

Learning-by-Doing in the Newsvendor Problem:

A Laboratory Investigation of the Role of Experience and Feedback

Gary E Bolton and Elena Katok
Smeal College of Business, Penn State University

We investigate learning-by-doing in the newsvendor inventory problem. An earlier study observed that decision makers tend to anchor their orders around average demand and fail to adjust sufficiently towards the expected profit-maximizing order. Principles of behavioral theory suggest some relatively simple interventions into the decision maker's experience and feedback that might improve performance, and these guide our investigation. The results imply that the institutional organization of experience and feedback may have a significant influence on whether inventory is stocked optimally.

Keywords: Newsvendor problem, feedback, behavioral operations management, supply chain management, experimental economics.

1. Introduction

1.1 Background

The newsvendor problem is a fundamental building block for models of inventory management in the face of stochastic demand (Porteus 1990), and at a broader level for models of supply chain systems (Cachon 2002). The newsvendor's problem is that he must stock his entire inventory prior to the selling season, knowing only the stochastic distribution from which the quantity demanded will be drawn; order too little, and he loses sales, order too much, and he must dispose of the excess stock at a loss. The organizational prescriptions that flow from newsvendor-based models typically assume that the newsvendor stocks optimally. Behavioral studies of the newsvendor problem, however, find that people often make suboptimal and biased choices. Schweitzer and Cachon (2000) found this to hold even for those who have been exposed to the optimal solution in an MBA classroom. They argue that, "new techniques may be required to optimize these systems. p. 420."

The experiments presented here investigate whether enhancements to experience and feedback can facilitate better newsvendor learning-by-doing. That classroom exposure falls short is not entirely surprising: Choices under uncertainty are known to be prone to a number of

judgment biases (Gilovich, Griffin and Kahneman 2002). To give an example with relevance here, even trained scientists are prone to the ‘law of small numbers’ bias, drawing conclusions on the basis of inappropriately small samples (Tversky and Kahneman 1971). That said, people can—given the right experience and feedback—learn to solve stochastic choice problems that resemble the newsvendor problem.¹ Principles of behavioral theory suggest factors that might curb judgment biases. These principles guide our study.

An investigation of this kind is a first step in a program to establish a more nuanced understanding of the organizational features that promote optimal behavior. Classic work in operations and inventory control suggests the potential gains from such a program. Based on production and scheduling problems reported by a number of firms, Holt, Modigliani, Muth and Simon (1960) derived an optimal linear decision rule for aggregate planning. Bowman (1963) showed that linear decision rules estimated from managers’ past decisions can outperform this rule—as well as the managers themselves. He speculated that experienced managers have refined information about their own operations not available to outsiders and make good decisions on average, but that they also tend to over-react to short-term fluctuations (the latter foreshadowing our findings). A carefully designed institution might curb short term over-reaction without impeding the manager’s ability to exercise valuable long-term judgment.

More recent studies also identify shortcomings where biased decision making likely plays a role. Firms are found to hold the wrong products in wrong locations or to over-react to random demand fluctuations (Lee, Padmanbhan and Whang 1997), to misunderstand inherent system delays (Sterman 2000) or to systematically order too much (Katok, Lethrop, Tarantino and Xu 2001). The firms in these studies include Hewlett-Packard, Procter & Gamble and IBM.

Boudreau, Hopp, McClain and Thomas (2003) make the broad case for pushing at simplifying behavioral assumptions, arguing that “Once a feature of human behavior has been recognized, incorporating it into the analysis can lead to better OM models” p. 185.² The same feature that makes the newsvendor problem an important theoretical paradigm—that it captures the critical decision parameters common to most inventory decisions—also makes it a promising

¹ Consider, for example, the probability learning problem in which a person guesses which door contains a prize. Each door contains the prize with a fixed but unknown probability. The optimal choice is the door with the highest probability. With experience—and material incentives—people solve the problem effectively (ex., Siegel (1964), Holt (1992), Shanks et al. (2002)). The newsvendor problem differs in that the prize is stochastic and that the option that maximizes expected profit may not be the risk averse choice. Both potential culprits are examined below.

