each bootstrap replicate) of τ and σ in the model (16)-(18) above. Using the 2.5-th and 97.5-th percentile of these estimates, we obtain bootstrap confidence intervals for our parameters estimates $\hat{\tau}, \hat{\sigma}$. Our results are summarized in Table 2. Given these estimates, the log-likelihoods of our fitted model are -78.81 and -86.15 for the high-profit and low-profit conditions respectively.

High-profit condition:

	Maximum-likelihood estimate	95% Confidence intervals
$\hat{\tau}$	28.84	(20.92,38.79)
$\hat{\sigma}$	0.0448	(0.00001,0.1177)

Low-profit condition:

	Maximum-likelihood estimate	95% Confidence intervals
$\hat{\tau}$	25.31	(17.66,35.33)
$\hat{\sigma}$	0.0791	(0.00001,0.1392)

Table 2: Maximum-likelihood estimates of τ and σ under both low-profit and high-profit conditions.

Next, as a benchmark for comparison, we fit the data to the reduced model with $\tau = 0$, which corresponds to perfect rationality. Using the same bootstrap samples generated above, we can obtain the maximum-likelihood estimate $\hat{\sigma}$ and then use it to compute the likelihood of our fitted model. In Table 3, we report the log-likelihood values of our full model (above) and the reduced model here, for both low-profit and high-profit conditions.

	High-profit condition	Low-profit condition
Log-likelihood of full model, l_{full}	-78.81	-86.15
Log-likelihood of reduced model, $l_{reduced}$	-81.92	-88.97
Difference, $l_{full} - l_{reduced}$	3.11	2.82

Table 3: Log-likelihoods of full and reduced models under both low-profit and high-profit conditions.

In terms of fit, we are interested in how well our model (16) performs compared to the reduced model with $\tau = 0$. Since our full model has one additional parameter, it naturally performs better, so we need to penalize the additional degree of freedom in some way. One common criterion for model selection is the Bayes Information Criterion (BIC)

$$BIC = l(\hat{\theta}) - \frac{d \log(n)}{2}, \tag{20}$$

where $l(\hat{\theta})$ is the log likelihood of the fitted model, $\hat{\theta}$ are the fitted parameters, d is the number of parameters, and n is the data set size. Compared to alternatives such as the Akaike Information Criterion (AIC) and Mallow's C_p criterion, the BIC is a relatively conservative model selection criterion that favors simpler models. In the present context, the BIC of our full model is

$$BIC_{full} = l(\hat{\tau}, \hat{\sigma}) - \frac{d \log(n)}{2}, \tag{21}$$

while the BIC of our reduced model is

$$BIC_{reduced} = l(\hat{\sigma}) - \frac{d \log(n)}{2}. \tag{22}$$

These values are reported in Table 4. Since $BIC_{full} > BIC_{reduced}$, this criterion, despite its conservative nature, chooses the full model over the reduced model, suggesting that our model with bounded rationality is preferred over the alternative model with perfect rationality.

	High-profit condition	Low-profit condition
BIC of full model	-81.70	-89.15
BIC of reduced model	-83.37	-90.47

Table 4: Bayesian Information Criterion of full and reduced models under both low-profit and high-profit conditions.

As a final test, we consider the likelihood ratio test. Here we wish to test the hypothesis that $\tau = 0$. Since the log likelihoods of the full and reduced models, l_{full} and $l_{reduced}$, are given in Table 3, we can directly compute the test statistic $\chi^2 = 2(l_{full} - l_{reduced})$. This yields $\chi^2 = 6.22$ and $\chi^2 = 5.64$ for the high-profit and low-profit conditions. Under the assumptions of our full and reduced models, the test statistic follows a χ^2-distribution with one degree of freedom, which has a critical value $\chi_1^2(0.95) = 3.84$. Since our test statistics exceed the critical value, we reject the hypothesis that $\tau = 0$. In other words, this suggests that there is significant evidence for bounded rationality (with $\beta > 0$) in our model.

6. Distortion in order quantities

In this section, we investigate the effect of bounded rationality on expected order quantities. How does the behavioral mean order quantity $EX^\flat$ differ from the optimal solution x^*?

We would like to distinguish the situations in which the boundedly rational newsvendor over-orders (i.e. when $EX^\flat < x^*$) from the situations in which he under-orders (i.e. when $EX^\flat > x^*$). We identify two underlying effects that may cause such distortions.

First, when the newsvendor is boundedly rational, order quantities tend to be biased toward the midpoint of the range of possible demand realizations. We call this the *midpoint bias*. The anchoring heuristic (see Tversky and Kahnemann, 1974) provides a behavioral rationale for this effect. When the demand density f has support $[a, b]$, so that a and b are the smallest and largest possible demand realizations, the midpoint is $m \equiv (a + b)/2$. The following result establishes this bias for the special case of uniform demand.

Proposition 3. *Suppose that the demand density f is constant over $[a, b]$. Then, there is under-ordering when $x^* > m$ and over-ordering when $x^* < m$.*

There is an equivalent way to describe this result. Let us say that the newsvendor's product is a *high-profit* product if the critical fractile $\xi \equiv c/p < 0.5$ and a *low-profit* product if the critical fractile $\xi \equiv c/p > 0.5$. Equivalently, $p > 2c$ for a high-profit product and $p < 2c$ for a low-profit product. Then, for uniformly distributed demand, there is under-ordering for high-profit products and over-ordering for low-profit products. This terminology has been introduced by Schweitzer and Cachon (2000), who also provide experimental evidence for this result. Specifically, in their study, demand was uniformly distributed between 0 and 300 and $p = 12$. For the low-profit condition, $c = 9$ (i.e. the critical fractile $\xi = 75\%$) and for the high-profit condition, $c = 3$ (i.e. the critical fractile $\xi = 25\%$). Using data from 33 subjects, each making 15 newsvendor decisions for each condition, they found that in the high-profit condition, the average order was significantly lower than the optimal order, and in the low-profit condition, the average order was significantly higher than the optimal order. Schweitzer and Cachon considered many explanations for their observations (such as risk aversion, loss aversion, waste aversion and stockout aversion), and identified two consistent explanations: preferences to reduce ex-post inventory error, and the anchoring and insufficient adjustment bias. Here, we show that under bounded rationality, the midpoint bias provides a possible alternative explanation.

Next, we describe the second decision bias. Consider the case where the demand density is monotone. An increasing density suggests that demand is more likely to be high, whereas a decreasing density suggests that demand is more likely to be low. Our next result shows that the behavioral solution is distorted in the direction of low-probability demand realizations.

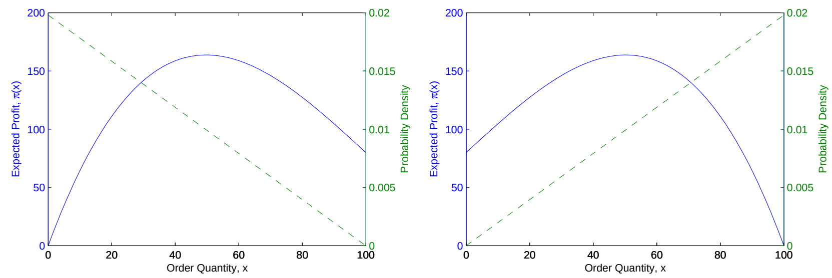

Figure 2: Newsvendor profit functions (solid lines) for linear probability densities (dashed lines) over [0,100]. On the left, $p = 10, c = 2.5$, so the optimal quantity is 50. On the right, $p = 10, c = 7.5$, so the optimal quantity is again 50.

We call this the *rare-event bias*. One potential behavioral explanation is that decision-makers tend to place excessive weight on rare occurrences (as in the probability weighing function of cumulative prospect theory in Tversky and Kahneman, 1992). In order to control for midpoint bias described earlier, we assume that the optimal solution occurs precisely at the midpoint m.

