

Impulse Balance in the Newsvendor Game

Axel Ockenfels, University of Cologne, and Reinhard Selten, University of Bonn*

2 August 2013

Abstract

One striking behavioral phenomenon is the "pull-to-center" bias in the newsvendor game: facing stochastic demand, subjects tend to order quantities between the expected profit maximizing quantity and mean demand. We show that the impulse balance equilibrium, which is based on a simple *ex-post* rationality principle along with an equilibrium condition, predicts the pull-to-center bias and other, more subtle observations in the laboratory newsvendor game.

* Corresponding author: Ockenfels, University of Cologne, Germany, Albertus-Magnus-Platz, D-50923 Cologne, Germany; ockenfels at uni-koeln.de. We thank participants in various seminars in Germany and in the US for helpful comments. We also thank Johannes Fendrich, Lea Pyhel, Andreas Pollak, and Johannes Wahlig for excellent research assistance and help with the collection of data. Ockenfels gratefully acknowledges financial support from the German Science Foundation (DFG) through the Gottfried Wilhelm Leibniz-Program and the Research Unit "Design & Behavior" (FOR 1371).

1. Introduction

In the newsvendor game, the newsvendor faces stochastic demand for a perishable product. Prior to seeing the actual demand draw, she must decide how much of the product to stock in inventory. The newsvendor model was first introduced and analyzed by Arrow et al. (1951). They show that, from a normative perspective, the expected profit maximizing order is straightforwardly computed. However, starting with the paper by Schweitzer and Cachon (2000), many researchers found that subjects in laboratory experiments order too much, relative to the expected profit-maximizing order, when the ordering cost is high, and order too little when the cost is low.¹ This pattern is called "pull-to-center" bias. The deviation from optimality is strong and robust: it persists for various feedback conditions (Bolton and Katok 2008) and demand distributions (Benzion et al. 2008), with extensive learning possibilities and strong incentives (Bostian et al. 2008), with different framings of the context (Kremer et al. 2010), and for different types of subject pools, including undergraduate and graduate students as well as experienced managers who are familiar with newsvendor-kind of situations (Bolton et al. 2012).

There appears to be no simple explanation for the pull-to-center bias. As noted by Schweitzer and Cachon (2000), the bias cannot be explained, for instance, by risk-aversion, risk-seeking, prospect theory, loss aversion, waste aversion, and stockout aversion. Each of these models would predict that newsvendors either always overorder or always underorder, regardless of whether the optimum is below or above the center (see also Schultz et al. 2007). However, based on their observation that "*Subjects were more likely to adjust their order quantities toward prior demand than away from prior demand. ... Most of the time, however, subjects did not adjust their decisions, and across rounds the average order quantity was relatively stable*" (p. 418), Schweitzer and Cachon (2000) propose two more potential explanations. For one, bounded rationality, including what they call the *chasing demand heuristic*, which "*assumes a decision maker anchors on a prior order quantity and adjusts towards prior demand.*" Second, because

¹ The newsvendor problem is related to the probability matching phenomenon, where the decision maker must press either a red or a green button. One of the buttons yields a given prize with a fixed but unknown probability. The optimal choice is the button with the highest probability. Unlike what we see in the newsvendor game, however, experienced and properly incentivized subjects solve the probability matching problem well (see Shanks et al. 2002 and the references cited therein). The newsvendor problem differs in a number of dimensions, such as that the underlying stochastics are known to the decision maker, the prize is stochastic, and that the option that maximizes expected profit may not be the risk-averse choice.

overall average behavior appears rather stable, Schweitzer and Cachon (2000) speculate that subjects may alternatively *"behave as if their utility function incorporates a preference to reduce ex-post inventory error."*

In a newsvendor game study which is probably most closely related to ours, Ho et al. (2010) investigate the potential of utility-based equilibrium explanations in more detail. Their model of reference-dependent preferences assumes that subjects maximize a utility function that includes psychologically costs of leftovers and stockouts. One advantage of their model is that it nests the standard model and Schweitzer and Cachon's (2000) idea of *ex-post* inventory error minimization as special cases, and so allows a simultaneous, structural estimate of all three models. The estimates suggest that their model explains newsvendor behavior and profits better than the other tested models.

Ho et al.'s (2010) model might be partly interpreted as an extension of Su's (2008) work, which assumes that individuals make utility-maximizing decisions with noise: while newsvendors choose the stock optimally with a higher probability than suboptimal orders, they do not so with certainty. His quantal response framework of noisy decision making can predict the pull-to-center bias, because there is more room to deviate from the optimum toward the mean demand (the center) than toward extreme demands. Ho et al's paper extends his approach by providing a psychological basis for why decision makers might appear to make errors.²

Econometric models by Bostian et al. (2008) investigate the adaptive nature of behavior as observed by Schweitzer and Cachon in more detail. Their study suggests that the demand-chasing heuristic performs well.³ However, the explanations put forward by Schweitzer and Cachon do not fully organize the learning and treatment effects observed in their data. Yet, a larger learning model, which permits noisy adjustments, recency and reinforcement effects, captures the dynamics well.

² Other utility-based models include Becker-Peth et al. (forthcoming) and Wu and Chen (2011). Ren and Croson (2012) demonstrate that overconfidence is also a consistent explanation for suboptimal stocking. Here, overconfident newsvendors are assumed to believe their information or their estimate to be more precise than it actually is.

