NHH

INSTITUTT FOR FORETAKSØKONOMI

DEPARTMENT OF BUSINESS AND MANAGEMENT SCIENCE

FOR 41 2014

December 2014

Discussion paper

Probabilistic cost efficiency and bounded rationality in the newsvendor model

BY

Jan Ubøe, Jonas Andersson, Kurt Jörnsten, Jostein Lillestøl AND **Leif K. Sandal**

Norges Handelshøyskole

NORWEGIAN SCHOOL OF ECONOMICS

Probabilistic cost efficiency and bounded rationality in the newsvendor model

Jan Ubøe*, Jonas Andersson, Kurt Jörnsten, Jostein Lillestøl, and Leif Sandal
Norwegian School of Economics,
Helleveien 30, N-5045 Bergen, Norway.

December 12, 2014

Abstract

In this paper we establish a link between probabilistic cost efficiency and bounded rationality in the newsvendor model. This establishes a framework where bounded rationality can be examined rigorously by statistical methods. The paper offers a relatively deep theoretical analysis of underorders/overorders in the newsvendor model. The theory is supported by empirical findings from our analysis of empirical data from laboratory experiments. In particular, we observe that underorders are systematically larger than overorders, an issue that our theoretical model explains. From statistical tests we conclude that all variability in our data can be explained by probabilistic cost efficiency and risk aversion.

Keywords: Behavioral economics, experimental economics, bounded rationality, probabilistic cost efficiency.

*Corresponding author: jan.uboe@nhh.no, phone: 004755959978, fax: 004755959650

1 Introduction

Gino and Pisano (2008) argue that greater emphasis should be placed on behavioral aspects of operations management. At the time of their writing, the majority of the operations management literature was concerned with fully rational agents, and departures from rationality assumptions were largely ignored. Today, however, the field has matured; see Croson et al. (2012) for a review. The present paper aims to contribute to this interesting stream of literature.

Discrete choice models emerged in the 1970's, with the pioneering work of D. McFadden on random utility maximization, see McFadden (1974) and Train (2003). The theory has been applied with success within several different fields in economics, and has obvious relevance to newsvendor behavior. Retailers often base their decisions on partial and incomplete information leading to a certain type of randomness in ordering. Managers should seek to understand the nature of this randomness and use their knowledge to improve performance. It is hence of some surprise that this approach is largely ignored in the literature on the newsvendor problem. A notable exception is Su (2008).

Our paper combines the framework of experimental economics with theory of bounded rationality. In an experimental study, Becker-Peth et. al (2013), the participants were asked to suggest order quantities in a standard newsvendor model. They were fully informed in the sense that a unique optimal order could be inferred from their information, but they seldom if ever suggested this optimal quantity. It is thus interesting to ask why this happened, and to relate our observations to theory of probabilistic choice. This approach was initiated by Su (2008), who obtained important new insights into this connection. In our paper we establish a more streamlined approach where these insights can be reached with a minimum of effort. Our more efficient framework enables us to extend the analysis of Su (2008) in several new directions.

A perfectly rational newsvendor orders a quantity q that maximizes expected profit. Su (2008) considers boundedly rational agents and discusses several alternative lines of enquiry. In Su's (2008) study, less well-informed agents can choose any order quantity, and the probability for choosing the size of an order is defined in terms of a multinomial logit (MNL) model.

An MNL model can be derived in many ways, some of which are discussed and referenced in Su (2008). The most common derivation is probably that based on random utility theory described by Manski (1977). A basic formulation used by Su (2008) is the following: "all alternatives are candidates for selection, but more attractive alternatives are chosen with larger probability". This statement is a necessary consequence of the multinomial logit model. What appears to be less well known is that a modified version of this statement is in fact sufficient; that is, if more attractive states are chosen with larger probability, then the model must be a multinomial logit model, see Erlander (2010). Here a state is referring to an allocation of choices made by several agents, and a state is more attractive if it leads to a larger aggregate utility.

In our paper we will use the following definition of bounded rationality: *Agents are boundedly rational if and only if more attractive states are chosen with larger probability.* In the end our definition of bounded rationality will lead us to a MNL model, and it might hence appear that our definition makes no difference. There is, however, a very good reason for taking that particular definition as a starting point; we want to use our definition to formulate a statistical test of bounded rationality. Su (2008), too, wanted to test for this, but had to settle for a test of the parameter in his model. Formally there is nothing wrong with his test, but it does not answer our main question; "are the observations consistent with bounded rationality?"

To our knowledge there is only one relevant test discussed in the literature; Erlander's graphical test for probabilistic cost efficiency, see Erlander (2010). Erlander takes the verbal statement "more attractive states are chosen with larger probability" as his starting point, and formalizes that statement into a direct statistical test of his definition. The reader should note that a similar approach fails if we instead start out with a definition based on random utility theory. From aggregate data we can never confirm that agents maximize random utility. Many other models lead to the same functional form, and data may be perfectly replicated by a MNL model even when random utility fails. In many contexts this is not at all a problem, but it effectively excludes the type of analysis that we are discussing here. That partly explains why it was difficult for Su (2008) to formulate a suitable test, and that very few such tests have been discussed in the literature.

In the paper we will correct Erlander' test and extend the test to a version which is applicable to small samples. The test has a compelling diagnostic part we can use to check if the agents behave according to our definition. In the theoretical part of the paper we study the classical pull-to-center effect extending the analysis in Su (2008). In particular we prove that when agents behave according to probabilistic cast efficiency, we can expect that the amount of underordering is typically larger than the amount of overordering. The results in the theory section are supported by an empirical analysis where we use the data from Becker-Peth et. al (2013). We suggest a simple likelihood ratio test which seems stronger than Erlander's test. The combination of the two tests appears to be very well suited for data of this kind.

The IIA property (independence of irrelevant alternatives) is a much debated issue in discrete choice theory and there exist several ways of testing if the IIA property is a problem in data. Problems with IIA typically occur when identical alternatives are listed multiple times, e.g., if yellow buses are painted red or blue, this should not make bus a more likely alternative for transport (color is irrelevant). In our paper the agents choose how much to order. As the order quantity is an ordinal variable with no special attributes, problems with multiple listings cannot occur and the IIA property is not something we will need to address here.

The paper is organized as follows. In Section 2, we briefly review some of the most relevant literature. In Section 3, we review the theory of probabilistic cost efficiency and discuss Erlander's graphical test in detail. We explain that Erlander's formula for the confidence band is in fact always wrong, and work out the correct version. We extend the test to a version that allows for small samples and also propose a new and very simple alternative to Erlander's test. In Section 4 we enter into a relatively deep theoretical discussion of underordering/overordering in the newsvendor model. In particular we prove that under certain conditions we can expect a systematic skewness in underorders versus overorders. In Section 5, we analyze our experimental data and conclude that all the variation in our data can be explained by a combination of probabilistic cost efficiency and risk aversion. A technical summary of the paper is provided in Section 6. Finally, in Section 7, we offer some concluding remarks. To enhance the readability of the paper, the major part of the technical proofs has been placed in the appendix.

2 Literature review

In the single-period newsvendor model, a retailer wishes to order a quantity q from a manufacturer. Demand D is a random variable, and the retailer selects an order quantity q maximizing his expected profit. When the distribution of D is known, the problem of determining an optimal quantity is easily solved. The basic problem is very simple, but it appears to have endless variations. There is now a very large body of literature on such problems; for further reading, refer to the reviews by Cachón (2003) and Qin et al. (2011) and the numerous references therein.

The analysis conducted in this paper relates to three main streams of literature.

- Discussions of probabilistic cost efficiency
- Discussions of bounded rationality in economics
- Discussions of the use of laboratory experiments to build better operations management models

In this section we provide a brief review of some of the literature related to the discussion in our paper.

2.1 Probabilistic cost efficiency

The notion of probabilistic cost efficiency was introduced by Smith (1978). The theory has been expanded and improved in several publications by S. Erlander and T. Smith, and a comprehensive discussion is provided in the monograph by Erlander (2010). The basic approach is to formulate a framework in which agents can choose from a list of alternative actions. Each action is associated with a cardinal utility, which in our context is interpreted as the cost of the action, i.e., a negative utility. If we assume that a pattern with higher total utility is always more probable than one with lower total utility, the resulting model will be a multinomial logit model.

The theory is very versatile, and admits generalizations where actions are constrained by K linear restrictions on the form $\mathbf{A}\,\mathbf{P}^{\perp} = \mathbf{B}^{\perp}$. In this case, an assumption of probabilistic cost efficiency implies a model formulation of the form

$$\mathbf{P} = \exp[(u_1, u_2, \ldots, u_{K+1})\mathbf{A} + \beta\mathbf{U}] \tag{1}$$

Here, $\mathbf{P} = (p_1, \ldots, p_M)$ are the probabilities of choosing actions $1, \ldots, M$. $\mathbf{A}$ is an $(K+1) \times M$ matrix and $\mathbf{B} = (b_1, \ldots, b_M)$ is a vector specifying the constraints on $\mathbf{P}$. Actions have utilities $\mathbf{U} = (U_1, \ldots, U_M)$ and the numbers $(u_1, u_2, \ldots, u_{K+1})$ and $\beta \geq 0$ are all constants. In the special case where $K = 0$, the constraint $p_1 + \cdots + p_M = 1$ leads to the multinomial logit model. See Jörnsten and Ubøe (2010) for a discussion and applications of the general framework.

2.2 Bounded rationality

Etzioni (1986) argues that natural human behavior is nonrational, largely governed by emotions and inconsistent values. Rational behavior is artificial in the sense that it results from a definition of cost, and what we define as rational behavior is hence largely a consequence of our definition of cost. Without a distinct definition of cost, no behavior would be rational.

From the above line of reasoning it comes as no surprise that agents are boundedly rational in the sense that they do not always choose the optimal, that is, the least costly, alternative. Only in cases where agents are fully informed and have a definite and indisputable definition of cost can we expect to observe fully rational behavior. In all other cases, there is a nonzero probability of mistakes.

The literature on boundedly rational agents is huge. An excellent survey of many streams is Conlisk (1996). Conlisk (1996) discusses four reasons for incorporating bounded rationality.