Proposition 4. *Suppose that $x^* = m$. Then, there is over-ordering when f is decreasing over $[a, b]$ and there is under-ordering when f is increasing over $[a, b]$.*

This result becomes quite intuitive when one visualizes the shape of the profit function $\pi(x)$. Let us consider the two examples in Figure 2, with linearly increasing or decreasing demand densities. Given that the optimal solution x^* occurs at the mid-point m, when f is decreasing (as in the left panel), profits fall from the optimal level faster when x is decreased below x^* compared to when x is increased above x^*. This implies that over-ordering is less costly compared to under-ordering, and hence occurs more frequently. The combined effect is that the expected behavioral order $EX^\flat$ exceeds the optimal quantity x^*, so there is over-ordering on average. Similarly, the reverse is true when f is increasing (as in the right panel). Notice that the rare-event bias may distort order quantities *away* from both the mean and the median. For example, in Figure 2, when the demand density is decreasing (on the left panel), the mean and median demands are both smaller than m, but the expected order quantity $EX^\flat$ is greater than m. This suggests that results of "regression toward the mean" (typically for uniform demand settings) do not capture the complete picture.

In the analysis above, we have made assumptions to isolate the two decision biases from each other. In Proposition 3, to control for the rare-event bias, we focus on uniform demand with all realizations equally likely. In Proposition 4, to suppress the midpoint bias, we assume that the optimal solution is located at the mid-point. However, in most instances of

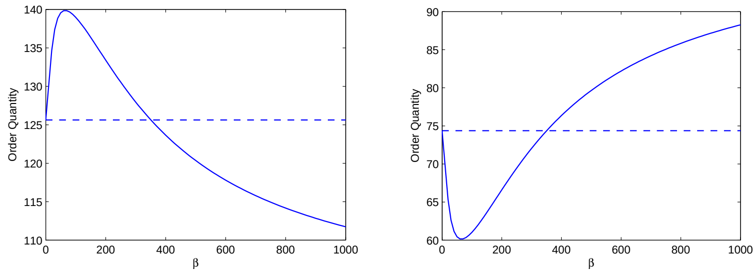

Figure 3: Behavioral (solid) and optimal (dashed) solutions for truncated normal demand with $\mu = 100, \sigma = 20$ and support $[0, 200]$, $p = 10$, and $c = 1$ (left), $c = 9$ (right).

the newsvendor problem, both effects are present and they may run in opposite directions.

Next, we provide a numerical example to illustrate the combined effect of both decision biases. In this example, we assume that the demand follows a truncated normal distribution with $\mu = 100$, $\sigma = 20$ and support $[0, 200]$. We assume that $p = 10$ and separately consider cases with $c = 1$ and $c = 9$. The optimal solution x^* is obtained using the critical fractile and the mean behavioral solution $EX^\flat$ is computed using the procedure outlined in Proposition 2. The results are summarized in Figure 3, with each panel corresponding to each value of c. In each panel, the solid lines represent the behavioral expected orders $EX^\flat$ and the dashed lines represent the optimal order quantities x^*. When $c = 1$, the optimal order quantity x^* exceeds μ, so the midpoint bias distorts decisions downwards; however, since the normal density is decreasing at x^*, the rare-event bias distorts decisions upwards. When these two effects are put together, we observe over-ordering for small values of β and under-ordering for larger values of β. This suggests that the rare-event bias is dominant for small values of β while the mid-point bias is dominant for large values of β. When $c = 9$, the directions of both decision biases are reversed. Nevertheless, a similar argument leads to the same conclusion: as β increases, decisions are first governed by the rare-event bias and then the midpoint bias takes over at larger values of β. The reason becomes clear when one recalls that as $\beta \to \infty$, the choice distribution becomes uniform over the choice domain, and the expected behavioral order $EX^\flat$ coincides with the midpoint m in the limit. It is therefore not surprising to find that the midpoint bias dominates for large values of β.

7. Supply Chain Coordination

The newsvendor model is an indispensable building block in the operations literature on supply chain coordination and contracting. This literature recognizes the problems caused by decentralized decision-making in a supply chain: when different parties act according to their own interests, system-wide optimal performance can not be attained. The general paradigm to solve these double marginalization problems is to design coordinating contracts that align the incentives of individual parties with the objectives of the entire supply chain. In this section, we will see that in the presence of bounded rationality, supply chain coordination can no longer be achieved in this manner.

The basic model of a decentralized supply chain consists of a single manufacturer and a single retailer. The manufacturer produces the good at unit cost c and sells it to the retailer at the wholesale price $w \geq c$. The retailer decides how many units x to procure before selling to the market at price p. His profit function is $\pi_R(x) = pE\min(D, x) - wx$, while total profits for the supply chain is given by $\pi_S(x) = pE\min(D, x) - cx$. It is well-known that the order quantity x_S^* that maximizes total supply chain profits satisfies $\bar{F}(x_S^*) = c/p$ but the retailer's profit-maximizing order quantity x_R^* satisfies $\bar{F}(x_R^*) = w/p \geq c/p$. Therefore, the retailer orders too little. Although standard theory concludes that system profits are lost, under bounded rationality, it is possible to construct examples in which double marginalization is beneficial. (Intuitively, this may occur when the centralized seller systematically over-orders and double marginalization helps to correct for these errors.)

How can the supply chain optimum be achieved? The literature has proposed many alternative solutions. First, consider the buy-back contract. Under this contract, the manufacturer agrees to buy back unsold units at the buy-back price b. In this case, the retailer faces the profit function $\pi_R(x) = (p - b)E\min(D, x) - (w - b)x$. When the contractual parameters w, b are chosen such that for some $\lambda \in [0, 1]$, $b = (1 - \lambda)p$ and $w = (1 - \lambda)p + \lambda c$, the retailer's profit function becomes $\pi_R(x) = \lambda[pE\min(D, x) - c] = \lambda\pi_S(x)$, which is a constant fraction of total supply chain profits. Therefore, by maximizing his own profits, the retailer is also maximizing supply chain profits because his payoff is always a constant fraction λ of total profits. In this way, the social optimum can be attained even though the ordering decision is made by the retailer considering only his own profits. Another possible alternative is the revenue sharing contract. Under this contract, the retailer agrees to share his sales revenue with the manufacturer; specifically, he pays the manufacturer r for every unit sold. Then, the retailer's profit function becomes $\pi_R(x) = (p - r)E\min(D, x) - wx$, which reduces to $\pi_R(x) = \lambda[pE\min(D, x) - c] = \lambda\pi_S(x)$ when we choose $r = (1 - \lambda)p$ and $w = \lambda c$. As before, the retailer now enjoys a fixed λ share of total system profits. The supply chain is thus similarly coordinated.

As we have seen, one general approach in supply chain coordination is to align individual incentives with social objectives by devising contractual transfers so that the decision-maker's payoff function $\pi_R(x)$ is a constant proportion λ of the social welfare function $\pi_S(x)$. While this approach is valid under perfect rationality, the next result shows that it does not always achieve coordination under bounded rationality.

To be precise, under boundedly rational decision-making, we say that the system is "coordinated" when decentralized control yields centralized profit levels. In other words, verifying system coordination involves comparing: (i) a decentralized system in which inventory decisions are made by a boundedly rational retailer, and (ii) a centralized system controlled by a boundedly rational newsvendor; further, the decision-makers in both cases share the same bounded rationality parameter β. When decentralized expected profits reach the same level as centralized profits, we say that system coordination is achieved.