³ Nelson and Bearden (forthcoming) show that some of the measures used in the literature to identify demand chasing are prone to false positives and so tend to overestimate demand chasing. They argue that a simple correlation measure does not suffer from this problem. Ockenfels and Selten (2005) point to a similar problem in a related context.

These approaches are useful in showing that models of motivation and adaptation can go a long way to capture certain aspects of newsvendor behavior. At the same time there are limitations. For instance, the equilibrium models to explain newsvendor behavior are all static in nature and so do not capture the chasing demand pattern. Pure adaptation models such as the demand chasing heuristic, on the other hand, cannot easily explain treatment effects. Our study complements approaches based on limited motivation and adaptation by a *limited cognition* approach, the "impulse balance equilibrium" (IBE). While IBE has been developed in other contexts, we demonstrate that it predicts the pull-to-center bias along with other observations made in the laboratory newsvendor game.⁴ It is based on a simple principle of *bounded (ex-post) rationality* that guides how decision makers adjust behavior over time, along with a straightforward *equilibrium* condition. The basic assumption of IBE, when applied to the newsvendor game, is that newsvendors respond to "impulses", which occur if a larger order would have been better in the last period (upward impulse) or a lower order would have been better (downward impulse). Assuming that newsvendors have a tendency to move in the direction of the impulse and to balance upward and downward impulses (as specified in Section 2), IBE predicts the central tendency of the stationary distribution of the newsvendors' orders. It does so without *ex-ante* parameter estimations.

Section 2 derives the IBE in the context of a typical laboratory newsvendor game with uniformly distributed demand, which is the distribution most often applied in laboratory settings. Section 3 shows that IBE predicts various behavioral phenomena, including the pull-to-center bias, and provides additional experimental evidence. Section 4 discusses the evidence and the limitations of IBE, and concludes.

2. The impulse balance equilibrium of the newsvendor game

IBE makes quantitative predictions about the central tendency of the stationary distribution of a decision parameter on the basis of the principle of *ex-post* rationality. IBE is applicable to the repeated decision on the same parameter in situations in which

⁴ See Selten and Buchta (1999) for learning direction theory, which is the basis for impulse balance, and Selten, Abbink, and Cox (2005), Avraami, Güth, and Kareev (2005), Ockenfels and Selten (2005), Camerer et al. (2011), and Selten et al. (2011) for the performance of IBE in other games, such as in auction games and various 2x2 games. Crawford (2013) discusses IBE in comparison to other models of boundedly rational behavior.

the decision maker receives feedback not only about the payoff for the decision taken, but also about the payoffs connected to alternative decisions. If a higher parameter would have brought a higher payoff, we speak of an upward impulse, and if a lower parameter would have yielded a higher payoff, we speak of a downward impulse. The decision maker is assumed to have a tendency to move in the direction of the impulse.⁵ Moreover, it is assumed that the decision maker tends to adjust the parameter such that imbalances between upward and downward impulses are mitigated.

The application of IBE to the newsvendor game captures Schweitzer and Cachon's (2000) basic idea how to organize newsvendor behavior: the idea that subjects are concerned about *ex-post* inventory error. *Ex-post* inventory error implies that a smaller (or larger) order would have been more profitable, yielding downward (or upward) impulses. Our IBE concept assumes that the newsvendor tends to move in the direction of that impulse. In order to get quantitative predictions, IBE needs to be explicit about the strength of an impulse, which is assumed to be proportional to the amount of forgone profit. An impulse balance equilibrium is reached when expected upward impulses are equal to expected downward impulses. Here, in line with previous applications of IBE, and following prospect theory (Kahneman and Tversky 1979) and other evidence from experimental behavioral research, losses count twice in the computation of impulses.⁶

More specifically, assume that a newsvendor repeatedly faces a stochastic demand q , which is uniformly and independently distributed on the normalized interval $[0, 1]$. In each round, the unit retail price is p , and the unit ordering cost is c , with $p > c > 0$. Unit price and unit cost are assumed to be exogenous to the newsvendor, and independent of order and sale volumes. Because a newsvendor cannot sell more than ordered, the actual sales s of the newsvendor are the minimum of order quantity $x \geq 0$ and realized

⁵ Unlike reinforcement learning, where an order that yields a larger payoff has an increased probability of being chosen in the future, our approach requires a cognitive model about what would have been a better choice.

⁶ The assumption of losses counting twice is made *ex-ante*, as it is part of other applications of impulse balance theory to comparable laboratory contexts; see Selten and Chmura (2008) for more details. For evidence on the asymmetric effect of losses and gains across a wide variety of laboratory games see, e.g., Abdellaoui et al. (2007) and the references cited therein, and, for evidence within the newsvendor game context, Becker-Peth et al. (forthcoming). However, because different treatments of losses are theoretically conceivable, we will, as a robustness check, also derive predictions without the loss impulse. It turns out that the treatment of losses in IBE as applied to the newsvendor game is critical only for a predicted asymmetry which is sometimes found in the data (Section 3.2).

demand $q : s(x, q) = \min[x, q]$. The newsvendor's task is to choose the order x^* that maximizes her expected profit

$$P(x) = \int_0^1 (ps(x, q) - cx) dq \quad (1)$$

It can readily be seen that, for the uniform distribution, in each round the optimal order is given by

$$x^* = (p - c)/p \quad (2)$$

In order to compute the impulse balance equilibrium x^{IBE} we now define the relevant impulses for any order x . An upward impulse occurs if a higher profit could have been made by a larger order; that is, if demand turns out to be higher than the ordered quantity. In our newsvendor game, the upward impulse for order x and a given demand realization q is thus defined as $(p - c) \cdot \max[0, q - x]$. Correspondingly, a downward impulse occurs if the ordered quantity exceeds realized demand; for order x and a given demand realization q the downward impulse is given by $c \cdot \max[0, x - q]$. Because an impulse counts twice in case of losses, we also define a loss impulse as $\max[0, -(ps(x, q) - cx)]$.