- There is empirical evidence for bounded rationality
- Models of bounded rationality are useful
- The logic of unbounded rationality is sometimes flawed
- Suboptimal decisions incur less cost

These reasons are discussed in detail and are supported by a long list of references. A paper with an interesting relation to the fourth reason is one by Mattsson and Weibull (2002).

Mattsson and Weibull (2002) assume that agents have a set of deterministic preferences over a set of alternatives. Agents are fully rational, and can solve any relevant maximization problem. However, a higher probability of choosing an alternative requires more effort, and in their model the marginal disutility of always choosing the optimal alternative is assumed to be infinite. There is thus a situation with well-defined costs, but in which the optimal decision for any decision maker is nevertheless to choose positive mistake probabilities. The option of never making mistakes is simply too costly/time-consuming, and the resulting choice between alternatives is not deterministic. In this kind of setting, there is hence a rational bound on how rational the agents can be.

Matejka and McKay (2013) take the rational inattention approach (Sims (1998, 2003)) to model how information frictions influence the behavior of utility-maximizing agents. When agents have no a priori preferences, choices are distributed in accordance with a standard multinomial logit model. Choice probabilities are systematically shifted, however, under nonuniform priors. The basic idea of the rational inattention approach is that information is costly to acquire, a point of view shared by Mattsson and Weibull (2002).

2.3 Laboratory experiments

In a classical laboratory experiment, Schweitzer and Cachón (2000) observe that agents order too little in cases where the profit is high and too much when the profit is low, the so-called "pull-to-center" effect. They offer two alternative explanations for this. The first is that agents seek to minimize the absolute difference between realized demand and quantity ordered. The second is that the decision making is biased because of comparisons with previous situations that may not be relevant to the present situation. They explain this by three heuristics by which the agents adapt by anchoring to one quantity and adjusting toward another, for example, anchoring to the previous order and adjusting toward previously observed demand.

Bostian et al. (2008) investigate these heuristics through a laboratory experiment involving a learning model inspired by Camerer and Ho (1999) in which the agents adaptively learn from their ordering decision. Like Schweitzer and Cachón (2000), they observe orders that are too small in high-profit situations and too large when profit is low. Their comparison supports the learning model in terms of fit to data. In this model, the agents learn adaptively which orders yield high profits and which yield low profits.

Bolton et al. (2012) compare a group of students with a group of experienced managers and, like Schweitzer and Cachón (2000), observe that the subjects too often order too little in high-profit situations and too much in low-profit situations. They conclude that the managers do not use the information or task training any more efficiently than the students.

Wachtel and Dexter (2010) review studies that largely confirm the findings of Bolton et al. (2012). Their focus is on staffing of operating theaters at hospitals. The order in this context is the number of staff needed for a surgeon to perform the tasks required in an operating room efficiently. Random demand is the number of patients. They conclude that both voluntary students as well as operating theater managers systematically allocate too many staff members to surgeons who do not need them and too few to the ones that do. Because the students have no reason to take organizational aspects into account but still make the same systematic error as the managers, the authors argue that this is evidence of an innate psychological bias.

Another possible explanation for the ordering bias proposed by Su (2008) is that newsvendors simply make random errors in ordering. Bias would then occur because there is more room to err toward the mean than away from it. Kremer et al. (2010) investigate this idea by allowing one group of subjects to place an order in a standard newsvendor problem and another to participate in a game that is identical in probabilistic terms but presented as a pure lottery. They conclude that these results are inconsistent with the random error model, and that the explanation for this is that the ordering strategies for the newsvendor group are based on biased order-to-demand mapping.

Rudi and Drake (2013) introduce demand censoring in the context of the “pull-to-center”effect. Demand censoring, that is, a situation where subjects of the experiment cannot observe demand when it exceeds the order quantity, is shown to lead to lower order quantities. In the case of a high-profit situation, this magnifies the distance between optimal and observed order quantities. Conversely, for cases with low profit it reduces this distance.

These explanations for the “pull-to-center” effect, interesting and sensible as they are, cannot apply to our results. The reason for this is that our experimental data are obtained from subjects that place only one order for each given set of parameters. In addition, we observe skewness in the distribution of ordered quantities that is not predicted by the explanations mentioned above. We discuss these issues from a theoretical perspective in Section 4.

Becker-Peth et al. (2013) construct a 3-parameter behavior model assuming that people consider the upside and downside potential of their order decisions separately in line with the mental accounting arguments (Thaler 1999) resulting in two separate accounts, one for sales and one for leftovers. The different values associated with income from sales and returns (buyback) are modeled by multiplying the income from return with a parameter larger than 1 to higher the values of income from returns. The other parameters represent anchoring and chasing effects. We use the same experimental data and offer what we think is a simpler explanation, i.e., that the variation we see in the data can be explained by probabilistic cost efficiency combined with risk aversion.

3 Probabilistic cost efficiency and statistical testing

The basic idea of probabilistic cost efficiency can be described as follows. Assume that N agents choose between K alternatives with costs $c_1, \ldots, c_K$. Consider two independent random samples of the same size N, and let

$$z_k^{(i)} = \text{number of times alternative } k \text{ is chosen in sample } i \qquad k = 1, \ldots, K, i = 1, 2.$$

When a sample of length N is drawn, we assume that the probability of choosing alternative k is the same for each individual decision, and that all individual decisions are independent. A probability distribution $p = (p_1, \ldots, p_K)$ is *probabilistically cost-efficient* if and only if , for any sample size N and for any pair of samples

$$\sum_{k=1}^{K} c_k z_k^{(1)} \geq \sum_{k=1}^{K} c_k z_k^{(2)} \;\Rightarrow\; \prod_{k=1}^{K} p_k^{z_k^{(1)}} \leq \prod_{k=1}^{K} p_k^{z_k^{(2)}}. \tag{2}$$

That is, if a sample has greater total cost, it is always less probable. The interesting point here is that if a probability distribution satisfies (2) for any pair of samples of arbitrary length, the probability p must satisfy

$$p_k = \frac{e^{-\beta c_k}}{\sum_{j=1}^{K} e^{-\beta c_j}}, \tag{3}$$

where $\beta \geq 0$ is a constant; see Erlander (2010) Chapter 4. The constant β measures agents' sensitivity to utility. If β is very large, alternatives with maximum utility are chosen with probability 1 at the limit. If β is very small, utility does not matter, and alternatives are equally probable at the limit.

3.1 Erlander's graphical test of probabilistic cost efficiency

We now discuss Erlander's graphical test for probabilistic cost efficiency in detail. Assuming any multinomial distribution (cost-efficient or not), we define the likelihood function $\mathcal{L}(z)$ and the average cost $\overline{c}$ by

$$\mathcal{L}(z) = \prod_{k=1}^{K} p_k^{z_k} \qquad\qquad \overline{c} = \frac{1}{N}\sum_{k=1}^{N} z_k c_k. \tag{4}$$

If we replace p by its maximum likelihood estimate $\overline{p} = (z_1/N, \ldots, z_K/N)$, we obtain

$$\text{loglikelihood} = \mathcal{LL}(\overline{p}) = \sum_{k=1}^{K} z_k \log[z_k/N]. \tag{5}$$

where terms with $z_k = 0$ are ignored because of the continuity extension $\lim_{z\to 0^+} z \log[z] = 0$. We define bounded rationality by the statement: *Agents are boundedly rational if and only if more attractive states are chosen with larger probability.* Assume that we observe M empirical samples each with N agents, and make a plot of the log likelihood function against the observed cost in each sample. If agents largely behave according to our definition, we would expect to see a falling pattern like the one shown in Figure 1.

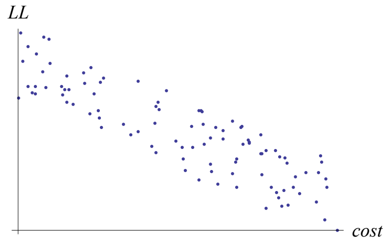

Figure 1: Log likelihood values as a function of $\bar{c}$

The rather informal diagnostic plot above is easily formalized into a rigorous statistical test, and Erlander (2010) suggests an explicit formula for the confidence band in that test. There is, however, a subtle oversight in Erlander's proof, and this error causes his formula to be wrong for any N. Erlander's proof for the confidence band progresses through a sequence of asymptotic approximations which are all correct until we reach the statement

$$-Ent(\overline{p}) \approx -Ent(p) + \beta c - \beta \bar{c} + \frac{1}{2N} H. \tag{6}$$

In this formula, Ent is the entropy and the random variable H is given by the expression $H = \sum_{k=1}^{K} (z_k - Np_k)^2 / Np_k$. Erlander assumes, correctly, that H is approximately χ^2 when N is large. In the formula for the confidence band, however, we need the distribution of H conditional on the event $\sum_{k=1}^{K} \frac{z_k}{N} c_k = \bar{c}$. If $\bar{c} \neq c = \sum_{k=1}^{K} c_k p_k$, we are conditioning on an event with probability zero in the limit. It is then not clear what happens, and more work is needed to compute the limit.

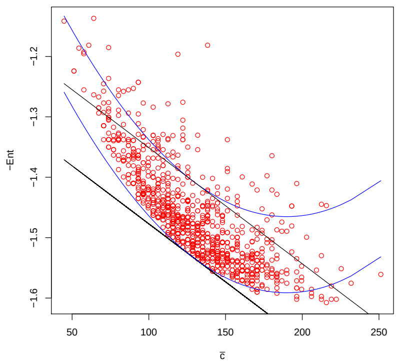

Figure 2: Confidence bands for Erlander's test. The curved lines show the correct band.

In Figure 2, we have drawn samples z of size $N = 31$ from the distribution in (3), assuming that $\beta = 0.002$ and that the costs are $(c_1, c_2, c_3, c_4, c_5) = (612.5, 312.5, 112.5, 12.5, 12.5)$. For each such sample, we have plotted the points $(\bar{c}, \frac{1}{N}\mathcal{LL}(z))$. Figure 2 shows the observed pattern from 1000 such samples. The straight lines are computed using Erlander's formula for a 95% confidence band, while the curved lines are using a correct formula. The problem is that far too many points lie above the upper straight line, and this problem do not vanish if we increase N.