Proposition 5. *Let $X_1^\flat$ and $X_2^\flat$ denote the behavioral decisions of two decision-makers with the same bounded rationality parameter β but facing different utility functions $u_1(x) = \lambda_1\pi(x)$ and $u_2(x) = \lambda_2\pi(x)$, where $\lambda_1 > \lambda_2$. Then, $E\pi(X_1^\flat) > E\pi(X_2^\flat)$.*

Corollary 2. *Consider a supply chain facing newsvendor profit function $\pi(x) = pE\min(D, x) - cx$. Let $X_i^\flat$ denote the behavioral solution when the decision-maker enjoys λ_i share of the total profits. Then, when $\lambda_1 > \lambda_2$, we have $E\pi(X_1^\flat) > E\pi(X_2^\flat)$.*

In a centralized system, the newsvendor's share of total profits is $\lambda = 1$. However, in a decentralized system, the decision-maker (retailer) enjoys only a reduced share of $\lambda < 1$. Assuming the same bounded rationality parameter β for both decision-makers, our result indicates that total expected profits are lower in the decentralized system. This shows even when incentives are perfectly aligned, coordination can not be attained when decision-makers are boundedly rational.

This theoretical result is consistent with recent experimental findings in the literature. Katok and Wu designed experiments on supply chain contracting and demonstrate that under contracts that are theoretically proven to coordinate the system, human subjects make ordering decisions that do not lead to perfect coordination. Their findings are based on buy-back contracts and revenue-sharing contracts. They not only discover that both contracts do not achieve coordination, but also identify systematic differences between these two contractual forms. Our current model of bounded rationality explains why coordination does not occur, but does not distinguish between different contractual forms that align incentives in the same way.

The novel finding herein is that it is not sufficient to align the *ratio* of marginal costs and benefits; the *actual* margins must be aligned to achieve coordination. In other words, for the newsvendor problem, it is important to align the actual overage and underage costs (two numbers) instead of simply aligning the retailer's critical fractile (one number). In particular, when the retailer receives a fixed share of total system profits (this aligns the ratio of margins but scales down both underage and overage costs), there is an efficiency loss. One possible explanation is that when payoffs are scaled down, the decreased stakes held by the decision-maker makes his choices more prone to errors and biases. This suggests that in order for an agent's decentralized decisions to coincide with the system's centralized decisions, the agent must be the sole stakeholder of the system. This points to the strategy of "selling the firm" to the agent, which can be implemented using a two-part tariff. Here, the manufacturer charges the retailer a fixed fee T and then sells to him at cost. In this case, the manufacturer's transfer payment T does not affect the retailer's decisions, which are thus made from the perspective of a sole owner. Therefore, two-part tariffs can coordinate the supply chain under our model of bounded rationality.

We may even go one step further and argue that in our model, there is potential for *super*-coordination. That is, expected profits may be higher in a decentralized system (under bounded rationality) relative to a centralized system (also under bounded rationality). By Proposition 5, this may occur if each ordering decision made by the retailer accounts for $\lambda > 1$ times of total system profits. This can be implemented in several ways. First, suppose that a manufacturer sells to a retailer who is constrained to place the same orders over n periods. Then, under some contract that gives the retailer λ share, his payoff from each decision is essentially $n\lambda$ times of the system's per period profits. Decentralization is thus beneficial as long as $n\lambda > 1$. Bolton and Katok (2005) experimentally demonstrates the benefits of having retailers commit to their standing orders over multiple periods. However, we stress that these benefits should be attributed to commitment (to placing the same orders several times) rather than decentralization per se, although the latter may facilitate such commitment.

Another way to achieve super-coordination is to use sales rebates. Suppose that the manufacturer charges a wholesale price w but offers a sales rebate of r per unit (so the retailer makes $p + r$ from each unit sold). For some $\lambda > 1$, let $r = (\lambda - 1)p$ and $w = \lambda c$. Let us also use a fixed transfer T to allocate some surplus to the manufacturer (so this is a two-part tariff coupled with a sales rebate). Then, the retailer's payoff function is $\lambda\pi_S(x) - T$.

Since the transfer T does not affect retailer decisions, and since $\lambda > 1$, expected profits are higher here than in the centralized system with profit function $\pi_S(x)$.

In general, increasing the monetary stakes associated with each decision epoch generates better decisions. From a behavioral standpoint, this is intuitive since people tend to allocate more cognitive and computational effort into more important tasks. In our model, notice that increasing λ achieves the same effect as decreasing β; essentially, a decision-maker facing increased stakes is akin to one who is "more rational." Since λ can not increase indefinitely (i.e., $\lambda \to \infty$), this suggests that all the scenarios discussed in this section can not match the ideal centralized benchmark with a perfectly rational decision-maker (with $\beta = 0$). Therefore, the reader should be cautious when interpreting our notion of super-coordination.

In summary, we have seen how bounded rationality can enrich the supply chain coordination framework. While the conventional normative approach predicts that perfect coordination is attained as long as the decision-maker's payoff function is a fixed λ share of total profits, our descriptive model of bounded rationality shows that the magnitude of λ also plays an important role. Specifically, when $\lambda < 1$, as in most cases studied in the literature, perfect coordination is not achieved. In contrast, contractual arrangements may give rise to individual decisions with $\lambda > 1$, under which there is super-coordination.

8. Bullwhip Effect

In this section, we discuss the relationship between bounded rationality and the bullwhip effect. The bullwhip effect refers to a commonly observed phenomenon in supply chains: the variance of orders tends to increase dramatically as we move upstream along the supply chain. In a seminal analysis, Lee et al. (1997) identify four different sources of the bullwhip effect: demand forecasting, inventory rationing due to supply constraints, order batching, and price fluctuations. They show that each of these factors can independently generate the bullwhip effect.

We shall analyze a model that eliminates the physical causes of the bullwhip effect that have been identified in the literature. In other words, in our setup, a standard normative analysis following the conventional paradigm of perfect rationality would not yield the bullwhip effect. However, we find that once bounded rationality is introduced, the bullwhip effect emerges. This suggests that apart from the physical causes of the bullwhip effect, which have been well-studied, there are also behavioral causes such as bounded rationality.

The model consists of a serial supply chain indexed by $i = 1, \ldots, n$. Market demand is fulfilled at stage $i = 1$, each member at stage i procures supply from the adjacent upstream member at stage $i+1$, and production is initiated at stage $i = n$. We assume that market demand D, with distribution F, is independent across periods. We consider an infinite-horizon, discrete-time model, with the following sequence of events in each time period. First, units previously shipped from the upstream neighbor are received (at stage $i = n$, units entered into production are completed). Let x_i denote the inventory position at stage i at this point. Second, demand D_i is realized at each stage i; this refers to market demand D at stage $i = 1$ and orders from downstream neighbors at stage $i = 2, \ldots, n$. Third, units are shipped off to fulfill demand, lowering the inventory position at stage i to $x_i - D_i$. If this is positive, leftover inventory is carried over to the next period, incurring a holding cost of h_i per unit. If this is negative, there is a back-ordering cost of b_i per unit. (Alternatively, we may assume that there is an alternative supply source, from which units can be borrowed at a cost of b_i per unit.) Finally, units are ordered from the upstream neighbor. We use O_i to denote the order submitted by member i to member $i+1$ at the end of the period.

Let us see how this model eliminates all the four physical causes of the bullwhip effect. First, there is no demand forecasting effects because market demand is i.i.d. across periods with a commonly known distribution. Second, there is no inventory rationing, which eliminates strategic gaming effects. Third, since there is no fixed cost, order batching is irrelevant. Finally, there is no price fluctuations in this model.