We can now formally define an IBE equilibrium as the x^{IBE} that solves the impulse balance equation:

$$E_U(x^{IBE}) = E_D(x^{IBE}) + E_L(x^{IBE}) \quad (3)$$

Adding the expected loss impulse to the downward impulse on the right hand side of the equation makes sure that downward impulses count twice in case of losses.⁷ Now,

$$E_U(x) = \int_x^1 (p - c)(q - x) dq \quad (4)$$

is the expected upward impulse,

$$E_D(x) = \int_0^x c(x - q) dq \quad (5)$$

is the expected downward impulse, and

$$E_L(x) = \int_0^{\hat{q}} (cx - pq) dq \quad (6)$$

⁷ There can be no losses in case of upward impulses, but there can be downward impulses without losses.

is the expected loss impulse, with \hat{q} defined as the critical quantity at which profits become zero ($\hat{q}(x) = cx/p$). Calculating the integrals and applying the impulse balance equation yields

$$\frac{(p-c)(1-x)^2}{2} = \frac{cx^2}{2} + \frac{c^2x^2}{2p} \quad (7)$$

Defining $z := c/p$ and rearranging yields the impulse balance order:

$$x^{IBE} = \frac{1}{1 + \sqrt{z \frac{1+z}{1-z}}} \quad (8)$$

Figure 1 shows the IBE prediction for z between 0 and 1, and the expected profit maximizing order x^* as defined in equation (2), which can be rewritten as $x^* = (1 - z)$. (The data points shown in Figure 1 will be discussed in the next section.)

Figure 1: Predictions and behavior in the newsvendor game




Figure 1 illustrates that the IBE is generally different from expected profit maximization. There are three exceptions. When z approaches either zero or one, the IBE predictions approach the respective optimum. That is, IBE agrees with the standard prediction that

whenever the ordering cost becomes irrelevant (c goes to zero), subjects order the maximal demand ($= 1$). The reason is that if the cost is zero, downward and loss impulses are zero, too. If, however, the ordering cost becomes very large (c goes to p), subjects stop ordering positive quantities. The reason is that if the cost equals the price, upward impulses are zero.

Finally, the figure shows that the IBE prediction cuts the optimum once at some z with $0 < z < 1$. Straightforward computations show that the intersection g is given by $g = \frac{(-1+\sqrt{5})}{2} \approx 0.618$.⁸ When the ordering cost is small compared to the price ($z < g$), IBE predicts order quantities that are below the optimum, and if the ordering cost is high ($z > g$), IBE predicts order quantities that are larger than the optimum.

We finally remark that without loss impulses, our IBE prediction would be equal to $\frac{\sqrt{(1-z)/z}}{1+\sqrt{(1-z)/z}}$, and thus shaped similarly like the one shown in Figure 1, but perfectly symmetric above and below $z = 1/2$.

3. Evidence

This section shows in three subsections that IBE captures key phenomena in the newsvendor game.

3.1 Pull-to-center bias

According to the pull-to-center bias, newsvendor orders are between the expected profit maximizing quantity and the center ($=$ mean demand, $1/2$ in the figure). As illustrated by Figure 1, the pull-to-center bias is explained by IBE in all cases with high costs (defined by $z > g$) and with low costs ($z < \tilde{z}$, where $\tilde{z} \approx 0.414$ is implicitly defined by $x^{IBE}(\tilde{z}) = 1/2$). In fact, starting with Schweitzer and Cachon, almost all previous newsvendor experiments with repeated choices and feedback between rounds (those cases where IBE is applicable) that found the pull-to-center bias employed $z > g$ or $z < \tilde{z}$.

⁸ This turns out to be the so-called *golden ratio*: Two quantities are in golden ratio if the ratio of the sum of the quantities to the larger quantity is equal to the ratio of the larger quantity to the smaller one.

That said, we also find that for z close to the center, IBE does *not* necessarily predict the pull-to-center bias (see Figure 1). The underlying reason is the asymmetry due to the loss impulse, which implies that for $z \in (\bar{z}, 0.5)$ IBE predicts not only that orders are smaller than optimal orders, but that they are smaller than the center. For $z \in (0.5, g)$, IBE even predicts that the order moves into the opposite direction of the center. We found only one relevant observation in the literature with z close to the center, and we get back to this and discuss new evidence regarding predictions for z close to 0.5 in Section 3.3 below.

3.2 Asymmetry

Ho et al. (2010) emphasize that in their data the deviation from the optimal order is stronger in those cases where the optimal order is above the center than when it is symmetrically below the center. Two optimal orders are called symmetric if they have the same distance to the center (see Ho et al., p. 1895). In our setting, this observation is implied when the following inequality holds:

$$|x^{IBE}(0.5 - e) - x^*(0.5 - e)| > |x^{IBE}(0.5 + e) - x^*(0.5 + e)| \text{ for all } e \in (0, 0.5).$$

Looking at Figure 1 suggests that this is the case for the IBE prediction in our setting. Straightforward calculations confirm this. The underlying reason for the asymmetry in IBE is that losses count twice for the impulses. (Ho et al. 2010 capture the asymmetry by imposing a relatively large psychological cost parameter of overstocking.)