The key to the correct confidence band is the following proposition:

Proposition 3.1
Let $z = (z_1, z_2, ..., z_K)$ be multinomial $(N, p_1, p_2, ..., p_K)$, and $\mathcal{X}^2 = \sum_{k=1}^K (z_k - Np_k)^2/(Np_k)$ be the common χ^2 expression. The distribution of $\mathcal{X}^2 - \rho_1^2$, conditionally on $\sum_{k=1}^K \frac{z_k}{N} c_k = \bar{c}$, is approximately (the error goes to zero as $N \to \infty$) χ^2 with $K-2$ degrees of freedom, where

$$\rho_1^2 = N(c-\bar{c})^2 / \sum_{k=1}^K (c_k - c)^2 p_k \qquad\qquad c = \sum_{k=1}^K c_k p_k.$$

Proof
The proof consists of two parts. The first and difficult part is to prove that the conditional asymptotic distribution of $\mathcal{X}^2$ minus a deterministic term is χ^2. The details are technical and are provided in the appendix. Once we know that the difference is deterministic, it is clear that the difference must equal the conditional minimum of the expression (the minimum of the χ^2 is zero). The details are straightforward and are omitted.

□

Since (6) in our notation is equivalent to to the statement

$$\sum_{k=1}^K \frac{z_k}{N} \log[\frac{z_k}{N}] \approx \sum_{k=1}^K p_k \log[p_k] + \beta c - \beta\bar{c} + \frac{1}{2N} \sum_{k=1}^K (z_k - Np_k)^2/(Np_k). \tag{7}$$

the following theorem follows directly from Proposition 3.1.

Theorem 3.2
Assume that N samples are drawn from a probabilistically cost-efficient distribution with costs $(c_1, c_2, \ldots, c_K)$ and parameter β. If $\bar{c} = \sum_{k=1}^K \frac{z_k}{N} c_k$ is the observed average cost and $c = \sum_{k=1}^K p_k c_k$ is the expected cost, there is an approximately (the error goes to zero as $N \to \infty$) 1-α percent probability that the observed log likelihood value $\sum_{k=1}^K \frac{z_k}{N} \log[\frac{z_k}{N}]$ is between

$$\text{Upper}[\bar{c}] = \sum_{k=1}^K p_k \log[p_k] + \beta c - \beta\bar{c} + \frac{(\bar{c}-c)^2}{2\sum_{k=1}^K p_k(c_k-c)^2} + \frac{1}{2N}\mathcal{X}_\alpha^2, \tag{8}$$

and

$$\text{Lower}[\bar{c}] = \sum_{k=1}^K p_k \log[p_k] + \beta c - \beta\bar{c} + \frac{(\bar{c}-c)^2}{2\sum_{k=1}^K p_k(c_k-c)^2}, \tag{9}$$

where $\mathcal{X}_\alpha^2$ denotes the $1-\alpha$ percentile in the χ^2 distribution with $K-2$ degrees of freedom.

Using the result in Theorem 3.2, we can easily establish a statistical test for probabilistic cost efficiency. The idea is to make a few independent observation sets with the same cost structure, and then, simply count how many of these have an observed log likelihood value exceeding the upper limit of the confidence band given by (8). The number of such sets can be relatively small as the test can be executed via resampling, see Section 3.2.

In Figure 2 we used the value $N = 31$. To check for accuracy, we constructed 100 000 independent samples (each sample with $N = 31$), and found that a total of $6,468$ samples had a log likelihood value exceeding the upper limit of the 95% confidence band given by (8). This shows that the approximation may be reasonably good even when N is moderate in size. In comparison, the linear bounds suggested by Erlander (2010) imply a total of $11,778$ points above the upper limit.

3.2 Resampling and an alternative test

In many cases we need to carry out the test based on a single observation. This is not necessarily a problem as the test can be carried out via resampling. If the resampling seed is constructed from several independent subsamples, the resampling error is small and the test can be executed via Theorem 3.2. If the number of observations is very small, however, the original data cannot be split into independent subsamples in a meaningful way. All resamples must then be drawn from the original seed. In this case the resampling error has the same order of magnitude as the original variation. As Theorem 3.2 is an asymptotic result, the expressions in (8) and (9) are subject to error when N is small. We have tested several cases numerically, and resampling together with the expressions (with Upper and Lower defined via (8) and (9))

$$\text{Adjusted upper}[\bar{c}] = \text{Upper}[\bar{c}] + \frac{\Delta}{2N}\mathcal{X}_\alpha^2 \tag{10}$$

$$\text{Adjusted lower}[\bar{c}] = \text{Lower}[\bar{c}] - \frac{\Delta}{2N}\mathcal{X}_\alpha^2 \tag{11}$$

can be used when N is small. The exact value of Δ depends on N and the parameters in the problem, and can be found by numerical simulation. The dependence is very slight, however, and we found that the approximation $\Delta \approx 1$ works well over a wide variety of cases. The downside with resampling from the original seed is that it makes the confidence band broader, leading to a rather weak test. To deal with this problem we suggest an alternative test. Our new test will be based on the equivalence:

$$\text{Probabilistic cost efficiency} \Leftrightarrow \text{Choices are drawn from an MNL model}$$

While it would be very artificial to use the MNL logit model as a definition of bounded rationality, we can stick with our original definition and use the equivalence above to formalize a statistical test. The reader should note that this idea fails in the random utility framework as

$$\text{Random utility theory} \not\Leftarrow \text{Choices are drawn from an MNL model}$$

Using the equivalence above, our new null hypothesis can be formulated as follows.

H_0: There exists a constant $\beta \geq 0$ such that $p_k = \frac{e^{-\beta c_k}}{\sum_{j=1}^K e^{-\beta c_j}}$.
H_A: The distribution is not of this kind.

The test is conducted as follows. We first find a value of $\hat{\beta}$ such that our model (under H_0) obtains the best possible fit in the sense of maximum likelihood. We define

$$\hat{p}_k^{(0)} = \frac{e^{-\hat{\beta} c_k}}{\sum_{j=1}^{K} e^{-\hat{\beta} c_j}}, \qquad \ln[L_0] = \sum_{k=1}^{K} y_k \ln[\hat{p}_k^{(0)}], \tag{12}$$

$$\hat{p}_k^{(1)} = \frac{y_k}{\sum_{i=1}^{K} y_i}, \qquad \ln[L_1] = \sum_{k=1}^{K} y_k \ln[\hat{p}_k^{(1)}]. \tag{13}$$

Here, y_k refers to the observed values. With these definitions, under H_0, the difference

$$\mathcal{X}^2 = 2(\ln[L_1] - \ln[L_0]) \tag{14}$$

is approximately χ^2 with $K-2$ degrees of freedom according to the standard maximum likelihood theorem.

While we would be the first to admit that this test is very simplistic, the test is rigorously supported by our original definition. We have never seen a test of this sort executed on data in our particular context, and as we will demonstrate in Section 5, it provides definitive answers to several empirical questions. The simplicity appears to be a strength not a weakness. We hence believe that our new test is an interesting alternative to Erlander's graphical test, which as remarked in the introduction, is the only relevant test previously discussed in the literature.

4 Boundedly rational agents in the newsvendor model.

We will now consider a setting where the agents choose how much to order in a newsvendor setting. The newsvendor model is specified as follows.

$W =$ wholesale price per unit (fixed)

$q =$ order quantity (rate chosen by the retailer)

$R =$ retail price per unit (fixed)

$D =$ demand (random rate)

$S =$ salvage price per unit (fixed)

A retailer is trading a commodity and orders q units from a manufacturer. He hopes to sell enough of these units to make a profit. We assume that the manufacturer offers a wholesale price W, and that the retail price R is exogenously given. Unsold items can be salvaged at the exogenously given salvage value S. A straightforward computation shows that the retailer maximizes expected profit when

$$P(D \leq q) = \frac{R-W}{R-S} \;\Rightarrow\; q_{\text{opt}} = F_D^{-1}\left[\frac{R-W}{R-S}\right], \tag{15}$$

where F_D denotes the cumulative distribution of D. A perfectly rational newsvendor will hence hence always order the quantity given by (15). Any deviation from the quantity given by (15) will incur a cost which is the loss in expected profit in comparison with the optimal choice. A boundedly rational newsvendor can order any quantity, but should have an inclination to make orders leading to small costs. If $\Pi(q)$ denotes the expected profit if the agent order q unit, the cost $c(q)$ is given by $c(q) = \Pi(q_{\text{opt}}) - \Pi(q)$. If we define bounded rationality in terms of probabilistic cost efficiency, our definition leads to a MNL model with density

$$\psi_Q(q) = \frac{e^{-\beta c(q)}}{\int_{d_{\min}}^{d_{\max}} e^{-\beta c(u)} du} \qquad \beta \geq 0. \tag{16}$$

As $\Pi(q_{\text{opt}})$ does not depend on q, it is easy to see that the particular value cancels in (16) and that we might just as well work with the expression

$$\psi_Q(q) = \frac{e^{\beta \Pi(q)}}{\int_{d_{\min}}^{d_{\max}} e^{\beta \Pi(u)} du} \qquad \beta \geq 0. \tag{17}$$

In this section there are different levels of randomness, and it is important to keep these apart. On the first level, we have randomness in the demand D. If we assume that D has a distribution with density $\phi_D(x)$ on the interval $[d_{\min}, d_{\max}]$, the expected profit in the newsvendor model is a function $\Pi = \Pi(q)$, and is given by

$$\Pi(q) = (R-S)\mathrm{E}_D[\min[D, q]] - (W-S)q, \tag{18}$$

Here, the subscript E_D is used to emphasize that we have this expectation at the first level. At the second level, the order quantity Q is a random variable with a multinomial logit density $\psi_Q(q)$ given by

$$\psi_Q(q) = \frac{e^{\beta \Pi(q)}}{\int_{d_{\min}}^{d_{\max}} e^{\beta \Pi(u)} du} \qquad \beta \geq 0. \tag{19}$$

We use the notation E_Q to emphasize that we have this expectation for this density.