Under perfect rationality, this model can be solved using a standard newsvendor analysis. Consider the decision problem at stage $i = 1$ of choosing the inventory level x_i. The trade-off is between ordering too much (incurring excessive holding costs) and ordering too little (incurring excessive borrowing costs). The problem can be formulated as

$$\min_{x \in S} \quad hE[x-D]^+ + bE[x-D]^- \tag{23}$$

$$\Leftrightarrow \min_{x \in S} \quad (h+b)E[x-D]^+ - bE(x-D) \tag{24}$$

$$\Leftrightarrow \max_{x \in S} \quad -(h+b)E[x-D]^+ + bx \tag{25}$$

$$\Leftrightarrow \max_{x \in S} \quad (h+b)E\min(D,x) - hx, \tag{26}$$

where the subscript i has been omitted for brevity. This is a newsvendor problem with $p = h + b$ and $c = h$, so the optimal solution x^* satisfies $F(x^*) = b/(h+b)$. The optimal order-up-to level x^* represents a target inventory position that the decision-maker would like to maintain. At stage $i = 1$, for a given order-up-to level x_1^*, orders submitted at the end of each period is always equal to the demand realization in that time period, so that inventory would be brought back up to the target level x_1^* in the next period. This implies that demand faced by the adjacent upstream member $i = 2$ is equal to market demand, but lagged by one time period. This upstream member thus faces the same demand distribution, solves a similar newsvendor problem, and submits orders that are equal to market demand (but now lagged by two periods). Using an inductive argument, it then follows that the orders placed by each supply chain member are all lagged versions of the same market demand, so these orders follow the same distribution F. In other words, the variance of orders remains constant throughout the supply chain, and there is no bullwhip effect.

Now let us see what happens under bounded rationality. For simplicity, we begin with a two-echelon model ($n = 2$), although the analysis carries forward to the general case. We also assume that the decision domain $S = \mathbb{R}$ is stationary. This yields the following result.

Proposition 6. *Suppose that the decision-makers at stage $i = 1, 2$ are boundedly rational with parameter β_i. Then, there exists independent random variables ϵ_1, ϵ_2 such that*

$$O_1 =^d D + \epsilon_1, \tag{27}$$

$$O_2 =^d D + \epsilon_1 + \epsilon_2, \tag{28}$$

and $Var(\epsilon_i) = 0$ if and only if $\beta_i = 0$.

Corollary 3. *(i) $Var(O_i) = Var(D) + \sum_{j=1}^{i} \nu_j$, where $\nu_i = 0$ if and only if $\beta_i = 0$.*
(ii) $Var(O_i) \geq Var(O_{i-1})$, with equality if and only if $\beta_i = 0$.

This result demonstrates that the bullwhip effect can persist even when its physical causes have been removed. As long as some agents in the supply chain are boundedly rational (with $\beta_i > 0$), the variance of orders will strictly increase upstream along the supply chain.

There is ample experimental evidence in the literature showing that the bullwhip effect persists even when its physical causes have been controlled for. In a seminal study that predates Lee et al. (1997)'s taxonomy of the causes of the bullwhip effect, Sterman (1989) controls for three causes (inventory rationing, order batching, and price fluctuations), and leaves demand signal processing as a potential cause since subjects were not informed of the demand distribution. Subsequently, Croson and Donohue (2006)'s experiment also controls for the fourth cause by using a publicly known (uniform) demand distribution. In both studies, the bullwhip effect is observed. Furthermore, both studies identified the failure to account adequately for the supply line as an important cause of the bullwhip effect; that is, subjects do not fully account for orders that have been placed but have not yet arrived. The framework of analysis is due to Sterman (1989). Subjects are assumed to make ordering decisions using the anchor and adjustment heuristic (see Tversky and Kahnemann, 1974): they first anchor on the expected demand rate, and then adjust their order quantities to correct discrepancies between desired and actual stock, both on hand and in the supply line. In both studies, fitting experimental data to this heuristic decision rule yields parameter estimates that demonstrate underweighting of the supply line, which leads to overordering and system instability. In contrast to these studies, there is *no supply line* in our setup here. Orders placed at the end of each period will always arrive at the beginning of the next period. Yet, our analysis suggests that the bullwhip effect remains. What, then, is causing the bullwhip effect?

Our results provide an alternative behavioral explanation, and we demonstrate it analytically. We point out that the bullwhip effect is potentially a phenomenon that arises whenever individuals attempt to guard against and correct for the mistakes that *others* may make. The underlying mechanism is made transparent in Proposition 6. In the normative model, orders at every stage should follow the same distribution as the market demand, so $Var(O_i) = Var(D)$ for each i. Now, suppose that the decision-maker at $i = 1$ is boundedly rational but the one at $i = 2$ is perfectly rational, i.e. $\beta_1 > 0$ and $\beta_2 = 0$. Then, while it is clear that $Var(O_1) > Var(D)$, we also end up with $Var(O_2) > Var(D)$. The order variance at $i = 2$ has been increased from the normative benchmark of $Var(D)$ even though the decision-maker there is perfectly rational, and this increase arises solely because of the need to recognize the decision biases of downstream members. Of course, if the decision-maker at $i = 2$ is also boundedly rational, the order variance will be increased further. This in turn increases the burden of upstream members to account for his decision biases, and the variance increase is propagated upstream. The bullwhip effect thus arises.

There is experimental evidence supporting this explanation. Croson, Donohue, Katok, and Sterman (2006) conduct an experiment in which the demand is constant (four units per period) and commonly known, and the system begins in equilibrium. With this setup, it is optimal to order four units every period at each stage. Yet, in their experiment, this does not happen and the bullwhip effect is still observed. Post-experimental questionnaires suggest that although subjects realized what the optimal policy was, they were uncertain whether the other players understood it. Lack of trust in others' actions may generate suboptimal order quantities. Once initial deviations from optimal orders occur, the system is knocked into disequilibrium and the bullwhip effect eventually occurs. The authors call this "coordination risk." Our model of bounded rationality complements this work by offering a quantitative framework to model coordination risk. For example, applying our framework to their experimental setup with $D \equiv 4$, we can use Proposition 2 to iteratively compute the behavioral solutions $X_1^\flat, X_2^\flat, X_3^\flat, X_4^\flat$, using the order distribution at each stage as the demand distribution at the next upstream stage. Experimental results can then be used to estimate the values of β_i at each stage.

In summary, our model of bounded rationality complements the recent behavioral operations literature by showing that the bullwhip effect can arise even in the absence of its physical causes. Furthermore, even in the absence of a supply line (i.e. zero lead-time), our model still generates the bullwhip effect. This brings out another relevant behavioral phenomenon (apart from underweighting of the supply line): decision biases may not be errors in themselves, but rather, they are appropriate safeguards that are taken when one is not sufficiently confident in others' actions.

Nevertheless, the conclusion that "there are behavioral causes of the bullwhip effect" stops short of a more general treatise. One of the fundamental goals of any supply chain is to match supply with demand, and in this regard, orders generated within the system are driven by both demand and supply processes. Demand processes trigger order incidence, and corresponding supply processes trigger order fulfillment. Naturally, there is uncertainty in both demand and supply processes, and this generates variance amplification upstream, along the direction of order flow. Although the causes of demand uncertainty, particularly those pertinent to the bullwhip effect, have been well-studied, the causes of supply uncertainty have received less attention. Bounded rationality and its associated behavioral phenomena are one such cause because they introduce noise into an otherwise deterministic decision-making process. Similarly, other sources of supply uncertainty, such as random yield, have analogous interpretations. In general, our results suggest that, after controlling for the physical causes of demand uncertainty, the bullwhip effect may still persist as a result of supply uncertainty. Interestingly, recent experimental findings by Rong, Shen, and Snyder (2006) suggest that under supply disruptions, there may be a reverse bullwhip effect that causes variance amplification downstream rather than upstream.