² For example, Schultz, Juran, Boudreau, McClain and Thomas (1998) find that processing times in serial production systems depend on inventory levels, with the amount of idle time less than traditional theory suggests.

test case for identifying broadly applicable behavioral insights. By the same token, it is a logical starting place in a program to investigate, in dialogue with theory, more complex settings. This building-up-with-theory approach is common to experimental economics (Kagel and Roth 1995). Papers in a special issue of *Interfaces* describe how this approach has been applied to a number of institutions and problems in business practice (Bolton and Kwasnica 2002).

1.2 A Previous Experiment and an Overview of Our Study

Schweitzer and Cachon (2000) is the first laboratory study of the newsvendor problem. They examined both a high and low safety stock condition in which the optimum inventory order was above (below) average demand. The game was repeated and subjects were provided feedback on realized demand and profitability at the end of each round. The data showed that subjects, including second year MBAs who received classroom training in the newsvendor solution in their first year, “consistently ordered amounts lower than the expected profit-maximizing quantity for high-profit products and higher than the expected profit-maximizing quantity for low-profit products.” The authors demonstrate that,

This too low/too high pattern of choice cannot be explained by risk aversion, risk-seeking preferences, loss avoidance, waste aversion, or underestimating opportunity costs. Preferences consistent with Prospect Theory (risk aversion over gains and risk seeking over losses) can explain some, but not all, of the data in our experiments. P. 418

Hence Schweitzer and Cachon observed a pattern of behavior that is at odds with expected profit maximization as well as with alternative risk profiles. The authors note that one explanation for the data is anchoring and insufficient adjustment; that is, subjects “anchor” around average demand in the early rounds of the game and insufficiently adjust in subsequent rounds towards the expected profit-maximizing order quantity. In essence, they fail to learn.³ Two new studies evidence the robustness of Schweitzer and Cachon’s results (as will ours). Benzion et al. (2005) vary the demand distribution and find that orders are affected by both the average demand and the demand in the previous round; this bias weakens slowly over time, but not enough to move newsvendors to optimal behavior. Lurie and Swaminathan (2005) find that more frequent feedback sometimes actually degrades performance and slows down learning.⁴

³ A second explanation Schweitzer and Cachon identify, ex-post inventory error minimization, tends to pull orders towards the average demand. We will see this behavior pattern in our data as well.

⁴ Earlier studies by Rapoport (1966), (1967) are suggestive of the same pattern of behavior. He found that decision-makers in a stochastic multistage inventory task generally under-control the system, and while demand draws are independent, orders are correlated with past demand.

Anchoring and insufficient adjustment are consistent with two robust findings from the behavioral decision literature, both dating to the early days of the field; these are, first, that, compared to the complexity of many tasks, people have limited information processing capacity, and, second, that people are adaptive (ex., Hogarth (1987) and references therein). Our experiment investigates modifications to feedback and experience known to improve adaptation or information processing in other contexts. The focus of the experiments is newsvendor performance, and we will take measures of performance, such as profitability, as a proxy for learning (that is, learning is presumed to be indicated by changes in performance levels).

The experiment is organized into three studies, summarized in Table 1. The hypotheses are fully developed, with appropriate references, as they are investigated below. Study 1 focuses on two hypotheses concerning adaptive learning. Schweitzer and Cachon provided subjects with 30 rounds of experience. We provide *extended experience* of 100 rounds to see what improvement this might make. A second hypothesis concerns the relatively *flat maximum* at the peak of the newsvendor’s expected profit function. To sharpen payoff differentials, we thin the set of ordering options from 100 to 9 or 3.

Table 1: Roadmap to the three studies. Each row of the table describes a study treatment.