This kind of asymmetry in newsvendor behavior can also be found, e.g., in Becker-Peth et al. (forthcoming) as well as in our data (see below). Yet we caution that the effect appears less robust than the pull-to-center bias. In particular, Schweitzer and Cachon (2000) observe in their data that the average deviation is somewhat stronger when the optimum is below the center.

3.3 Non-linearity

The IBE order for uniformly distributed demand is a non-linear function in z , as illustrated by Figure 1. In contrast, expected profit maximization as well as the utility framework in Ho et al. (2010) imply a linear relationship for this context (this follows from equation (2) on p. 1894 in Ho et al., 2010). To our knowledge, the functional form

of x as a function of z has not been analyzed so far, and there have been no newsvendor game studies with more than three z -values in a unified framework with feedback between rounds. Thus, we conducted a new experiment to investigate newsvendor behavior across the full range of z -values.

In total we had 340 participants, each playing a newsvendor game with demand independently and with equal probabilities drawn from $\{0,1,2, \dots,100\}$ in each of 200 rounds. In our 11 treatments we fixed $p = 10$, and varied the cost c from 0 to 10 in steps of 1. Each subject faced only one z -value. Our predictions regarding the actual order as a function of z was derived before we run the experiments, as summarized in Figure 1.⁹

Figure 1 shows the normalized order for each z , averaged across rounds and subjects ("Data").¹⁰ The figure shows that IBE captures average behavior much better than expected profit maximization. Rank sum tests (based on individuals' independent average orders) yield significant differences between actual and profit maximizing orders for all z , with the exception of $z = 0.4$ and 0.5 (for $z = 0.6$ and 0.7 we have $p = 0.0024$; in all other significant cases $p < 0.0001$). At the same time, the rank sum tests yield significant differences between actual and IBE orders for $z = 0.5$ and 0.6 ($p < 0.0001$ and 0.0103 , respectively), and for both extreme values $z = 0$ and 1 , where actual behavior can deviate from predictions in only one direction. Thus, the only z -value where the IBE order is significantly different from the actual order, but the optimal order is not, is $z = 0.5$. That is, only when the optimum coincides with the center is the optimum a better predictor than IBE.¹¹ This confirms the impression from Figure 1 that IBE is a relatively good quantitative predictor in our experiment, in particular for those

⁹ The experiment took place in the Cologne Laboratory for Economic Research using ORSEE (Greiner 2004) and zTree (Fischbacher, 2007). See Appendix A for the instructions for $c = 1$; other instructions were adjusted according to the different cost parameters. The numbers of independent observations are 32, 32, 31, 30, 30, 32, 31, 31, 32, 32, 27 for $c = 0, 1, 2, \dots, 10$, respectively. The different numbers reflect differences in show-up rates.

¹⁰ To keep our exposition brief, we concentrate here on testing the non-linearity prediction. Our data are fully consistent with the phenomena described above. In particular, there is an asymmetry in our data: the absolute deviation from the optimum is generally larger when the optimum is above the center than when it is symmetrically below. Two-sided rank sum tests based on individuals' average orders yield $p = 0.0005, 0.0044, 0.0481, \text{ and } 0.0796$ for the comparisons between $z = 0.1$ and $0.9, 0.2$ and $0.8, 0.3$ and $0.7, \text{ and } 0.4$ and 0.6 , respectively. The only case where the deviation is smaller when the optimum is above the center is the comparison between $z = 0$ and 1 , where deviations are very small (as described below).

¹¹ That actual orders are optimal if the optimum equals the center is roughly in line with an observation by Bostian et al. (2008), which is the only other relevant observation with z equal to the center that we are aware of. Becker-Peth et al. (forthcoming) is another study that tests parameter values such that the mean demand equals the profit maximizing order. Yet, their study does not provide feedback between rounds, so that IBE is not applicable. In their study, the mean order for $z = 0.5$ is 0.33 .

cases where the optimum is different from the center – which are those cases that have attracted most interest of researchers.¹²

The exercise above does not rule out that another linear relationships between x and z , which may be the result of maximizing a psychologically richer utility function (Ho et al. 2010), outperforms IBE. However, when we apply a Ramsey RESET test of misspecification based on individuals' independent average orders, the null-hypothesis of a linear relationship is strongly rejected ($p < 0.0001$). More importantly, the *ex-ante* non-linear IBE yields a better fit of the data than the *best ex-post* linear fit of our data, which we obtain as the result of a simple OLS regression: the sum of squared deviations of all 340 data points from the IBE is 3.753, and the sum of squared deviations from the *best ex-post* linear fit is 4.423.¹³ The underlying reason is that actual orders exhibit a strong pull-to-center bias for intermediate values of z , and *at the same time* are close to the optimum for extreme values of z . The overall average order is with 99.54% very close to the maximum demand when orders can only yield upward impulses ($z = 0$), and it is with 5.08% close to the minimum demand when orders can only yield downward impulses ($z = 1$). No model that yields a linear relationship can simultaneously capture both phenomena.

A prominent newsvendor model that also predicts a non-linear relationship between x and z is the quantal response model by Su (2008), which allows decision makers to adopt a probabilistic choice rule such that more attractive alternatives are chosen more often. Su's (2008) analysis is based on previous work by McKelvey and Palfrey (1995), Anderson et al. (1992, 2004), and Goeree and Holt (2001, 2005), among others.