9. Inventory Pooling

The goal of this section is to investigate the impact of bounded rationality on the benefits of inventory pooling. This refers to a commonly-used strategy in multi-location inventory problems. When demand occurs at different locations, instead of holding separate stocks for each source of demand, firms may alternatively pool inventory at some centralized location. Most of the literature on inventory pooling is based on the newsvendor model. The general conclusion is that centralization generates pooling economies that lead to lower costs (specifically, holding costs and backordering costs) for the system. We are interested in whether this result continues to hold when the decision-makers are boundedly rational.

We shall use the following setup in our analysis. This is similar to the setting used in Eppen (1979). Consider the following single-period multi-location newsvendor problem. There are n different sources of demand, each occurring at a different location. Let D_i be the demand at location i for $i = 1, \ldots, n$ and let F_i be its distribution. We assume that the demand vector follows a multivariate normal distribution; furthermore, let μ_i and σ_i^2 denote the mean and variance of D_i, and let σ_{ij}^2 and ρ_{ij} denote the covariance and correlation coefficient of D_i and D_j. First, we treat the decentralized case; that is, separate inventories are maintained at each location. In this case, the decision variables x_i are the quantities to hold on hand at each location i. There is a holding cost of h for each unit left unsold at the end of the period, and there is a backlogging penalty of b for each unit of demand that can not be fulfilled. In other words, the goal at each location i is to minimize the following cost function

$$\gamma_i(x_i) = hE[x_i - D_i]^+ + bE[x_i - D_i]^-. \tag{29}$$

This is equivalent to maximizing the following newsvendor profit function

$$\pi_i(x_i) = pE\min(D_i, x_i) - cx_i, \tag{30}$$

where $p \equiv h + b$ and $c \equiv h$. Clearly, the optimal solution x_i^* at each location satisfies $F_i(x_i^*) = 1 - (c/p)$. Next, we treat the centralized case. Here, instead of solving n separate newsvendor problems for the stocking levels x_i at each location, there is only one centralized stocking decision x to make. Let $D_T \equiv \sum_{i=1}^n D_i$ denote the total demand from all sources, and let F_T be its distribution. We use μ_T and σ_T^2 to denote the mean and variance of the total demand, so $\mu_T = \sum_{i=1}^n \mu_i$ and $\sigma_T^2 = \sum_{i,j=1}^n \sigma_{ij}^2$. This total demand is to be met from the same pool of inventory x_T. The goal is to minimize total cost

$$\gamma_T(x_T) = hE[x_T - D_T]^+ + bE[x_T - D_T]^-, \tag{31}$$

which is equivalent to maximizing the following newsvendor profit function

$$\pi_T(x_T) = pE\min(D_T, x_T) - cx_T. \tag{32}$$

As before, the optimal centralized solution satisfies $F_T(x_T^*) = 1 - (c/p)$. The benefits of inventory pooling can then be studied by comparing the optimal total costs in the two cases: $\sum_{i=1}^n \gamma_i(x_i^*)$ in the decentralized case and $\gamma_T(x_T^*)$ in the centralized case.

Under perfect rationality, when demands are normally distributed, it is well-known that optimal decentralized costs $\sum_{i=1}^n \gamma_i(x_i^*)$ and optimal centralized costs $\gamma_T(x_T^*)$ are respectively proportional to the sum of standard deviations $\sum_{i=1}^n \sigma_i$ and the standard deviation of total demand σ_T (with the same proportionality constant). Notice that $\sigma_T \leq \sum_{i=1}^n \sigma_i$, with equality if and only if all the demands are perfectly correlated (i.e. $\rho_{ij} = 1$ for all i, j). This implies the following two results. First, the total cost in a decentralized system is at least as high as that in a centralized system; in other words, inventory pooling saves costs. Second, these cost savings depend on the correlation between individual demands; in particular, there are no cost savings when all the demands are perfectly correlated with one another. These results are due to Eppen (1979).

We proceed to check whether these predictions are robust against bounded rationality. Let $X_i^\flat$ denote the behavioral solutions at each location i, and let $X_T^\flat$ denote the behavioral solution at the centralized location. The following result provides sufficient conditions for pooling benefits to persist.

Proposition 7. *Suppose that the following inequalities hold:*

(a) $\sigma_T \leq \sum_{i=1}^n \sigma_i$,

(b) $\sigma_T \geq \sigma_i$ *for every* $i = 1, \ldots, n$.

Then, we have $E\gamma_T(X_T^\flat) < \sum_{i=1}^n E\gamma_i(X_i^\flat)$. *In other words, inventory pooling leads to a strict reduction in total costs.*

Let us examine the two conditions of Proposition in greater detail. We already know that (a) always holds, with equality if and only if the demands at all locations are perfectly correlated with $\rho_{ij} = 1$. Next, condition (b) says that the standard deviation of aggregate demand is at least as large as the standard deviation in any one location, which is likely to hold in practical situations (unless substantial negative correlations exist between demand sources). Hence, under broad conditions, inventory pooling continues to generate cost savings.

In fact, Proposition 7 additionally shows that the reduction in total costs is strictly positive. For the extreme case with perfectly correlated demands, there is no cost reduction under perfect rationality. However, under bounded rationality, we observe strictly positive gains. This suggests that the benefits of inventory pooling extend beyond the physical benefit that is related to the reduction in variance resulting from summing separate random demands. On top of that, there may be *behavioral* benefits of inventory pooling.

To see the intuition behind the behavioral benefits of inventory pooling, let us consider the following example. Suppose that demand at each separate location is deterministic (i.e. $\sigma_i = 0$). In this case, it is trivially optimal to stock the quantity that matches demand exactly. Therefore, under perfect rationality, there is no difference between decentralization (with individual stocks $x_i^* = \mu_i^*$) and centralization (with total stock $x_T^* = \mu_T$) since zero cost is attained in both cases. In contrast, under our model of bounded rationality, decision errors create a disparity between these two cases. Under decentralization, decision errors $|X_i^\flat - \mu_i|$ at each location accumulate and separately contribute toward aggregate costs. However, under centralization, these decision errors may cancel out and the expected total costs are thus decreased. This example illustrates that in an environment with no demand uncertainty, inventory centralization helps by pooling decision errors across locations.

More generally, inventory pooling achieves two effects: it pools demand uncertainty as well as supply uncertainty. The former is well-understood and arises because of variance reduction. The latter, however, deserves elaboration. Bounded rationality (in particular, our model of decision noise) injects uncertainty into the supply process, and we have seen above that pooling reduces the aggregate impact of such decision errors. Similarly, inventory pooling serves to attenuate other sources of supply uncertainty, for instance, random yield, record inaccuracy, and processing errors. Therefore, it is no wonder that in Proposition 7, the benefits of pooling persist even when demands are perfectly correlated: these gains are the result of pooling supply (rather than demand) uncertainty.

10. Conclusion

The classic quantal choice paradigm posits that people do not make the best decision all the time, but they make good choices more often than worse ones. In this paper, we use this framework to capture bounded rationality and apply it to several newsvendor-type inventory settings. Our analysis generalizes existing results and reconciles them with empirical observations. This suggests that accounting for decision noise and optimization error is one possible way to enhance the predictive accuracy of theoretical models. We hope that our modeling approach serves to connect the theoretical and experimental literatures, and in so doing, stimulate future research on "behavioral theory" in operations management.