When researchers design economic experiments of the kind discussed in this paper, the parameters in the experiment should be carefully selected to avoid bias. By an experimental design $\mathcal{ED}$, we mean the collection of parameters used in the experiment. To examine overall tendencies in our experiment, we average our results over all the cases. We use the notation $\mathrm{E}_{\mathcal{ED}}$ to signify the average value over all the experiment in our experimental design $\mathcal{ED}$, i.e., an expectation where each particular case has uniform weight.

4.1 Over/underordering

Su (2008) examined the sign of the error in ordering. In the case where D is uniform, he provided a rigorous proof that agents overorder if $q_{\text{opt}} < \frac{d_{\max}+d_{\min}}{2}$ (low profit case) and underorder if $q_{\text{opt}} > \frac{d_{\max}+d_{\min}}{2}$ (high profit case). This corresponds to the classical pull-to-center effect discussed by many authors. He also obtained some partial results for the non-uniform case. This analysis can, however, be compressed to only a few lines when it is done efficiently. The key is the following observation: By a linear change of variables using the density specified by (19), we see that

$$\mathrm{E}_Q[Q - q_{\text{opt}}] = \int_{d_{\min}}^{d_{\max}} (q - q_{\text{opt}})\psi_Q(q)dq = \frac{\int_{d_{\min}-q_{\text{opt}}}^{d_{\max}-q_{\text{opt}}} q e^{\beta\Pi(q_{\text{opt}}+q)} dq}{\int_{d_{\min}-q_{\text{opt}}}^{d_{\max}-q_{\text{opt}}} e^{\beta\Pi(q_{\text{opt}}+u)} du}. \tag{20}$$

Proposition 4.1
Assume that D has arbitrary distribution, and let $\Pi(q)$ be the expected profit when the retailer orders q units.

- *If for all $q \in [d_{\min} - q_{\text{opt}}, d_{\max} - q_{\text{opt}}]$, the function Π satisfies*

$$\Pi(q_{\text{opt}} + q) \geq \Pi(q_{\text{opt}} - q) \qquad \text{“overordering is better than underordering”}, \tag{21}$$

  then we expect overorders in low profit cases, i.e., $q_{\text{opt}} < \frac{d_{\max}+d_{\min}}{2} \Rightarrow \mathrm{E}_Q[Q - q_{\text{opt}}] > 0$.

- *If for all $q \in [d_{\min} - q_{\text{opt}}, d_{\max} - q_{\text{opt}}]$, the function Π satisfies*

$$\Pi(q_{\text{opt}} + q) \leq \Pi(q_{\text{opt}} - q) \qquad \text{“underordering is better than overordering”}, \tag{22}$$

  then we expect underorders in high profit cases, i.e., $q_{\text{opt}} > \frac{d_{\max}+d_{\min}}{2} \Rightarrow \mathrm{E}_Q[Q - q_{\text{opt}}] < 0$.

Proof
In Figure 3 we have plotted $q \mapsto q e^{\beta\Pi(q_{\text{opt}}+q)}$ between $d_{\min} - q_{\text{opt}}$ and $d_{\max} - q_{\text{opt}}$ in the two principal cases. Note that the shaded areas are equally wide.

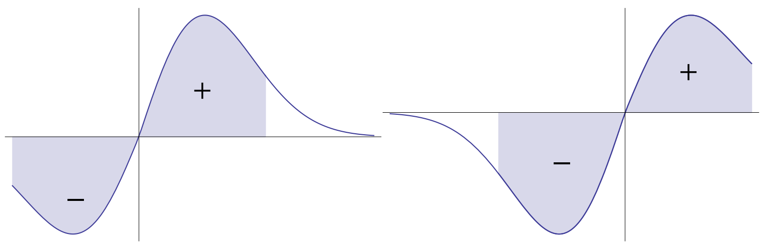

Figure 3: $q_{\text{opt}} < \frac{d_{\max}+d_{\min}}{2}$ (left) $\qquad q_{\text{opt}} > \frac{d_{\max}+d_{\min}}{2}$ (right)

If (21) is satisfied, then the shaded area under the axis to the left is smaller than or equal to the shaded area over the same graph. The positive values are hence at least as big and are integrated over a strictly longer interval. This implies a net positive value in (20). If (22) is satisfied, then the shaded area under the graph to the right is bigger than or equal to the shaded area over the same graph. The negative values are hence at least as big and are integrated over a strictly longer interval. This implies a net negative value in (20).

□

Proposition 4.1 generalizes Proposition 4 in Su (2008). In the particular case where the distribution of D is uniform, it is evident that there exist constants C_1, C_2 such that

$$\Pi(q) = C_1 - C_2(q - q_{\text{opt}})^2. \tag{23}$$

In this case, conditions (21) and (22) are satisfied for an arbitrary q_{opt}, and it follows that we have overordering if $q_{\text{opt}} < \frac{d_{\max}+d_{\min}}{2}$ and underordering if $q_{\text{opt}} > \frac{d_{\max}+d_{\min}}{2}$. This gives a new proof of the pull-to-center effect in the uniform case.

As we can see, the analysis of the sign in the pull-to-center effect is straightforward. In the next few sections, however, we will use the expression in (20) to examine the rate of change of this effect. Even though some partials are discussed in Su (2008), these are alternative expressions for the expected order, and are not related to the rate of change effect. Proposition 5 in Su (2008) discusses changes in expected profit, but this result follows directly from the relation

$$\frac{\partial \mathrm{E}_Q[c(Q)]}{\partial \beta} = -\mathrm{E}_Q[c(Q)^2] + \mathrm{E}_Q[c(Q)]^2 = -\mathrm{Var}_Q[c(Q)] < 0. \tag{24}$$

and is not relevant to us. In our newsvendor problem, it is of interest to examine what happens to the expected order when we change β. If we differentiate the expression in (20) w.r.t β, see the appendix for details, we can see that

$$\frac{\partial \mathrm{E}_Q[Q - q_{\text{opt}}]}{\partial \beta} = \mathrm{E}_Q[Q \cdot \Pi(Q)] - \mathrm{E}_Q[Q] \cdot \mathrm{E}_Q[\Pi(Q)] = \mathrm{Cov}_Q[Q, \Pi(Q)]. \tag{25}$$

In general, $\frac{\partial \mathrm{E}_Q[Q-q_{\text{opt}}]}{\partial \beta}$ can be either negative or positive, depending on the relative strength of the two terms in the middle of (25). The final result can be stated as follows.

Proposition 4.2
Assume that D is uniformly distributed on $[d_{\min}, d_{\max}]$ and that the (sensitivity) parameter β in (19) increases. Then the expected order decreases in low profit cases and the expected order increases in high profit cases, i.e., the expected error in ordering decreases.
Proof
The proof is technical, and is shown in the appendix. The basic idea is to use the expectation format in (25) to rewrite the expression to a form more suitable for analysis.

□

Proposition 4.2 has an intuitive interpretation. If the retailer is more concerned about costs (which is reflected in a larger β parameter), he is less inclined to deviate from the optimal order.

4.2 Skewness of underorders/overorders

In the empirical part of the paper we find that the overall size of the underorders is larger than the overall size of the overorders. This happens even though the critical fractiles are unbiased, i.e., the mean critical fractile is $\frac{d_{\min}+d_{\max}}{2}$. The purpose of this section is to explain that, under certain conditions, this effect is what we expect when orders are selected via probabilistic cost efficiency. The main result can be stated as follows:

Theorem 4.3
Assume that D is uniformly distributed on $[d_{\min}, d_{\max}]$, and that $\mathcal{ED}$ is an experimental design where given R, S, the selected values of W are always symmetric about $\frac{R+S}{2}$. Then, if orders are chosen from a cost efficient distribution with parameter β and $\beta = \beta(R, S, W)$ is a strictly increasing function of W, an overweight of underorders is expected, i.e., the average over all the experiments in $\mathcal{ED}$ satisfies

$$\mathrm{E}_{\mathcal{ED}}[Q - q_{\mathrm{opt}}] < 0 \tag{26}$$

Proof
The formal proof is technical, and all the details are shown in the appendix.

□

Even though the formal proof is somewhat elaborate, the essence of the proof is quite easy to understand. If we choose two values $W_1 < W_2$ symmetric about $\frac{R+S}{2}$, the smaller value will lead to underorder while the bigger value leads to overorder. When the the agents are boundedly rational with the same value of β in the two cases, it is possible to show that the size of the underorder will exactly match the size of the overorder. Under the conditions stated in the theorem, however, the β value used with W_1 is strictly smaller than the one used with W_2. According to Proposition 4.2, the error in ordering is reduced when we increase β. Therefore the error in ordering using W_2 is smaller than the error in ordering using W_1, i.e., the agents make larger errors when they underorder than when they overorder.

5 Analyzing empirical data

In this section we will use the theoretical machinery from Section 3 and 4 to analyse empirical data. Ulrich Thonemann has kindly given us access to the data used in the paper Becker-Peth et. al (2013). As Becker-Peth et. al (2013) contains the protocol and all specific information related to the experiment, we will only provide a minimum of detail. The experiment can be described (very roughly) as follows:

31 persons participated in the experiment. After a 15 minutes briefing on the newsvendor problem and a warm-up phase presenting 5 different contracts, the data collection started with

the following case: Assume that demand D is uniformly distributed on the interval $[0, 100]$, and that $R = 100, W = 80, S = 75$. On the basis of this information the participants should suggest order quantities, and the suggested order quantities were:

$$75, 82, 30, 50, 80, 100, 85, 45, 50, 100, 95, 70, 60, 80, 50, 50, 20,$$
$$90, 45, 70, 87, 50, 80, 50, 80, 20, 50, 55, 100, 80, 65. \tag{27}$$

Our main research question can be phrased as follows: Is the observed set of orders consistent with bounded rationality? According to our definition, agents are boundedly rational iff less costly orders are more probable, i.e, the orders must be drawn from a cost-efficient distribution. To carry of Erlanders graphical test, we need to sort the observations into K bins. With only 31 observations $K = 5$ is a more or less canonical choice, so we will start our analysis with that case.