¹² Moreover, we note that IBE predicts the opposite of a pull-to-center bias for $z \in (0.5, g)$. Such pattern has not been observed before, and no other study has tested a z -value in this range before. In our experiment, we have $z = 0.6$, which falls into this range. And, indeed, for $z = 0.6$ the average order is on the opposite side of the optimum order than predicted by the pull-to-center bias. Clearly more evidence is needed for robust conclusions. However, the results seem to suggest that if we are to expect a deviation from the optimum that is inconsistent with the pull-to-center bias, we should probably expect it in the range where this is predicted by IBE.

¹³ The sum of squared deviations from the profit-maximizing order is 6.096. The same conclusion holds if we compare the sums of *absolute* deviations, which is smallest for IBE (25.918), larger for the best linear fit of our data (30.666), and again larger for the optimum (32.464). The best fit is obtained by a straightforward OLS regression based on all 340 observation, which yields a constant of 0.874 and a slope of -0.788 . The better performance of IBE is not driven by how IBE treats losses: the sum of squared (absolute) deviations from IBE predictions computed without any additional downward impulse due to monetary losses is 4.049 (27.820), and thus still smaller than the corresponding values for the best *ex post* linear fit of the data.

More specifically, Su's analysis (2008) employs a standard quantal response model, where logit choice probabilities γ , as applied to our context, are given by


$$\gamma(x) = \frac{e^{P(x)/\beta}}{\int_{y=0}^1 e^{P(y)/\beta} dy} \quad (9)$$

Here, $P(x)$ is the expected profit as a function of the order x , as defined in equation (1). The parameter β determines the extent of noise in ordering decisions: as β approaches infinity, the distribution of chosen orders approaches the uniform distribution, and as β approaches 0, the distribution of orders concentrates on the profit maximizing order. With uniform demand, and for the parameter values in our experiment (normalized so that minimum demand = 0, maximum demand = 1, $p = 1$), one can calculate the expected quantal response order x^{QR} as:

$$x^{QR} = x^* - \beta \left[\frac{f((1-x^*)/\beta) - f(-x^*/\beta)}{F((1-x^*)/\beta) - F(-x^*/\beta)} \right] \quad (10)$$

Here, $f(\cdot)$ and $F(\cdot)$ denote the standard normal density and distribution function, respectively (see Su, 2008, p. 573, for the derivation). x^{QR} is equal to the optimum x^* minus β times a term that may be positive or negative. Figure 2 gives an example of x^{QR} as a function of z and for a given parameter β .

Figure 2: Best *ex-post* quantal response prediction and behavior



The figure illustrates Su's (2008) finding that the quantal response model captures the pull-to-center bias. The underlying reason for the bias is that there is 'more room' to deviate downwards when the optimum is above the center, and 'more room' to deviate upwards when the optimum is below the center.¹⁴

However, the figure also suggests that the model does not accurately reflect the non-linear shape of actual average orders as we vary z . For one, because of the symmetry properties of the standard normal density and distribution functions, x^{QR} as defined in (10) implies that the deviation from the optimal order is the same when the optimal order is above the center and when it is symmetrically below the center. This contradicts our finding of asymmetric deviations, as defined and described in Section 3.2. More importantly, as illustrated by Figure 2, the quantal response solution (10) implies that starting at $z = 1 - x^* = 0.5$, the deviation from the optimum convexly increases as z goes to 1, and concavely increases as z goes to 0, with the respective

¹⁴ Whenever the optimum is above the center, we have $|1 - x^*| < |-x^*|$, so that $f((1 - x^*)/\beta) > f(-x^*/\beta)$, which implies that $x^{QR} < x^*$. A similar argument holds for the case that the optimum is below the center (see Su 2008, p. 586).

maximal deviations for the most extreme values of z (0 and 1). One reason is that deviations from the optimum are least costly if they can only bring orders more closely to the center.

To the contrary, IBE predicts that orders coincide with the optimum for extreme values of z , because there are then no impulses that could justify deviations. In fact, as pointed out before, deviations are small for extreme z -values. This contributes to the observation that the *ex-ante* IBE prediction also outperforms the best *ex-post* fit of the non-linear quantal response model (which is shown in Figure 2): the minimal sum of squared deviations of all 340 data points from the quantal response prediction is 5.230 (3.753 for IBE).¹⁵

4. Discussion, limitations and conclusion

We complement Arrow et al.'s (1951) normative model by proposing impulse balance equilibrium as a descriptive model for newsvendor behavior. The model predicts the pull-to-center bias, as well as other, more subtle phenomena observed in laboratory experiments.¹⁶ Our model builds upon Schweitzer and Cachon's (2000) idea in their classic paper on what drives newsvendor behavior, a concern for *ex-post* errors: Impulse balance equilibrium quantifies the impulses generated by *ex-post* errors and, in equilibrium, balance them out.

Our model is probably closest related to Ho et al.'s (2010) model. In both models, realized demand serves as an anchor that guides behavior. Deviations from demand generate impulses or psychological costs, which depend on the direction and size of the deviation. Both models capture important phenomena, including the pull-to-center bias

¹⁵ The β in equation (10) that minimizes squared deviations is 0.19. Our statements above assume that the price p is fixed for all z . A price increase would increase the stakes and thus reduce deviations from the optimum (Su 2008, p. 573). In our experiment, the price was fixed only nominally, in terms of the 'Experimental Currency Unit (ECU)'; exchange rates between ECU and EURO differed across treatments (see Appendix A). However, exchange rates cannot explain the failure of the quantal response solution to predict behavior for extreme z -values. The reason is that for $z < 0,5$ and $z > 0,5$, *real* prices are lowest for $z = 0$ and 1, respectively, so that quantal response would suggest large deviations for extreme z -values even if we take real prices into account.