We conclude with some suggestions for future research. The first and most impending direction is experimental. While some of the findings herein have been experimentally validated by previous studies, many others remain untested. For example, studies of newsvendor decision biases under non-uniform demand would be a good starting point. The second direction is applications-oriented. There is a wide variety of situations in which decision-makers (whether they are firms, workers, or customers) may display some extent of bounded rationality. In these cases, how are the current theoretical results affected? This is a broad question that equally applies across different areas, such as operations strategy with boundedly rational firms, revenue management with boundedly rational customers, and staffing and human resource management with boundedly rational workers. Next, another research direction is to understand the fundamental behavioral mechanisms responsible for biases and errors. Specifically, how does the bounded rationality parameter β depend on the nature of the operational task (e.g., complexity and context)? What are the effects of learning on β? How does individual heterogeneity in β affect the aggregate? From an experimental viewpoint, how can β be manipulated? Finally, it is also worthwhile to study the combined effects of bounded rationality (decision noise) and other behavioral regularities. We believe that the quantal choice framework, being general yet parsimonious, is well-suited to complement other behavioral theories.

Appendix

Proof of Lemma 1 Since $\tilde{y} \equiv ay + b$, we have $d\tilde{y}/dy = a$. Therefore, for any $y_0, \tilde{y}_0$ satisfying $\tilde{y}_0 \equiv ay_0 + b$, we have

$$\int_{-\infty}^{\tilde{y}_0} e^{\tilde{u}(\tilde{y})/\beta} d\tilde{y} = a \int_{-\infty}^{y_0} e^{u(y)/\beta} dy, \tag{33}$$

which implies that

$$\tilde{\Psi}(\tilde{y}_0) = \frac{\int_{-\infty}^{\tilde{y}_0} e^{\tilde{u}(\tilde{y})/\beta} d\tilde{y}}{\int_{-\infty}^{\infty} e^{\tilde{u}(\tilde{y})/\beta} d\tilde{y}} = \frac{a \int_{-\infty}^{y_0} e^{u(y)/\beta} dy}{a \int_{-\infty}^{\infty} e^{u(y)/\beta} dy} = \frac{\int_{-\infty}^{y_0} e^{u(y)/\beta} dy}{\int_{-\infty}^{\infty} e^{u(y)/\beta} dy} = \Psi(y), \tag{34}$$

as desired. ■

Proof of Lemma 2

$$\tilde{\Psi}(y) = \frac{\int_{-\infty}^{y} e^{\tilde{u}(v)/\beta} dv}{\int_{-\infty}^{\infty} e^{\tilde{u}(v)/\beta} dv} = \frac{\int_{-\infty}^{y} e^{(u(v)+a)/\beta} dv}{\int_{-\infty}^{\infty} e^{(u(v)+a)/\beta} dv} = \frac{e^{a/\beta} \int_{-\infty}^{y} e^{u(v)/\beta} dv}{e^{a/\beta} \int_{-\infty}^{\infty} e^{u(v)/\beta} dv} = \frac{\int_{-\infty}^{y} e^{u(v)/\beta} dv}{\int_{-\infty}^{\infty} e^{u(v)/\beta} dv} = \Psi(y). \blacksquare \tag{35}$$

Proof of Proposition 1 The density of the behavioral solution, from (4) and (5), is given by

$$\psi(x) = \frac{e^{(Ax^2+Bx+C)/\beta}}{\int_a^b e^{(Av^2+Bv+C)/\beta} dv}. \tag{36}$$

The density $\zeta(x)$ of a truncated normal random variable over $[a, b]$ with mean μ and variance σ^2 is

$$\zeta(x) = \frac{e^{-\frac{(x-\mu)^2}{2\sigma^2}}}{\int_a^b e^{-\frac{(v-\mu)^2}{2\sigma^2}} dv}. \tag{37}$$

Therefore, over domain $[a, b]$, we have

$$\psi(x) \quad \propto \quad e^{(Ax^2+Bx)/\beta}, \tag{38}$$

$$\zeta(x) \quad \propto \quad e^{-\frac{1}{2\sigma^2}x^2 + \frac{\mu}{\sigma^2}x}, \tag{39}$$

which implies that the behavioral solution has the truncated normal distribution with parameters μ and σ^2 satisfying

$$A \quad \equiv \quad -\frac{p}{2(b-a)} = -\frac{1}{2\sigma^2}\beta, \tag{40}$$

$$B \quad \equiv \quad \frac{pb}{b-a} - c = \frac{\mu}{\sigma^2}\beta. \tag{41}$$

Solving these two equations yields the desired values of μ and σ^2. ■

Proof of Proposition 3 When the demand density is constant, we have uniformly distributed demand. Recall from Corollary 1 that we have

$$EX^\flat = \mu - \sigma \cdot \frac{\phi(\frac{b-\mu}{\sigma}) - \phi(\frac{a-\mu}{\sigma})}{\Phi(\frac{b-\mu}{\sigma}) - \Phi(\frac{a-\mu}{\sigma})}, \tag{42}$$

where $\mu = b - \frac{c}{p}(b-a)$, $\sigma^2 = \beta\frac{b-a}{p}$, and $\phi(\cdot)$ and $\Phi(\cdot)$ denote the standard normal density and distribution functions. When $x^* = \mu > m$, we have $|b-\mu| < |a-\mu|$, so $\phi(\frac{b-\mu}{\sigma}) > \phi(\frac{a-\mu}{\sigma})$, implying that $EX^\flat < \mu = x^*$, so there is under-ordering. The same argument shows over-ordering when $x^* < m$. ■

Proof of Proposition 4 We first consider the case where f is decreasing over $[a, b]$. Denote $r \equiv (b-a)/2$, so $b = m + r$ and $a = m - r$. Thus, for any $y \in (0, r]$, we have $f(m+y) < f(m-y)$. Now, the assumption $x^* = m$ implies that $c = p\bar{F}(m)$, since $\bar{F}(m) = \bar{F}(x^*) = c/p$ is satisfied at the optimal solution x^*. It then follows that for any $y \in (0, r]$,

$$
\begin{aligned}
& \pi(x^*+y) - \pi(x^*-y) && (43)\\
= & \int_{m-y}^{m+y} \pi'(v)dv = \int_{m-y}^{m+y} (p\bar{F}(v) - c)dv = p\left\{\int_{m-y}^{m+y} \bar{F}(v) - \bar{F}(m)dv\right\} && (44)\\
= & \ p\left\{\int_{m-y}^{m+y} F(m) - F(v)dv\right\} && (45)\\
= & \ p\left\{\int_{-y}^{y} F(m) - F(m+v)dv\right\} && (46)\\
= & \ p\left\{\int_{0}^{y} F(m) - F(m+v)dv + \int_{0}^{y} F(m) - F(m-v)dv\right\} && (47)\\
= & \ p\left\{\int_{0}^{y} \left[\int_{m+v}^{m} f(z)dz + \int_{m-v}^{m} f(z)dz\right] dv\right\} && (48)\\
= & \ p\left\{\int_{0}^{y} \left[\int_{0}^{v} -f(m+z)dz + \int_{0}^{v} f(m-z)dz\right] dv\right\} && (49)\\
= & \ p\int_{0}^{y}\int_{0}^{v} [f(m-z) - f(m+z)]dzdv > 0. && (50)
\end{aligned}
$$

Since the density of the behavioral solution $X^\flat$ satisfies $\psi(x) \propto e^{\pi(x)/\beta}$, which is increasing in $\pi(x)$, we have $\psi(x^*+y) > \psi(x^*-y)$ for every $y \in [0, r]$. Together with the fact that

$$
\begin{aligned}
EX^\flat &= \int_a^b x\psi(x)dx = \int_{x^*-r}^{x^*+r} v\psi(v)dv = \int_{-r}^{r} (x^*+v)\psi(x^*+v)dv && (51)\\
&= x^* + \int_{-r}^{r} v\psi(x^*+v)dv && (52)\\
&= x^* + \int_{0}^{r} v[\psi(x^*+v) - \psi(x^*-v)]dv, && (53)
\end{aligned}
$$

we conclude that $EX^\flat > x^*$, so there is over-ordering when f is decreasing over $[a, b]$. The case with increasing f is treated similarly. ■