The optimal order is $q_{\text{opt}} = 80$ in this case. With $K = 5$ we partition the interval into 5 bins:

$$[0, 20] - [21, 40] - [41, 60] - [61, 80] - [81, 100]$$

The observed frequencies are

$$2 \qquad 1 \qquad 11 \qquad 9 \qquad 8. \tag{28}$$

The cost associated with each bin can be computed in several different ways, and we will start with costs defined in terms of the midpoint in each interval. The midpoint orders lead to the costs

$$612.5 \qquad 312.5 \qquad 112.5 \qquad 12.5 \qquad 12.5. \tag{29}$$

The next step is to draw resamples from the observed frequency distribution in (28). Using $N = 31$, we compute the pair $(\bar{c}, \mathcal{LL}(\bar{p}))$ from the formulas (4) and (5). The plot in Figure 4 shows a plot of these pairs after 200 resamples. In Figure 5 we have also drawn the upper and lower limits for the 95% confidence band defined via (8), (9), (10), and (11).

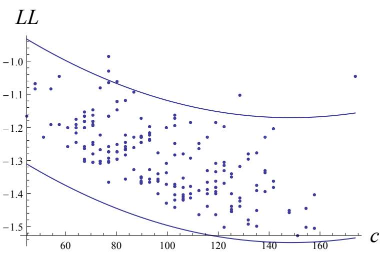

Figure 4: Likelihood values as a function of costs

Using 10000 resamples we recorded 9838 points within the 95% confidence band, and there is no reason for reject our null hypothesis: *All variation in the data can be explained from probabilistic cost efficiency.*

In this particular case the finding is confirmed by the likelihood ratio test suggested in Section 3.2. The best fit is obtained using $\hat{\beta} = 0.00323$, and using (12), (13), and (14) we report $\mathcal{X}^2 = 5.26$ which is well within the non-rejection region for a χ^2 variable with 3 degrees of freedom (7.81).

Non-rejection of the null-hypothesis does not in itself provide any strong support for our model. The interesting feature with Erlander's test is that the diagnostic plot shown in Figure 4 reveals a distribution where the likelihood value falls with increasing costs. This suggests that the agents, broadly speaking, are boundedly rational in the sense of our definition. This is supported by the relatively low value of χ^2 which indicates a fairly good model fit.

5.1 Robustness

The choice of $K = 5$ bins and costs specified via the mid-point in this bins is somewhat arbitrary. To check if the conclusion is robust with respect to the number of bins, we carried out the same analysis using $K = 4$ and $K = 6$ bins. The diagnostic plots for Erlanders graphical test are shown in Figure 5.

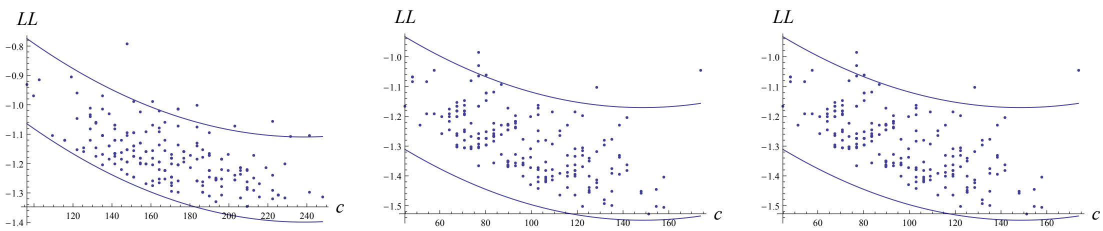

Figure 5: Likelihood plots using $K = 4, 5, 6$ bins

We see that all 3 likelihood plots are falling, and with 10000 resamples the recorded values within the bands: 9588, 9838, 9922 give no reason to reject that all variation can be explained via probabilistic cost efficiency. The maximum likelihood test, too, offers the same conclusion. The reported values are: $\chi^2 = 4.25, 5.26, 8.43$ and should be compared with the levels for rejection $5.99, 7.81, 9.49$, respectively. All reported values hence lead to non-rejection.

The mid-point specification of cost is also somewhat arbitrary, but other specifications do not seem to change the overall picture. Using the observed average cost within bins instead of the midpoint specification, the reported values for the maximum likelihood test changes to: $\chi^2 = 3.66, 4.51, 7.06$. The numerical values change, but the overall conclusion, non-rejection, does not change.

5.2 Analysis of the experimental data

In Becker-Peth et. al (2013) the experiment above was repeated 28 times with different combinations of W, S. The value $R = 100$ was fixed throughout the experiments. We examined the data in search of learning effects, but we could not find any noticeable development over time. This is what we would expect as one would probably need a very large number of repetitions to notice a learning effect. We hence assume that the 28 different cases can be handled as separate experiments.

Even though we noted a few exceptions, the specification of costs is relatively unimportant. As a clear majority leads to the same overall conclusion, we only report the findings with a mid-point specification of costs.

In summary, the likelihood ratio test appears to be stronger than Erlander's graphical test, and we first take a look at the χ^2 values for the likelihood ratio test. These values are reported in Figure 6.

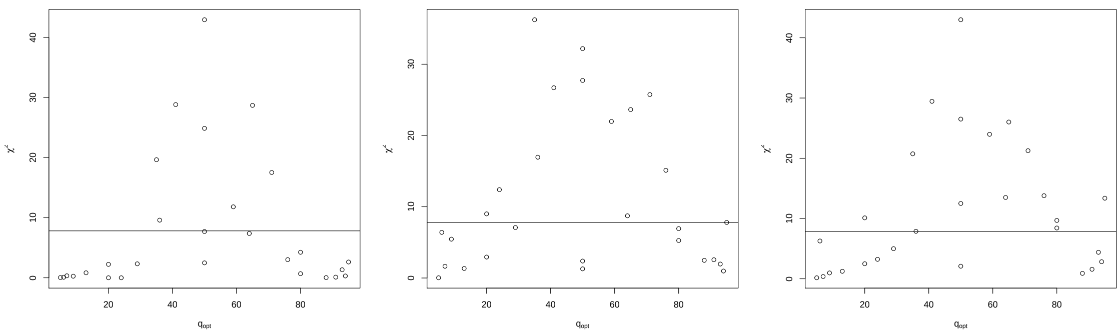

Figure 6: χ^2 values as a function of q_{opt}, 4 (left), 5 (middle), and 6 (right) bins. The solid lines show the rejection level.

As we can see from Figure 6, cases with small or big critical fractiles typically lead to non-rejection, while the results in the middle range typically lead to rejection. Before we enter into a general discussion, we consider similar plots for the β parameter. In this theory the β parameter measures how sensitive the agents are with respect to costs. For each case a value of β was fitted in the sense of maximum likelihood, and the reported values are shown in Figure 7.

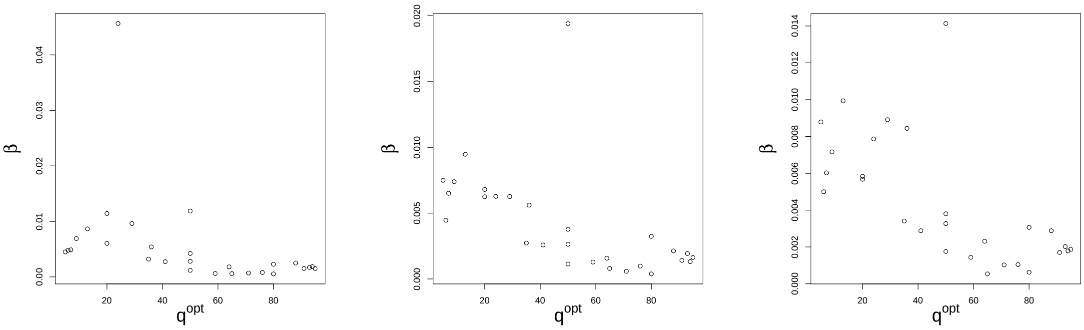

Figure 7: Sensitivity to cost $\hat{\beta}$ as a function of q^{opt}, 4 (left), 5 (middle), and 6 (right) bins.

The 3 versions $K = 4, 5, 6$ are consistent and we conclude that the sensitivity to costs falls with decreasing wholesale cost (significant by any standard). This makes good sense as it seems logical that agents are more concerned about costs when goods are expensive. Note that the β values are roughly equal in the 3 plots; the apparent difference is due to different scales on the axes.

5.3 Small or big critical fractiles

When the critical fractile is smaller than 0.2 or bigger than 0.8, all findings are consistent. With a small number of mild exceptions (which is expected at the 5% level), the likelihood ratio test leads to non-rejection. The likelihood plots from Erlander's graphical test are all falling, and all cases lead to non-rejection. When the likelihood plots are falling, we are able to conclude that agents, at least in a broad sense, behave according to our definition of bounded rationality. This impression is strengthened by the relatively small χ^2 values, which can be interpreted as good model fits.

Skewness of overorders/underorders

According to the theory in Section 4.1, the agents can be expected to underorder when $q_{\text{opt}} \geq 80$ and to overorder when $q_{\text{opt}} \leq 20$. In our data set there are 7 cases of each kind. The differences between the optimal and the average observed order is shown below.

−16	−19	−21	−17	−16	−11	−16
6	8	11	−2	5	9	7

The first row shows the 7 cases where $q_{\text{opt}} \geq 80$, while the second row shows the 7 cases with $q_{\text{opt}} \leq 20$. We note that with one exception, the value -2 in the second row, the findings are consistent with our theory. Moreover we see that the underorders are (by any standard) larger than the overorders. This is just what we expect from Theorem 4.3. The plots in Figure 7 suggests that agents are more concerned about costs in low profit cases. When we increase

W, the profitability goes down, leading to a larger value of β. The observed skewness is hence consistent with probabilistic cost efficiency.

5.4 The middle range

When the critical fractile is in the interval $(0.2, 0.8)$ the situation is less supportive. As we can see from Figure 6, our null hypothesis is usually rejected in such cases. This means that our model is unable to explain all the variation in the data. Note that rejection of the null hypothesis does not render the model useless. It is still possible that the model explains a major part of the variation and a good model fit is not excluded. The diagnostic plots from Erlander's graphical test provide some interesting insights. Figure 8 shows the likelihood plots for the four cases with the highest values of χ^2 in the likelihood ratio test.