¹⁶ In a companion paper (Ockenfels and Selten, mimeo), we extend the *ex post* rationality principle to standing orders and multiple-round feedback. In particular, we show that IBE also captures that constraining newsvendors to make a standing order for a sequence of periods highly significantly moves the average of submitted orders toward the optimum, as observed by Bolton and Katok (2008; Bostian et al. 2008 and Lurie and Swaminathan 2009 report similar findings).

and an asymmetry in the data. However, it is useful to emphasize the key differences in the underlying approaches, which make them complementary.

For one, Ho et al. assume that deviations from the optimum are due to *motivational* limits: the psychological costs are parameters of the newsvendors' utility function, suggesting that newsvendors deviate because they do *not like* deviations from actual demand – even when this is costly. This captures that the treatment effects observed in newsvendor experiments sometimes seem to materialize already in the very first rounds, and thus may not only rely on adaptation. On the other hand, Ho et al.'s model does not capture that there is substantial adjustment towards previous demand in newsvendor experiments as emphasized by various researchers. Moreover, a study by Kremer et al. (2011) suggests that judgmental time-series forecasting in situations that include newsvendor environments place too much emphasis to signals they receive relative to the system that generates the signals. Also, a newsvendor study by Bolton et al. (2012) suggests that if subjects are explicitly told the optimal solution, newsvendors are much closer to the optimum. This kind of evidence suggests that the biases are not (only) caused by motivational but (also) by *cognitive* limitations. IBE is an attempt to complementarily capture such cognitive limitations.

One advantage of IBE (and standard theory) is that it does not require parameter estimations: IBE yields predictions without previous knowledge about newsvendor behavior. The fact that IBE outperforms the best *ex-post* linear fit of our newsvendor data as well as the best *ex-post* fit of a non-linear quantal response model, and that it has been shown to perform well also across many other games, increases our confidence that the underlying mechanism captures an important driver of behavior.

There are also important limitations of the predictive value of our IBE model. One is that our IBE model cannot predict the degree of heterogeneity in laboratory newsvendor behavior: While heterogeneity of newsvendor behavior appears to be a robust phenomenon across experiments (e.g., Moritz et al. 2011; Figure B.1 in Appendix B shows that there is also non-negligible heterogeneity of average orders in our experiment), our IBE's prediction is only concerned with the central tendency of the stationary distribution of the newsvendors' order decisions, and neither with the variance in individual orders nor with the distribution of average orders across individuals. However, we note that IBE is not *per se* inconsistent with individual

differences. Ordering is affected by impulses, which depend on the specific path of orders and demand realizations, which can differ across individuals. Future research should look at the kind of order distributions one might expect to see with impulse balancing newsvendors, and whether incorporating additional sources of noise is necessary to capture observed heterogeneity.¹⁷

IBE (like other equilibrium models) also does not predict the dynamics observed in newsvendor behavior. In particular, many newsvendor game studies starting with Schweitzer and Cachon (2000) report demand chasing behavior. That is, if newsvendors adjust their order quantity from round-to-round, they are more likely to adjust their order toward prior demand than away from prior demand. However, this observation supports the cognitive mechanism that is *assumed* by IBE: newsvendors adjust behavior towards what would have been the better choice in the previous round. In this sense, demand chasing is consistent with IBE's notion of *ex-post* rationality.

Finally, many newsvendor game studies observe that there is only a very slow convergence, if at all, toward the optimum when the newsvendor game is played over many rounds (Schweitzer and Cachon 2000, Bolton and Katok 2008, Bolton et al. 2012, among others). While deviations from the optimum in our experiment are mostly persistent, even after 200 rounds, there is some general trend toward the optimum in our data (see Figure B.2 in Appendix B). Specifically, the average orders for all eleven z -values move into the direction of the profit maximal order when we compare behavior in the first 100 and the second 100 rounds, while in only seven of those cases the change is also consistent with moving into the direction of the IBE prediction (see Table B.1 in Appendix B). This, together with the round-by-round dynamics, shows that a full picture of newsvendor behavior cannot neglect the adaptive nature of boundedly rational ordering, such as analyzed by Bostian et al. (2008).

Behavioral approaches such the Cognitive Hierarchy Model by Camerer, Ho and Chong (2004) and the Analogy-based Expectation Equilibrium by Jehiel (2005) have demonstrated that models that attempt to capture 'bounded economic cognition' can systematically and robustly capture behavior in a variety of situations. In this paper we

¹⁷ That said, IBE cannot capture pure presentation effects, as observed by Kremer et al. (2010). Here, models like Ho et al.'s (2010) motivational model and Su's (2008) quantal response model of noisy decision making can be useful to measure and organize differences in behavior by allowing differences in parameter values (although, these models, too, typically cannot predict differences *ex-ante*).

add to this rather small literature by showing that the very parsimonious IBE model successfully predicts various non-trivial behavioral phenomena. The way IBE captures cognitive limitations does not suppose that behavior is dumb or just noisy. Rather, subjects follow their own, *bounded* rationality. To us, using mathematical tools to describe how limited cognition affects information processing and decision making is a promising endeavour in current behavioral research, complementing the more standard approaches based on limited motivation and pure adaptation, and ultimately helping us to design better decision support systems and to make better decisions.