<i>Study</i>	<i>Treatment label</i>	<i>Order Options</i>	<i>Tracking Information</i>	<i>Operational Interventions</i>
1	100-option	100		
1	9-option	9		
1	2 3-option	3		
	2 MAVG	3	Forgone	Moving average
	2 3 FORE	3	Forgone	
	3 10P	3	Forgone	Standing order
	3 UPFRONT	3	Forgone	Upfront info

Forgone = information about payoff for each option (including those not taken)

Moving average = information about 10-round moving average payoff for each option (low safety stock only)

Standing order = standing order restriction for 10 demand periods

Upfront info = expected profit information for each option provided prior to play (low safety stock only)

Study 2 focuses on improving forward looking learning. The study introduces, in the context of the 3-option design, *tracking information* on the profit of both forgone and taken decisions (FORE). Another treatment adds information on 10-round moving averages (MAVG). Study 3 takes a more invasive approach. There is evidence in our first two studies that newsvendors fall victim to the law of small numbers. We constrain newsvendors, in the context of the 3-option design, to making standing orders, fixed for 10 demand-periods at a time (10P). As a

comparison, we conduct a treatment in which newsvendors order for one demand-period at a time but are given descriptive statistics at the beginning of the session (UPFRONT).⁵

Section 2 reviews preliminaries concerning the newsvendor problem and the experiment. Sections 3 through 5 describe the studies, including formulating specific hypotheses and summarizing the supporting literature. Section 6 summarizes main results. Section 7 draws conclusions and discusses managerial implications.

2. Laboratory Implementation of the Newsvendor Problem

In this section, we describe the implementations and methods common to all three studies. Features specific to individual studies are discussed in later sections, as each study is introduced.

2.1 The Newsvendor Problem and Solution

The newsvendor must place an order q before knowing the actual demand, D . The set of feasible order quantities will be a variable in our experiment. Each unit is sold at a price p and costs c . If the amount ordered, q , exceeds D , then exactly D units are sold, and $q - D$ units are discarded. If D exceeds q then q units are sold and potential profit from selling $D - q$ units is forgone. Additionally, our setting includes a fixed rent of R that is subtracted from the total profit each round. If D is a random variable with distribution function F and density function f , the profit when q is ordered and the demand is D can be written as

$$\pi(q, D) = p \min(q, D) - cq - R$$

and expected profit is

$$E[\pi(q, D)] = \int_0^q f(x) \pi(q, x) dx + (p - c)q(1 - F(q)) - R(1 - F(q)).$$

It is well-known that the order quantity q^* that maximizes the expected profit must satisfy

$$F(q^*) = \frac{p - c}{p}.$$

For the remainder of this paper we refer to q^* as the *optimal order*.

⁵ We use the ‘round’ to refer to a decision making opportunity. Thus, a round corresponds to a demand period in all treatments with the exception of 10P. In the 10P treatments a round corresponds to 10 demand periods.

2.2 Laboratory Design: All Studies

Our experiment considers both a high and low safety stock condition. In the *low safety stock condition* $p = 12$, $c = 9$, $R = 50$, and $D \sim U(50,150)$ and integer, implying an optimal order of 75. In the *high safety stock condition* $p = 12$, $c = 3$, $R = 200$, and $D \sim U(0,100)$ and integer, also implying an optimal order of 75.⁶ All monetary quantities are in units of laboratory francs. We vary the fixed cost R , so that expected monetary payments to subjects are similar across conditions (ex., the expected total profit from placing the optimal order of 75 is \$13.60 in the low safety stock condition and \$14.20 in the high safety stock condition).

A total of 234 people participated in the experiments we report here. Each subject participated in exactly one session. Cash was the only incentive offered. Subjects were students, mostly undergraduates, from various fields of study, recruited through a computerized recruitment system. The one exception, the 100-option high safety stock treatment in study 1, was conducted with executive MBA students as part of a class. This treatment permits us to benchmark our other results against a subject pool with management experience.

Sessions were conducted at the Laboratory for Economic Management and Auctions (LEMA) at the Smeal College of Business, Penn State University, except for the executive MBA session conducted in the classroom. Prior to the study, we piloted the experiment in undergraduate and MBA classrooms, with subjects debriefed both before and after play on the clarity of instructions and software. Subjects first read instructions (on-line Appendix A1). After completing the instructions we invited subjects to ask questions and answered any questions that were asked before starting the session.