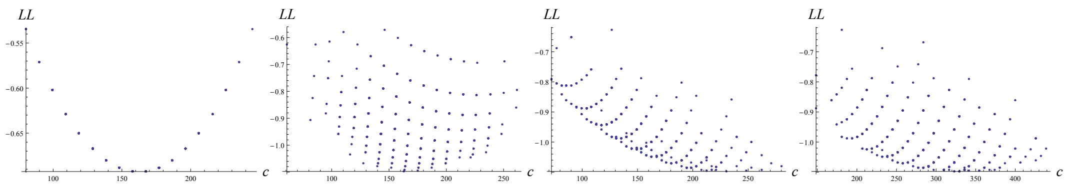

Figure 8: likelihood plots for the four worst cases (highest values of χ^2)

While the third plot indicates a falling pattern, the other plots are questionable. If we compare with the two other cases with critical fractile 0.5, the difference is striking, see Figure 9.

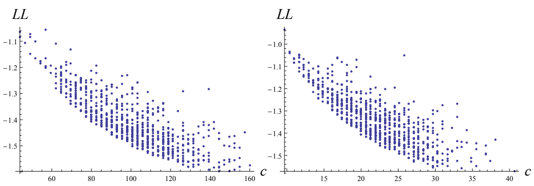

Figure 9: likelihood plots for two cases with critical fractile 0.5

Inspection of the parameters used in these experiments offer a simple explanation to the problem. In Figure 8 the salvage parameter S has low values, while the values are high in Figure 9. This principle appears to apply to most cases in the middle range. If the salvage value is high, the likelihood plots are clearly falling combined with a fair model fit. If the salvage value is low, the likelihood plots are questionable and the overall model performance is poor.

5.5 Modelling risk aversion

While parameters sets with the same critical fractile lead to the same optimal order and expected profit, cases where S is small are subject to more risk. We have

$$\mathrm{Var}\Pi(q) = \mathrm{Var}[\text{"random profit when ordering } q \text{ units"}] = (R-S)^2\mathrm{Var}[\min[D,q]] \tag{30}$$

It is then natural to ask if risk aversion could be an issue. When the agents are risk averse, the perceived cost may be different from the loss in expected profit. A model taking this issue into account could be

$$c(q) = \Pi(q_{\mathrm{opt}}) - \Pi(q) + \lambda \cdot \mathrm{Var}\Pi(q) \tag{31}$$

where λ is a constant parameter reflecting the level of risk aversion, i.e., the shadow cost of a variance constraint. Our new MNL model is hence equipped with two parameters, β measuring the agents sensitivity to costs and λ measuring the amount of risk aversion. To each set of observations we can fit parameter values in the sense of maximum likelihood. With D uniformly distributed on $[0, 100]$, we get

$$\mathrm{Var}\Pi(q) = \frac{(R-S)^2 q^3(400-3q)}{120\,000} \tag{32}$$

With our new specifications of costs, performance is greatly improved. All the likelihood plots in Erlander's graphical test are falling, and no cases are rejected. The same thing happens in the likelihood ratio test, i.e., the cases leading to rejection in Figure 6 are now much better, see Figure 10. In 28 experiments we should of course tolerate a small number of mild rejections, so the overall conclusion is that our new model works surprisingly well. Moreover, the small χ^2 values imply that a good model fit is obtained in all cases. We are hence able to conclude that the variation in all experiments can be explained by a combination of probabilistic cost efficiency with risk aversion.

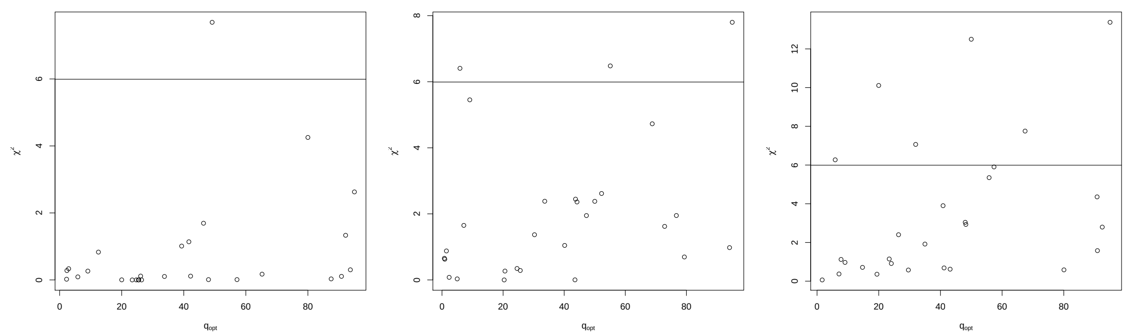

Figure 10: χ^2 values as a function of q_{opt}, 4 (left), 5 (middle), and 6 (right) bins. The solid lines show the rejection level.

5.6 Model fit in terms of AIC

In our new model we have included a new parameter λ in comparison with the original model. It is hence of interest to see how much an extra parameter improves performance in terms of the Akaike information criterion AIC. In Table 1 we have computed AIC values for our original 1-parameter model (AIC1par), for the new 2-parameter model including risk aversion (AIC2par), and for a saturated model with 4 parameters (AIC4par), for the case $K = 5$ bins. In general we should prefer a model with low AIC.

q_{opt}	AIC4par	AIC2par	AIC1par
5	17.86	13.87	11.87
6	25.06	24.26	22.26
7	25.93	22.76	20.76
9	27.05	25.78	23.78
13	32.50	28.94	27.17
20	45.58	41.91	41.05
20	28.72	25.03	27.22
24	30.74	26.78	30.94
29	37.26	33.40	34.80
35	29.47	25.47	41.59
36	38.17	34.34	40.65
41	37.72	33.86	45.07
50	39.41	36.60	47.27
50	53.97	49.97	48.61
50	39.84	36.52	49.93
50	51.70	48.89	46.89
59	44.54	41.06	49.52
64	51.38	48.36	49.75
65	44.07	41.29	49.88
71	44.02	41.20	50.89
76	46.01	45.25	47.57
80	50.28	48.64	46.91
80	52.91	50.22	50.38
88	49.24	46.05	44.48
91	49.14	46.11	44.42
93	42.14	38.49	37.12
94	44.62	41.11	39.11
95	35.63	35.53	33.53

Table 1: AIC values for 3 different model specifications ($K = 5$ bins)

The findings are very consistent. When the critical fractile is small or big, we should prefer the simple 1-parameter model. In the middle range risk aversion is usually needed to explain the observed behavior, and the 2-parameter model is our preferred choice.

6 Technical summary

In this paper, we have established a link between probabilistic cost efficiency and bounded rationality in the newsvendor model. In our opinion, probabilistic cost efficiency is a superior approach to bounded rationality, and we believe that this paper is the first to focus on this connection in the context of newsvendor models.

In the paper we have made Erlander's graphical test fully operational, and we have extended the test such that it can be applied to small samples. The main advantage with Erlander's test is that it provides a diagnostic tool we can use to see if the agents largely behave according to our definition. We also propose a likelihood ratio test that can be used for additional information.

The empirical analysis in Section 5 reveals several interesting insights that we believe are new.

- All the variation in our data can be explained by a combination of probabilistic cost efficiency with risk aversion. When the critical fractile is small (< 0.2) or big (> 0.8), a simple model without risk aversion is sufficient. In the middle range, a model with risk aversion is needed to explain the data.
- The value of the β parameter (which measures sensitivity to costs) appears to decrease when the critical fractile increases. Costs thus appear to be more important when goods are expensive, which seems intuitive. It would be of interest to check whether this is true for data sets other than ours.
- In our data set, we find that the underorders are consistently larger than overorders, leading to an overall majority of underorders. We have proved that this is just what we would expect from theory, see Theorem 4.3.

7 Concluding remarks

We believe that the theory in Section 4 is built on so much common sense that we would be surprised if these effects are not present in almost any experiment on newsvendor behavior. Moreover, we think that a corresponding bias must be present in real life cases. Managers should understand the nature of this type of randomness, and use this knowledge to improve performance. In the introduction we discussed several papers focusing the "pull-to-center" effect. While we are in no position to refute any of these theories, we suggest that there might be a much simpler explanation, i.e., that the effect is driven by a systematic bias caused by random choices. In a system driven by discrete choice, e.g., via random utility maximization, the resulting distribution will be probabilistically cost efficient. Such choices are consistent with our theory and no additional effect is needed to explain the pull-to-center effect.

The successful application of discrete choice models in several fields in economics demonstrates that the effects that arise from pure randomness of this type are very real. Many applications report quite remarkable model fits, which in turn support a general belief that these models reflect major trends in real world systems. We see no reason why a similar approach should be less successful in the newsvendor context. In this paper we have demonstrated that these models are able to explain non-trivial economic effects that are consistent with empirical observations. Almost any problem in the newsvendor theory can be rephrased within this extended context, and in general one could ask how strategies, profits and equilibria change if the newsvendor choose orders from a cost efficient distribution. In that respect we feel this theory has a lot to offer in terms of problems for future research, and hope that other researchers will pursue this line of research.

8 Appendix

8.1 Proof of Proposition 4.1

To prove this, we use a result by Satterthwaite (1942) that the distribution of a sum of n squared standard normal variables with m inhomogeneous linear restrictions on the variables is distributed as a shifted χ^2 variable with $n-m$ degrees of freedom. For the asymptotics of the χ^2 statistic $\mathcal{X}^2$ in the common K category multinomial case, we essentially start with $n = K-1$, and the theorem states that we lose a further degree of freedom by imposing a linear restriction. Consequently, we must demonstrate that this restriction can be maintained in the proof of the asymptotic result.

We can write $z_k = \sum_{i=1}^{N} x_{ik}$ so that $(x_{i1}, x_{i2}, ..., x_{iK})$ for $i = 1, 2, .., N$ are independent multinomial $(1, p_1, p_2, ..., p_K)$. Then, the expectations, variances, and covariances are $\mathrm{E}[x_{ik}] = p_k$, $\sigma_{kk} = p_k(1-p_k)$, and $\sigma_{jk} = -p_j p_k$, respectively. Because $x_{iK} = 1 - \sum_{k=1}^{K-1} x_{ik}$, to avoid singularity we omit the last component and consider the column vector $\mathbf{x_i} = (x_{i1}, x_{i2}, ..., x_{iK-1})'$ with nonsingular $(K-1)\times(K-1)$ covariance matrix $\mathbf{\Sigma} = (\sigma_{jk})$ and inverse $\mathbf{\Sigma^{-1}} = (\sigma^{-}_{jk})$ where $\sigma^{-}_{jj} = 1/p_j + 1/p_K$ and $\sigma^{-}_{jk} = 1/p_K$ for $j \neq k$.