References

- Abdellaoui, M., H. Bleichrodt, and C. Paraschiv (2007). Loss Aversion Under Prospect Theory: A Parameter-Free Measurement. *Management Science* 53 (10): 1659-1674.
- Anderson, S. P., A. de Palma, and J. F. Thisse (1992). *Discrete Choice Theory of Product Differentiation*. MIT Press, Cambridge, MA.
- Anderson, S. P., J. K. Goeree, and C. A. Holt (2004). Noisy directional learning and the logit equilibrium. *Scandinavian Journal of Economics*, 106(3), 581–602.
- Arrow, K. J., T. Harris and J. Marschak (1951). Optimal inventory policy. *Econometrica* 19 (3) 250-272.
- Avrahami, Judith, Güth Werner, and Yaakov Kareev (2005). Games of Competition in a Stochastic Environment. *Theory and Decision*, 59(4): 255–94.
- Becker-Peth, Michael, Elena Katok, and Ulrich Thonemann (forthcoming). Designing Contracts for Irrational but Predictable Newsvendors. *Management Science*.
- Benzion, U., Y. Cohen, R. Peled and T. Shavit (2008). Decision-making and the newsvendor problem—an experimental study. *Journal of the Operational Research Society*, 59, 1281-1287.
- Bolton, Gary E., and Axel Ockenfels (2012). Behavioral economic engineering. *Journal of Economic Psychology*, 33 (3), 665-676.
- Bolton, Gary E., Axel Ockenfels, and Ulrich Thonemann (2012). Managers and Students as Newsvendors. *Management Science*, 58(12), 2225-2233.
- Bolton, Gary E. and Elena 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.
- Bostian, AJ A., Charles A. Holt and Angela M. Smith (2009). Newsvendor “Pull-to-center” effect: Adaptive Learning in a Laboratory Experiment. *Manufacturing & Service Operations Management*, 10(4), 590-608.
- Brunner, C., C. Camerer, and J. Goeree (2011). Stationary Concepts for Experimental 2×2 Games: Comment, *American Economic Review* 101(2), 1029-1040.
- Camerer, Colin F, Teck-Hua Ho, and Juin-Kuan Chong (2004). A cognitive hierarchy model of games. *Quarterly Journal of Economics*, 119, 861–898.
- Vincent P. Crawford (2013). Boundedly Rational versus Optimization-Based Models of Strategic Thinking and Learning in Games. *Journal of Economic Literature* 51:2, 512-527.
- Feng, Tianjun, L. Robin Keller and Xiaona Zheng (2011). Decision making in the newsvendor model: A cross national laboratory study. *Omega* 39, 4-50

- Fischbacher, U. (2007). z-Tree: Zurich Toolbox for Ready-made Economic Experiments. *Experimental Economics*, 10(2): 171–178.
- Goeree, J.K. and C.A. Holt (2001). Ten little treasures of game theory and ten intuitive contradictions. *American Economic Review* 91, 1402–22.
- Goeree, J.K. and C.A. Holt (2005). An explanation of anomalous behavior in models of political participation. *American Political Science Review* 99, 201–13.
- Greiner, B. (2004). An Online Recruitment System for Economic Experiments. In: Kurt Kremer, Volker Macho (eds.): *Forschung und wissenschaftliches Rechnen 2003*. GWDG Bericht 63, Göttingen: Ges. für Wiss. Datenverarbeitung, 79-93.
- Ho, T., N. Lim, T. Cui (2010). Reference dependence in multilocation newsvendor models: A structural analysis. *Management Science*, 56(11), 1891-1910.
- Hollander, M., and D. A. Wolfe (1999). *Non-parametric Statistical Methods*. Second Edition, Wiley.
- Philippe Jehiel (2005). Analogy-based expectation equilibrium. *Journal of Economic Theory*, 123(2), 81–104.
- Kremer, Mirko, and Stefan Minner (2008). The human element in inventory decision making under uncertainty: A review of experimental evidence in the newsvendor model, *Zeitschrift für Betriebswirtschaftslehre*, 4, 83-97.
- Kremer, Mirko, Stefan Minner, Luk N. Van Wassenhove (2010). Do random errors explain newsvendor behavior?. *Manufacturing & Service Operations Management*, 12(4), 673-681.
- Kremer, Mirko, Brent Moritz, and Enno Siemsen (2011). Demand Forecasting Behavior: System Neglect and Change Detection. *Management Science*, 57(10), 1827–1843.
- Lau, Nelson, and J. Neil Bearden (forthcoming). Newsvendor Demand Chasing Revisited." *Management Science*.
- Lurie, Nicholas H. and Jayashankar M. Swaminathan (2009). Is timely information always better? The effect of feedback frequency on decision making. *Organizational Behavior and Human Decision Processes*, 108(2), 315-329.
- McKelvey, Richard D., and Thomas R. Palfrey (1995). Quantal Response Equilibria for Normal Form Games. *Games and Economic Behavior*, 10(1), 6–38.
- Moritz, B., A.V. Hill and K. Donohue (2011). Individual differences in the newsvendor problem: Behavior and cognitive reflection. Working paper.
- Ockenfels, Axel, and Reinhard Selten (2005). Impulse Balance Equilibrium and Feedback in First Price Auctions. *Games and Economic Behavior*, 51(1), 155–70.
- Ockenfels, Axel, and Reinhard Selten (mimeo). Impulse Balance Equilibrium and Multiple-period Feedback in the Newsvendor Game, work in progress.