In all variations of the problem we consider, there are 100 consecutive inventory ordering decisions. Each round of the game began with the participant choosing an order quantity, after which the customer demand and realized profit were revealed. In all treatments except 100-option high safety stock, subjects faced the same sequence of demand draws (randomly drawn prior to the experiments). A snapshot of a typical newsvendor computer screen appears in on-line Appendix A2. The screens displayed information about p , c , R and the demand distribution, as well as historical information about the outcomes in prior rounds of the game, including demand realization, the order placed, and the resulting profit, as well as the current total profit

⁶ The average of 100 demand draws in the low safety stock condition was 100.2 and in the high safety stock condition it was 50.2 (mean of uniform distribution is 100 and 50, respectively). Standard deviation of demand draws was 27.7 in both conditions (standard deviation of a uniform distribution with a range of 100 is 28.9).

accumulated since the start of the session. The experiment's software was built from Microsoft™ Access with Visual Basic for Applications (VBA), and MySQL database software.

At the end of the session subjects were paid, in private, their total individual earnings from 100 decisions at a rate of 1000 lab francs = \$1 (to maintain comparability in value per decision, in the 10P treatments of study 3, 10000 lab francs = \$1; see section 5.2). Sessions lasted between 30 and 45 minutes. Actual average earnings, including a \$5 participation fee, were about \$17.

2.3 Metrics for the Analysis

Beyond the number of units ordered, we analyze the behavior in our experiments using two additional metrics. *Proportion of maximum expected profit achieved:* Our primary focus is on the financial optimality of choices. Focusing on *expected* profitability reduces the role of luck in comparisons. To compute the proportion of maximum expected profit captured by an order decision, we calculate the associated expected profit, and divide the result by the expected profit from the optimal order. *Search Pattern:* Averages and standard deviations can sometimes mislead about behavioral heterogeneity (Juran and Schruben 2004). Analysis of search patterns of individual newsvendors provides clues on the cognitive processes behind decisions.

3. Study 1: Extended Experience and Flat Maximums

3.1 Hypotheses

Relative to Schweitzer and Cachon's study, our 100-option treatments extend the length of the session from 30 to 100 rounds. There is a good deal of evidence that increased experience can lead to more frequent optimal behavior. For example, Siegel (1964) demonstrates substantial learning towards the optimum in the probability learning problem (see footnote 1), with learning continuing even after 50 rounds of experience. Prasnikar and Roth (1992) show that, with experience, people play the optimal strategy in a "best shot" public goods game; Roth and Erev (1995) show that adaptive learning implies that experience leads to optimal play in this game. Given these results, it seems plausible that extended experience might help the newsvendor.

H1: Experience hypothesis. *With increasing experience, newsvendors make decisions that achieve a higher proportion of maximum expected profit.*

At the same time, the newsvendor problem has a characteristic that may make it particularly difficult for experience to overcome anchoring and insufficient adjustment: From Figure 1, the newsvendor’s expected profit function is flat around the neighborhood of the maximum. Flat maximums impede learning in other kinds of games: Siegel and Fouraker (1960) found that small payoff differences between contracts made it less likely that bargainers would settle on the optimum contract. Harrison (1989) showed that the flat maximum problem extends to *expected* payoff differences between bidding strategies in first-price auctions; when these differences are increased, learning and performance improve. Erev and Roth’s 1998 adaptive learning model provides a theoretical explanation for these effects. In their model, sharper differences in expected payoffs provide greater differential reinforcement, and this leads to faster adoption of the highest payoff option. It seems plausible then, that the flat maximum in Figure 1 might impede whatever positive influence experience has on newsvendor decisions.