Let $\mathbf{p} = (p_1, p_2, ..., p_{K-1})$. In terms of the defined quantities, our $\mathcal{X}^2$ may be expressed as $\mathcal{X}^2 = N(\bar{\mathbf{x}}-\mathbf{p})'\Sigma^{-1}(\bar{\mathbf{x}}-\mathbf{p})$. By the central limit theorem the distribution limit of $\mathbf{u} = \sqrt{N}\Sigma^{-1/2}(\bar{\mathbf{x}}-\mathbf{p})$ is $N(\mathbf{0}, \mathbf{I})$ as $N \to \infty$. This means that $\mathcal{X}^2$ behaves as a sum of squares of $K-1$ asymptotically independent standard normal variables, and therefore behaves asymptotically as a χ^2 variable with $K-1$ degrees of freedom. Henceforth, we assume we have a large N leading to a good normal approximation. We can now write $\bar{\mathbf{x}} = \mathbf{p} + \frac{1}{\sqrt{N}}\mathbf{B}\cdot\mathbf{u}$ where $\mathbf{B} = \mathbf{\Sigma}^{1/2}$ so that in terms of components we have $z_k = Np_k + \sqrt{N}\sum_{j=1}^{K-1} B_{kj}u_j$ for $k = 1, 2, ..., K-1$. The omitted component is then $z_K = N - \sum_{k=1}^{K-1} z_k = Np_K - \sqrt{N}\sum_{k=1}^{K-1}\sum_{j=1}^{K-1} B_{kj}u_j$. By inserting the expressions for all z_k's into the restriction $N\bar{c} = c_K z_K + \sum_{k=1}^{K-1} c_k z_k$, we see that the restriction simplifies to $\sum_{j=1}^{K-1} b_j u_j = (\bar{c} - c)\sqrt{N}$ where $b_j = \sum_{k=1}^{K-1}(c_k - c_K)B_{kj}$ and $c = \sum_{k=1}^{K} c_k p_k$. Consequently, we have shown that the restriction is carried over to a linear restriction for the u_k variables.

To apply Satterthwaite's formula, we must normalize the coefficients so that their sum of squares equals one. This is done by taking $b^2 = \sum_{j=1}^{K-1} b_j^2$ and $a_j = b_j/b$. Then, $\rho_1 = \sum_{j=1}^{K-1} a_j u_j = (\bar{c}-c)\sqrt{N}/b$. The limiting distribution of $\mathcal{X}^2$ is therefore a χ^2 distribution with $(K-1)-1 = K-2$ degrees of freedom with its minimum shifted from zero to ρ_1^2. Tedious calculations show that $b^2 = \sum_{k=1}^{K}(c_k - c)^2 p_k$ so that $\rho_1^2 = N(c-\bar{c})^2/\sum_{k=1}^{K}(c_k - c)^2 p_k$. Alternatively, we can find the expression for ρ_1^2 solving for the minimum of the expression $\mathcal{X}^2$ with respect to the z_k's under the linear constraint.

$\square$

8.2 How (25) is derived from (20)

Note that

$$E_Q[Q - q_{\text{opt}}] = \frac{\int_{d_{\min}}^{d_{\max}} q e^{\beta\Pi(q)} dq}{\int_{d_{\min}}^{d_{\max}} e^{\beta\Pi(u)} du} - q_{\text{opt}}$$

Elementary calculus (noting that q_{opt} does not depend on β) gives

$$\begin{aligned}
\frac{\partial E_Q[Q - q_{\text{opt}}]}{\partial \beta} &= \frac{\int_{d_{\min}}^{d_{\max}} q\Pi(q) e^{\beta\Pi(q)} dq \cdot \int_{d_{\min}}^{d_{\max}} e^{\beta\Pi(u)} du - \int_{d_{\min}}^{d_{\max}} q e^{\beta\Pi(q)} dq \cdot \int_{d_{\min}}^{d_{\max}} \Pi(q) e^{\beta\Pi(u)} du}{\left(\int_{d_{\min}}^{d_{\max}} e^{\beta\Pi(u)} du\right)^2} \\
&= E_Q[Q\Pi(Q)] - E_Q[Q]E_Q[\Pi(Q)]
\end{aligned}$$

8.3 Proof of Proposition 4.2

To prove Proposition 4.2, we first need to prove the following non-trivial lemma.

Lemma A 1
Let a, b be any strictly positive real numbers, and consider the expression

$$\Phi(a,b) = \int_{-a}^{b} u^3 e^{-u^2} du \cdot \int_{-a}^{b} e^{-u^2} du - \int_{-a}^{b} u e^{-u^2} du \cdot \int_{-a}^{b} u^2 e^{-u^2} du. \tag{33}$$

- *If $a > b$, then $\Phi(a,b) < 0$.*
- *If $a = b$, then $\Phi(a,b) = 0$.*
- *If $a < b$, then $\Phi(a,b) > 0$.*

Proof
For fixed $b > 0$, define $\phi_b : [0,\infty) \to \mathbb{R}$ by

$$\phi_b(x) = \int_{-b}^{x} u^3 e^{-u^2} du \cdot \int_{-b}^{x} e^{-u^2} du - \int_{-b}^{x} u e^{-u^2} du \cdot \int_{-b}^{x} u^2 e^{-u^2} du. \tag{34}$$

If we differentiate this expression, we obtain

$$\phi_b'(x) = \int_{-b}^{x} (x^3 + u^3 - xu^2 - x^2u) e^{-u^2-x^2} du. \tag{35}$$

The key observation is the following

$$\phi_b'(x) = \int_{-b}^{x} (x-u)^2 (x+u) e^{-u^2-x^2} du. \tag{36}$$

Now if $x \geq b$, the integrand in (36) is strictly positive on $(-b, x)$, and it follows that ϕ_b is strictly increasing on (b, ∞). Because $\phi_b(b) = 0$, we obtain $\phi_b(x) > 0$ on the open interval (b, ∞), and we define a new function $\psi_b : [0, \infty) \to \mathbb{R}$ by

$$\psi_b(x) = \int_{-x}^{b} u^3 e^{-u^2} du \cdot \int_{-x}^{b} e^{-u^2} du - \int_{-x}^{b} u e^{-u^2} du \cdot \int_{-x}^{b} u^2 e^{-u^2} du. \tag{37}$$

Change variables by $v = -u$, to see that

$$\psi_b(x) = -\int_{-b}^{x} v^3 e^{-v^2} dv \cdot \int_{-b}^{x} e^{-v^2} dv + \int_{-b}^{x} v e^{-v^2} du \cdot \int_{-b}^{x} v^2 e^{-v^2} dv. \tag{38}$$

Hence, $\psi_b(x) = -\phi_b(x)$ for any $b > 0$ and any $x \in [0, \infty)$. Since by definition $\phi_b(x) = \psi_x(b)$, it follows that $\phi_b(x) = -\phi_x(b)$ which is strictly negative when $0 \leq x < b$ since $\phi_x(b)$ is strictly positive when $b \in (x, \infty)$ (from the first part of this proof).

□

We are now ready to prove Proposition 4.2. In general we have

$$\frac{\partial \mathrm{E}_Q[Q - q_{\mathrm{opt}}]}{\partial \beta} = \mathrm{E}_Q[Q \cdot \Pi(Q)] - \mathrm{E}_Q[Q] \cdot \mathrm{E}_Q[\Pi(Q)]. \tag{39}$$

A key step in the proof is to add and subtract terms to see that

$$\frac{\partial \mathrm{E}_Q[Q - q_{\mathrm{opt}}]}{\partial \beta} = \mathrm{E}_Q[(Q - q_{\mathrm{opt}}) \cdot \Pi(Q)] - \mathrm{E}_Q[Q - q_{\mathrm{opt}}] \cdot \mathrm{E}_Q[\Pi(Q)]. \tag{40}$$

When D is uniformly distributed, there exist constants C_1, C_2 where $C_2 > 0$ such that

$$\Pi(q) = C_1 - C_2(q - q_{\mathrm{opt}})^2. \tag{41}$$

If we insert that expression into (40), change variables as in (20), and simplify the expression, we obtain

$$\frac{\partial \mathrm{E}_Q[Q - q_{\mathrm{opt}}]}{\partial \beta} = -\frac{C_2 \int_{-c}^{d} q^3 e^{-\beta \cdot C_2 q^2} dq}{\int_{-c}^{d} e^{-\beta \cdot C_2 q^2} dq} + \frac{\int_{-c}^{d} q e^{-\beta \cdot C_2 q^2} dq}{\int_{-c}^{d} e^{-\beta \cdot C_2 q^2} dq} \cdot \frac{C_2 \int_{-c}^{d} q^2 e^{-\beta \cdot C_2 q^2} dq}{\int_{-c}^{d} e^{-\beta \cdot C_2 q^2} dq}. \tag{42}$$

Here, $c = q_{\mathrm{opt}} - d_{\mathrm{min}}, d = d_{\mathrm{max}} - q_{\mathrm{opt}}$. Because $C_2 > 0$, we see that the sign is determined by the expression

$$-\int_{-c}^{d} q^3 e^{-\beta\cdot C_2 q^2} dq \cdot \int_{-c}^{d} e^{-\beta\cdot C_2 q^2} dq + \int_{-c}^{d} q e^{-\beta\cdot q^2} dq \cdot \int_{-c}^{d} q^2 e^{-\beta\cdot C_2 q^2} dq. \tag{43}$$

If we make a second linear change of variables $v = \sqrt{\beta \cdot C_2} \cdot q$, the previous expression equals

$$-\frac{1}{\beta^2 C_2^2}\left(\int_{-a}^{b} v^3 e^{-v^2} dv \cdot \int_{-a}^{b} e^{-v^2} dv - \int_{-a}^{b} v e^{-v^2} dv \cdot \int_{-a}^{b} v^2 e^{-v^2} dv\right), \tag{44}$$

where $a = \sqrt{\beta \cdot C_2}(q_{\text{opt}} - d_{\text{min}})$ and $b = \sqrt{\beta \cdot C_2}(d_{\text{max}} - q_{\text{opt}})$. The conclusions in Proposition 4.2 then follow directly from Lemma A 1.