- Ren, Y. and R.T.A. Croson (2012). Explaining biased newsvendor orders: An experimental study. Working paper, University of Texas Dallas.
- Schweitzer, Maurice E. and Gérard P. Cachon (2000). Decision Bias in the Newsvendor Problem with a Known Demand Distribution: Experimental Evidence. *Management Science*, 46(3), 404-420.
- Selten, Reinhard, Klaus Abbink, and Ricarda Cox (2005). Learning Direction Theory and the Winner's Curse. *Experimental Economics*, 8(1): 5-20.
- Selten, Reinhard, and Joachim Buchta (1999). Experimental Sealed Bid First Price Auctions with Directly Observed Bid Functions. In *Games and Human Behavior: Essays in the Honor of Amnon Rapoport*, ed. David Budescu, Ido Erev, and Rami Zwick, 79-104. Mahwah NJ: Lawrence Erlbaum Associates.
- Selten, Reinhard and Thorsten Chmura. (2008). Stationary Concepts for Experimental 2x2-Games. *American Economic Review* 98(3), pp. 938-66.
- Selten, Reinhard, Thorsten Chmura, and Sebastian J. Goerg, (2011). Correction and Re-examination of Stationary Concepts for Experimental 2x2 Games: A Reply. *American Economic Review*, 101(2), 1041-1044.
- Shanks, D. R., R. J. Tunney, and J. D. McCarthy (2002). A re-examination of probability matching and rational choice. *J. Behavioral Decision Making* 15 233-250.
- Su, X. (2008). Bounded rationality in newsvendor models. *Manufacturing and Service Operations Management*, 10 (4), 566-589.
- Wu, D.Y and K.-Y. Chen (2012). Supply Chain Contract Design: Impact of Bounded Rationality and Individual Heterogeneity. Working paper.

Appendix A: Instructions for $c = 1$ (translation from German)¹⁸

Instructions

You will be able to earn money during this experiment. The amount you will earn depends among other things on your decisions taken over the course of this experiment.

The experiment consists of 200 rounds. During each round one decision is to be taken. All rounds are payoff relevant. During the experiment an Experimental Currency Unit (ECU) will be used. At the end the sum of all ECU-amounts is converted into EURO. The exchange rate is 4,400 ECU = 1 EUR. Additionally, each participant will receive a show-up fee of 2.50 EUR.

Please turn off your cellphone and abstain from communicating with other participants from now on. Please turn your full attention to the experiment. Please raise your hand if you have any questions concerning the experiment. We will come to you and answer your question.

All decisions made during the experiment as well as all payments at the end will be kept anonymous. Please also abstain from discussing these with other participants after the experiment.

The decision situation

You are a retailer who offers a single generic product. Each period of the game, you will order the quantity required of the product from an external supplier, in order to sell it on to the costumers. You can order any integer quantity between 0 and 100 (boundaries included).

You will pay the external supplier 1 ECU for each unit ordered. You will receive 10 ECU for each unit demanded.

Each period when placing your order, you do not know the customer demand for the respective period. You do however know that demand will lie within a certain interval. The computer will randomly generate the demand. In each round, every integer quantity between 0 and 100 (boundaries included) is equally likely.

¹⁸ With c we also changed the exchange rate of the experimental currency unit (ECU) to Euros to smooth differences in expected profits (1€ = 5400, 4400, 3500, 2700, 2000, 1400, 900, 500, 220, and 55 ECU for $c = 0, 1, 2, \dots, 9$, respectively; for $c = 10$, subjects could not increase payoffs beyond the fixed payment, so we added an initial endowment of 16.50 Euros to the show-up fee, which was 2.50 Euros in all treatments, and implemented an exchange rate of 1€ = 2000).

Profit calculation per round

Once you have placed your order, demand will be filled and your profit calculated. You will then receive information on customer demand, quantity sold, and your profit. Additionally, we inform you about how high your profit would have been would you have ordered exactly the quantity demanded.

Your payoff is calculated as follows: First, your profit for the units sold is calculated (10 ECU revenue minus 1 ECU costs per unit). Then the costs for units ordered in excess of demand are subtracted from your profit.


Please be aware, that you can also make a loss. Should you have accumulated losses after the 200 rounds, these will be set against your show-up fee of 2.50 EUR.

The end of the experiment

The results of all rounds are added up after the last round, converted to EURO and paid out to you in cash, including the show-up fee.

Appendix B: Heterogeneity and dynamics of newsvendor behavior

Figure B1: Heterogeneity of average orders




The boxplot shows, for each z , the mean and the quartiles of the respective individual orders (which are averaged across the 200 rounds). Dots denote outliers (individual averages, which are more than two standard deviations away from the mean), which are not taken into account for the computation of the boxplots.

Table B1: Predictions and dynamics of average orders

x^*	IBE	All rounds	Rounds 1-100	Rounds 101-200
1.000	1.000	0.995	0.993	0.998
0.900	0.741	0.727	0.682	0.772
0.800	0.646	0.645	0.624	0.667
0.700	0.573	0.603	0.571	0.635
0.600	0.509	0.563	0.549	0.578
0.500	0.449	0.506	0.508	0.503
0.400	0.392	0.369	0.367	0.372
0.300	0.334	0.332	0.332	0.331
0.200	0.271	0.295	0.333	0.256
0.100	0.195	0.184	0.197	0.170
0.000	0.000	0.051	0.067	0.035

Figure B2: Dynamics of average orders



The lines show, for each z (0, 0.1, ..., 1) the average order across all subjects in the respective treatment. We averaged orders across rounds of 10, so that period 1 in the figure corresponds to rounds 1-10 in the experiment, period 2 corresponds to rounds 10-20, etc.