Proof of Proposition 5 Let the density of $X_i^\flat$ be $\psi_i(x) = \frac{e^{u_i(x)/\beta}}{\int_S e^{u_i(v)/\beta}dv}$ for $i = 1, 2$. Let $K \equiv \frac{\int_S e^{u_1(v)/\beta}dv}{\int_S e^{u_2(v)/\beta}dv}$. Then, we have $\psi_1(x) > \psi_2(x)$ if and only if

$$
\begin{aligned}
& e^{(u_1(x)-u_2(x))/\beta} > K && (54)\\
\Leftrightarrow \quad & \pi(x) > \frac{\beta \ln K}{\lambda_1 - \lambda_2}. && (55)
\end{aligned}
$$

Therefore, for any $k \geq \frac{\beta \ln K}{\lambda_1 - \lambda_2}$, we have

$$P(\pi(X_1^\flat) \geq k) = \int_{\{v:\pi(v)\geq k\}} \psi_1(v)dv > \int_{\{v:\pi(v)\geq k\}} \psi_2(v)dv = P(\pi(X_2^\flat) \geq k). \quad (56)$$

Similarly, for any $k \leq \frac{\beta \ln K}{\lambda_1 - \lambda_2}$, we have

$$P(\pi(X_1^\flat) \geq k) = 1 - \int_{\{v:\pi(v)\leq k\}} \psi_1(v)dv > 1 - \int_{\{v:\pi(v)\leq k\}} \psi_2(v)dv = P(\pi(X_2^\flat) \geq k). \quad (57)$$

This shows that $\pi(X_1^\flat)$ stochastically dominates $\pi(X_2^\flat)$, so we have our result. ■

Proof of Proposition 6 Let $X_i^\flat$ and x_i^* denote the behavioral solutions and optimal solutions at stage $i = 1, 2$. Then, using time subscripts t to avoid ambiguity, we have

$$O_{1,t} = X_{1,t+1}^\flat - (X_{1,t}^\flat - D_t). \quad (58)$$

This is because after period t, the inventory position has been lowered from the previous target of $X_{1,t}^\flat$ by an amount equal to current demand D_t, and the order $O_{1,t}$ is placed to replenish inventory to the new target $X_{1,t+1}^\flat$. Note that although the targets $X_{1,t}^\flat$ and $X_{1,t+1}^\flat$ follow the same choice distribution, the actual realizations may differ across time periods. Now, let us define ϵ_1 as the difference between two independent realizations of the random variable $X_1^\flat$. Then, we have the first relation (27). Next, for stage $i = 2$, we similarly have

$$\begin{aligned} O_{2,t} &= X_{2,t+1}^\flat - (X_{2,t}^\flat - D_{2,t}) && (59) \\ &= X_{2,t+1}^\flat - (X_{2,t}^\flat - O_{1,t-1}) && (60) \\ &= (X_{2,t+1}^\flat - X_{2,t}^\flat) + (X_{1,t}^\flat - X_{1,t-1}^\flat) + D_{t-1}. && (61) \end{aligned}$$

Now, defining ϵ_2 as the difference between two independent realizations of the random variable $X_2^\flat$, we have (28). Finally, it is straightforward to show that

$$\beta_i = 0 \Leftrightarrow P(X_i^\flat = E(X_i^\flat)) = 1 \Leftrightarrow Var(X_i^\flat) = 0 \quad (62)$$

using the Chebyshev inequality. Therefore, it follows that $Var(\epsilon_i) = 0$ if and only if $\beta_i = 0$.
■

Proof of Proposition 7 In the proof, we shall maximize over profit functions $\pi_i(x_i)$ rather than minimize over cost functions $\gamma_i(x_i)$. For the stocking problem at each location i, consider the following transformation. Instead of choosing the stocking quantities x_i, we shall choose the *standardized* stocking quantities $z_i \equiv (x_i - \mu_i)/\sigma_i$. The profit function over z_i should satisfy $\tilde{\pi}_i(z_i) = \pi_i(x_i)$, so we have

$$\tilde{\pi}_i(z_i) = \sigma_i p E \min(Z, z_i) - \sigma_i c z_i + (p - c)\mu_i, \tag{63}$$

where Z is a standard normal random variable. Let $Z_i^\flat$ denote the behavioral solution of maximizing $\tilde{\pi}_i(z_i)$. By Lemma 1, we know that behavioral solutions are invariant against affine transformations of the decision domain, so $X_i^\flat = \mu_i + \sigma_i Z_i^\flat$ and $E\tilde{\pi}_i(Z_i^\flat) = E\pi_i(X_i^\flat)$. This justifies working in the standardized decision domain. Next, by Lemma 2, we know that behavioral solutions are not affected by translations in the utility function, so $Z_i^\flat$ is also the behavioral solution when maximizing z_i over the centered profit function

$$\tilde{\pi}_i^\circ(z_i) = \sigma_i p E \min(Z, z_i) - \sigma_i c z_i. \tag{64}$$

Therefore, at each location i, we shall use the standardized and centered profit function $\tilde{\pi}_i^\circ(z_i)$ to characterize the behavioral solution $Z_i^\flat$. This argument also applies to the centralized case: we can use the problem of maximizing

$$\tilde{\pi}_T^\circ(z_T) = \sigma_T p E \min(Z, z_T) - \sigma_T c z_T \tag{65}$$

over z_T to characterize the behavioral solution $Z_T^\flat$.

Let us define the following *canonical* newsvendor problem: choose z to maximize

$$\Pi(z) = pE(Z, z) - cz, \tag{66}$$

which we refer to as the canonical profit function. Let ϕ and Φ denote the standard normal density and distribution functions. Then, it is easy to see that the solution to the canonical newsvendor problem z^* satisfies $\Phi(z^*) = 1 - c/p$, and the optimal objective function satisfies $\Pi(z^*) = p\int_{-\infty}^{z^*} v\phi(v)dv \leq 0$. In other words, $\Pi(z) \leq 0$ for every $z \in \mathbb{R}$.

Now, observe that $\tilde{\pi}_i^\circ(z_i) = \sigma_i \Pi(z_i)$ and $\tilde{\pi}_T^\circ(z_T) = \sigma_T \Pi(z_T)$. Therefore, for each i, applying Proposition 5 and condition (b) yields the result that $E\Pi(Z_T^\flat) \geq E\Pi(Z_i^\flat)$, with equality holding if and only if (b) holds with equality.

Finally, we can put our conclusions together to write

$$\sum_{i=1}^{n} E\pi_i(X_i^\flat) = \sum_{i=1}^{n} E\tilde{\pi}_i(Z_i^\flat) = \sum_{i=1}^{n} E\tilde{\pi}_i^\circ(Z_i^\flat) + (p-c)\sum_{i=1}^{n}\mu_i \tag{67}$$

$$= \sum_{i=1}^{n} \sigma_i E\Pi(Z_i^\flat) + (p-c)\mu_T \tag{68}$$

$$\leq \sum_{i=1}^{n} \sigma_i E\Pi(Z_T^\flat) + (p-c)\mu_T \tag{69}$$

$$\leq \sigma_T E\Pi(Z_T^\flat) + (p-c)\mu_T \tag{70}$$

$$= E\tilde{\pi}_T^\circ(Z_T^\flat) + (p-c)\mu_T \quad = E\tilde{\pi}_T(Z_T^\flat) = E\pi_T(X_T^\flat). \tag{71}$$

Note that inequality (70) holds because of condition (a) and $\Pi(z) \leq 0$. Now, observe that (70) binds if and only if (a) holds with equality, and (69) binds if and only if (b) holds with equality. However, since $\sigma_i < \sum_{i=1}^{n}\sigma_i$, (a) and (b) can not both bind at the same time. Therefore, strict inequality must hold in our result. ■

Acknowledgments

The author is indebted to Elena Katok and Gary Bolton for sharing the data from their paper, Bolton and Katok (2005).