□

8.4 Details for the proof of Theorem 4.3

The key step in the proof of Theorem 4.3 is the following proposition:

Proposition A.2
Assume that D is uniformly distributed on $[d_{\text{min}}, d_{\text{max}}]$, and that $\mathcal{ED}$ is an experimental design where W, conditional on R, S, has a distribution that is symmetric about $\frac{R+S}{2}$. Then if β is a constant with respect to W

$$\mathrm{E}_{\mathcal{ED}}[\mathrm{E}_Q[Q - q_{\text{opt}}]] = 0. \tag{45}$$

Proof
If D is uniformly distributed, it is easy to verify that in (23), $C_2 = \frac{R-S}{2(d_{\text{max}} - d_{\text{min}})}$. Note that C_2 does not depend on W. It follows from (20) that

$$\mathrm{E}_Q[Q - q_{\text{opt}}] = \frac{\int_{d_{\text{min}}-q_{\text{opt}}}^{d_{\text{max}}-q_{\text{opt}}} q e^{-C_2 \beta\, q^2} dq}{\int_{d_{\text{min}}-q_{\text{opt}}}^{d_{\text{max}}-q_{\text{opt}}} e^{-C_2 \beta\, q^2} dq}. \tag{46}$$

If we define a function h by

$$h(x) = \frac{\int_{d_{\text{min}}-x}^{d_{\text{max}}-x} q e^{-C_2 \beta\, q^2} dq}{\int_{d_{\text{min}}-x}^{d_{\text{max}}-x} e^{-C_2 \beta\, q^2} dq}, \tag{47}$$

then $h(q_{\text{opt}}) = \mathrm{E}_Q[Q - q_{\text{opt}}]$. The key to the proof is to realize that h is antisymmetric around $x = \frac{d_{\text{max}} + d_{\text{min}}}{2}$, i.e., that

$$h\left(\frac{d_{\text{max}} + d_{\text{min}}}{2} - x\right) = -h\left(\frac{d_{\text{max}} + d_{\text{min}}}{2} + x\right). \tag{48}$$

The details are tedious but straightforward and are omitted. Note that the function h in (47) does not depend on W. In the newsvendor model with uniform demand D, we have

$$q_{\text{opt}} = d_{\min} + (d_{\max} - d_{\min})\frac{R - W}{R - S}. \tag{49}$$

We first assume that R, S are constants and that W has a continuous distribution with density $f_W(w)$ on the interval $[S, T]$. Then, by (47) and (49)

$$\mathrm{E}_W[\mathrm{E}_Q[Q - q_{\text{opt}}]] = \int_S^R h\left(d_{\max} + (d_{\max} - d_{\min})\frac{R - w}{R - S}\right) f_W(w)dw. \tag{50}$$

Now, change variables by $w = \frac{R+S}{2} + u$ to see that

$$\mathrm{E}_W[\mathrm{E}_Q[Q - q_{\text{opt}}]] = \int_{-\frac{R-S}{2}}^{\frac{R-S}{2}} h\left(\frac{d_{\max} + d_{\min}}{2} - u\left(\frac{d_{\max} - d_{\min}}{R - S}\right)\right) f_W\left(\frac{R + S}{2} + u\right) du. \tag{51}$$

Because the integrand is antisymmetric by (48) and our assumptions on f_W, it follows that $\mathrm{E}_W[\mathrm{E}_Q[Q - q_{\text{opt}}]] = 0$. If R, S are not constants, extra levels of expectation do not change the result and hence

$$\mathrm{E}_{\mathcal{ED}}[\mathrm{E}_Q[Q - q_{\text{opt}}]] = 0. \tag{52}$$

If W has a discrete distribution, exactly the same argument applies (just replace integrals with sums). The result in (52) holds when W has any discrete or continuous distribution, which proves the result in Proposition A.1 for any conceivable experimental design.

□

8.4.1 Proof of Theorem 4.3

It follows from Proposition A.2 that the net balance would have been zero if $\beta = \beta\left(\frac{R+S}{2}\right)$ for all W. If $W < (R + S)/2$, then $q_{\text{opt}} > \frac{d_{\max}+d_{\min}}{2}$. Because $\beta(W) < \beta\left(\frac{R+S}{2}\right)$, if follows from Proposition 4.2 that the order is decreased (leading to larger underorder) in comparison with the case $\beta = \beta\left(\frac{R+S}{2}\right)$. Conversely, if $W > (R + S)/2$, then $q_{\text{opt}} < \frac{d_{\max}+d_{\min}}{2}$. Because $\beta(W) > \beta\left(\frac{R+S}{2}\right)$, it follows from Proposition 4.2 that the order quantity is decreased (leading to a smaller overorder) in comparison with the case where $\beta = \beta\left(\frac{R+S}{2}\right)$. In essence, the overorders are reduced in size while the underorders are increased in size (in comparison with the neutral case in Proposition A.2), which proves the theorem. □

Acknowledgement
The authors are very grateful to Professor Ulrich Thonemann for giving us access to and permission to use his experimental data. The authors wish to thank the editors and the reviewers for several substantial suggestions on how to improve the paper.

References

[1] Becker-Peth, M., E. Katok, and U.W. Thonemann. 2013. Designing buyback contracts for irrational but predictable newsvendors. *Management Science* 59, 8, 1800–1816.

[2] Bolton, G., A. Ockenfels, and U. Thonemann. 2012. Managers and students as newsvendors: how out-of-task experience matters. *Management Science*, 58, 12, 2225–2233.

[3] Bolton, G., and E. Katok. 2008. Learning by doing in the newsvendor problem: a laboratory investigation of the role of experience and feedback. *Manufacturing & Service Operations Management*, 10, 3, 519–538.

[4] Bostian, A. J., C. A. Holt, and A. M. Smith. 2008. Newsvendor pull-to-center effect: adaptive learning in a laboratory experiment. *Manufacturing & Service Operations Management*, 10, 4, 590–608.

[5] Cachón G. P., 2003. Supply chain coordination with contracts. In: *The Handbook of Operations Research and Management Science: Supply Chain Management: Design, Coordination and Operation.* Chapter 6. A. G. de Kok and S. C. Graves (eds.), Elsevier, Amsterdam, pp. 229–340.

[6] Camerer, C., and T.- H. Ho. 1999. Experience-weighted attraction learning in normal form games. *Econometrica*, 67, 4, 827–874.

[7] Conlisk, J. 1996. Why bounded rationality? *Journal of Economic Literature*, 34, 2, 669–700.

[8] Croson, R., K. Schultz, E. Siemsen, and M. L. Yeo. 2012. Behavioral operations: the state of the field. *Journal of Operations Management*, 31, 1–2, 1–5.

[9] Etzioni, A. 1986. Rationality is anti-entropic. *Journal of Economic Psychology*, 7, 17–36.

[10] Erlander, S. 2010. *Cost-Minimizing Choice Behaviour in Transportation Planning*, Springer, Berlin Heidelberg.

[11] Gino, F., and G. Pisano. 2008. Toward a theory of behavioral operations. *Manufacturing & Service Operations Management*, 10, 4, 676–691.

[12] Jörnsten, K., and J. Ubøe. 2010. Quantification of preferences in markets. *Journal of Mathematical Economics*, 46, 4, 453–466.

[13] Jörnsten, K., I. Thorsen, and J. Ubøe. 2004. Replication/prediction problems in the journey to work. *Environment and Planning A*, 36, 347–364.

[14] Kremer, M., S. Minner, and L. N. Van Wassenhove. 2010. Do random errors explain newsvendor behavior? *Manufacturing & Service Operations Management*, 12, 4, 673–681.

[15] Mattsson, L.-G., and J. W. Weibull. 2002. Probabilistic choice and procedurally rationality. *Games and Economic Behavior*, 41, 61–78.

[16] Manski, C. 1977. The structure of random utility models. *Theory and Decision*, 8, 229–254.

[17] Matejka, F., and A. McKay. 2013. Rational inattention to discrete choices: a new foundation for the multinomial logit model. Manuscript.

[18] McFadden, D. 1974. Conditional logit analysis of qualitative choice behavior. In Zarembka, P., editor, *Frontiers in Econometrics*, pp105–142. Academic Press, New York.

[19] Qin, Q., R. Wang, A. J. Vakharia, Y. Chen, and M. Seref, 2011. The newsvendor problem: Review and directions for future research. *European Journal of Operational Research* 213, 361–374.

[20] Rudi, N., and D. Drake. 2013. Observation bias: the impact of demand censoring on newsvendor level and adjustment behavior. To appear in Management Science.

[21] Satterthwaite, F. E. 1942. Linear restrictions on chi square. *Annals of Mathematical Statistics*, 13, 3, 326–331.

[22] Schweitzer, M. E., and G. P. Cachón. 2000. Decision bias in the newsvendor model with a known demand distribution: experimental evidence. *Management Science*, 46, 3, 404–420.

[23] Sims, C. A. 1998. Stickiness. *Carnegie-Rochester Conference Series on Public Policy*, 49.

[24] Sims, C. A. 2003. Implications of rational inattention. *Journal of Monetary Economics*, 50, 3, 665–690.

[25] Smith, T. 1978. A cost-efficiency principle of spatial interaction behaviour. *Regional Science and Urban Economics*, 8, 313–337.

[26] Su, X. 2008. Bounded rationality in newsvendor models. *Manufacturing & Service Operations Management*, 10, 4, 566–589.

[27] Train, K. E. 2003. *Discrete Choice Methods with Simulation*. Cambridge University Press.

[28] Ubøe, J. 2004. Aggregation of gravity models for journeys to work. *Environment and Planning A*, 36, 715–729.

[29] Wachtel, R. E., and F. Dexter. 2010. Review of behavioral operations experimental studies of newsvendor problems for operating room management. *Anesthesia & Analgesia*, 110, 6, 1698–1710.