References

Anderson, S.P., A. de Palma, J.F. Thisse. 1992. Discrete choice theory of product differentiation. MIT Press, Cambridge, MA.

Anderson, S.P., J.K. Goeree, C.A. Holt. 2004. Noisy directional learning and the logit equilibrium. *Scandinavian Jour. Econ.* 106(3): 581-602.

Bendoly, E., K. Donohue, K. Schultz. 2006. Behavior in operations management: Assessing recent findings and revisiting old assumptions. *Jour. Oper. Mgt.* Forthcoming.

Benzion, U., Y. Cohen, R. Peled, T. Shavit. 2005. Decision-making and the newsvendor problem - an experimental study. Working paper.

Binmore, K., L. Samuelson. 1997. Muddling through: Noisy equilibrium selection. *Jour. Econ. Theory.* 74: 235-265.

Block, H.D., J. Marschak. 1960. Random ordering and stochastic theories of response. *Contributions to Probability and Statistics.* (I. Olkin, ed.) pp. 97-132. Stanford, CA.

Blume, L.E. 1993. The statistical mechanics of strategic interaction. *Games and Econ. Behavior.* 5: 387-424.

Bolton, G.E., E. Katok. 2005. Learning-by-doing in the newsvendor problem. *M&SOM.*

Bostian, J.A., C.A. Holt, A.M. Smith. 2007. The newsvendor "pull-to-center effect:" adaptive learning in a laboratory experiment. Working paper.

Cachon, G.P. 2003. Supply chain coordination with contracts. Steve Graves, Ton de Kok (eds.). *Handbooks in Operations Research and Management Science.*

Cachon, G.P., M. Lariviere. 2005. Supply chain coordination with revenue sharing: strengths and limitations. *Mgt. Sci.* 51(1): 30-44.

Chen, H.C., J.W. Friedman, J.F. Thisse. 1997. Boundedly rational Nash Equilibrium: A probabilistic choice approach. *Games and Econ. Behavior.* 18: 32-54.

Chen, J., D.D. Yao, S. Zheng. 2001. Optimal replenishment and rework with multiple unreliable supply sources. *Oper. Res.* 49(3): 430-443.

Conlisk, J. 1996. Why bounded rationality? *J. Econ. Lit.* 34(2): 669-700.

Croson, R., K. Donohue, E. Katok, J. Sterman. 2005. Order stability in supply chains: coordination risk and the role of coordination stock. Working paper.

Croson, R., K. Donohue. 2006. Behavioral causes of the bullwhip effect and the observed value of inventory information. *Mgmt. Sci.* 52(3): 323-336.

Dada, M., N.C. Petruzzi, L.B. Schwarz. 2006. A newsvendor's procurement problem when suppliers are unreliable. *Manufacturing & Service Oper. Mgmt.* Forthcoming.

Eppen, G.D. 1979. Effects of centralization on expected costs in a multi-location newsboy problem. *Mgmt. Sci.* 25(5): 498-501.

Geigerenzer, G., R. Selten. 2001. *Bounded rationality: the adaptive toolbox.* MIT Press, Cambridge, MA.

Gerchak, Y., D. Mossman. 1992. On the effect of demand randomness on inventories and cost. *Oper. Res.* 40(4): 804-807.

Gino, F., G. Pisano. 2006. Behavioral operations. Working paper.

Hofbauer, J., W. Sandholm. 2002. On the global convergence of stochastic fictitious play.

Econometrica. 70: 2265-2294.

Karlin, S. 1958. One stage inventory models with uncertainty. *Studies in the mathematical theory of inventory and production.* Arrow, Karlin, Scarf (eds.) Stanford University Press, Stanford, CA. 109-134.

Katok, E., D. Wu. 2005. Contracting in supply chains: a laboratory investigation. Working paper.

Keser, C., G. Paleologo. 2004. Experimental investigation of retailer-supplier contracts: the wholesale price contract. Working paper.

Kok, A.G., K.H. Shang. 2006. Inspection and replenishment policies for systems with inventory record inaccuracy. *Manufacturing & Service Oper. Mgmt.* Forthcoming.

Lee, H.L., V. Padmanabhan, S. Whang. 1997. Information distortion in a supply chain: the bullwhip effect. *Mgmt Sci.* 43(4): 546-558.

Lee, H.L., C.A. Yano. 1995. Lot sizing with random yields: A review. *Oper. Res.* 43(2): 311-334.

Loch, C.H., Y. Wu. 2007. Behavioral operations management. Working paper.

Luce, R.D. 1959. Individual choice behavior: a theoretical analysis. Wiley, New York, NY.

Lurie, N.H., J.M. Swaminathan. 2005. Is timely information always better? The effect of feedback frequency on performance and knowledge acquisition. Working paper.

McFadden, D. 1981. Econometric models of probabilistic choice. C.F. Manski and D. McFadden (eds.) *Structural Analysis of Discrete Data with Econometric Applications.* MIT Press, Cambridge, MA. 198-272.

McKelvey, R.D., T.R. Palfrey. 1995. Quantal response equilibria for normal form games. *Games and Econ. Behavior.* 10: 6-38.

Pasternack, B.A. 1985. Optimal pricing and return policies for perishable commodities. *Mktg. Sci.* 4(2): 166-176.

Porteus, E.L. 2002. Foundations of stochastic inventory theory. Stanford University Press, Stanford, CA.

Rong, Y., Z.J. Shen, L.V. Snyder. 2006. The impact of ordering behavior on order-quantity variability: A study of forward and reverse bullwhip effects. Working paper.

Rubinstein, A. 1998. Modeling bounded rationality. MIT Press, Cambrige, MA.

Schweitzer, M.E., G.P. Cachon. 2000. Decision bias in the newsvendor problem with a known demand distribution: experimental evidence. *Mgmt. Sci.* 46(3): 404-420.

Simon, H.A. 1955. A behavioral model of rational choice. *Quart. J. Econ.* 69(1): 99-118.

Simon, H.A. 1982. Models of bounded rationality. MIT Press, Cambridge, MA.

Sterman, J.D. 1989. Modeling managerial behavior: misperceptions of feedback in a dynamic decision making environment. *Mgmt. Sci.* 35(3): 321-339.

Taylor, T. 2002. Coordination under channel rebates with sales effort effect. *Mgmt. Sci.* 48(8): 992-1007.

Thurstone, L.L. 1927. A law of comparative judgment. *Psych. Rev.* 34: 273-286.

Tomlin, B., Y. Wang. 2005. On the value of mix flexibility and dual sourcing in unreliable newsvendor networks.*Manufacturing & Service Oper. Mgmt.* 7(1): 37-57.

Tsay, A. 1999. Quantity-flexibility contract and supplier-customer incentives. *Mgmt. Sci.* 45(10): 1339-58.

Tsay, A. 2001. Managing retail channel overstock: markdown money and return policies. *J. of Retailing.* 77. 457-492.

Tversky, A., D. Kahneman. 1974. Judgment under uncertainty: heuristics and biases. *Science.* 185: 1124-1131.

Tversky, A., D. Kahneman. 1992. Advances in prospect theory: cumulative representation of uncertainty. *J. Risk and Uncertainty.* 5(4): 297-323.

Wu, Y., C.H. Loch. 2007. Social preferences and supply chain performance: an experimental study. Working paper.

Young, P. 1993. The evolution of conventions. *Econometrica.* 61: 57-84.