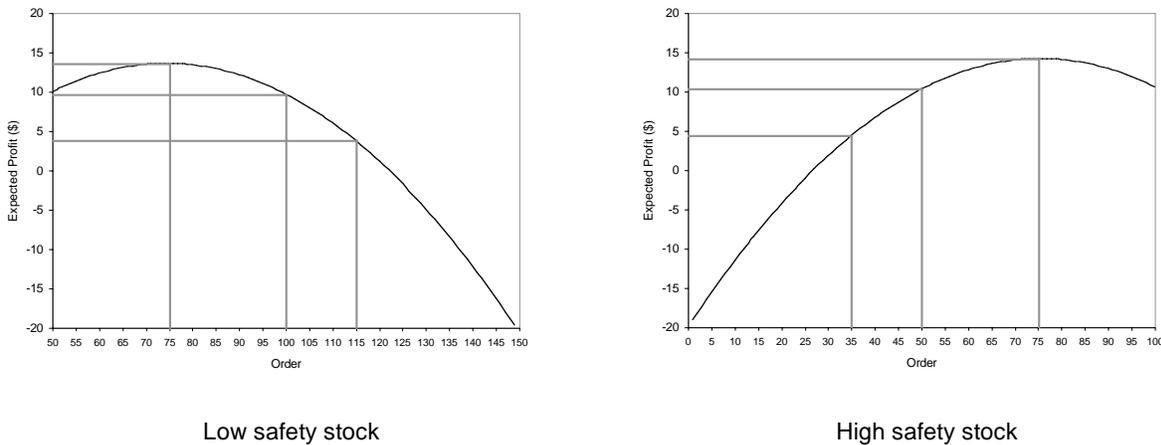


Figure 1: Expected profit as a function of order quantity. The gray lines mark the three order options in the 3-option treatments of study 1, where the optimal order is 75, and the middle option corresponds to the average demand.

A second, related reason that the flat maximum might slow learning is that actual profit draws for any given order quantity tend to be quite variable around the expected profit. This variability tends to make performance comparisons between ordering quantities with similar expected profitability unreliable; individual draws, for example, may differ in the opposite direction from that of expected profitability, thereby delivering the wrong lesson.

A simple way of testing the flat maximum hypothesis is to thin the ordering space to a sparser set of options more or less equidistant from one another (in this sense, maintaining the

representation of the entire space). We compare treatments in which decision makers have 100 ordering options to treatments with 9 and 3 ordering options (see Table 1). Thinning the space in this way both sharpens expected payoff differences between “neighboring” order quantities, and makes comparison of draws for neighboring order quantities more reliable. Both effects might lead to better performance.

H2: Flat maximum hypothesis. *Thinning the set of order options leads to newsvendor decisions that achieve a higher proportion of maximum expected profit.*

3.2 Method

We implement 100, 9 and 3-option treatments for both high and low safety stock conditions. For the low safety stock condition, the 100 options were the integers from 51 to 150, and the 9 options were 75, 80, 85, 90, 95, 100, 105, 110, and 115. In the 3-option case, newsvendors chose between 75, 100 and 115; the latter because 125, the midway between 100 and 150, yields an easily detected negative expected profit, making for what is effectively a 2-option game. For similar reasons we avoided the order quantities at the end points of the demand range.⁷ The high safety stock cases are analogous: the 100 options were the integers from 1 to 100, the 9 options were 35, 40, 45, 50, 55, 60, 65, 70, and 75, and the 3 options were 35, 50 and 75. Three is the minimum number of options that satisfy criteria important to distinguishing behavioral hypotheses: For example, distinguishing between average demand matching and a preference for minimizing variability requires a third option. Also, the optimal order is an extreme choice in the 3 and 9-option treatments but not in the 100-option treatment.

The 100-option, high safety stock treatment used executive MBAs, and interpreting these results needs to take into account this difference among subject groups. All treatments had 20 subjects, save 100-option, high safety stock which had 18. Procedures common to all three studies are stated in section 2.2.

3.3 Results: Experience Hypothesis and the 100-Option Treatments

Results for the 100-option treatments, both low and high safety stock conditions, are plotted in Figure 2. We observe the same pattern of anchoring and insufficient adjustment

⁷ So a change in performance could conceivably be attributable to range reduction instead of to the reduction in the number of options. As we will see, the implemented reduction has little effect on performance, rendering the point moot.

reported by Schweitzer and Cachon: Recalling that 75 is the optimal order in both conditions, plots for both conditions show average orders falling between average demand and the optimal order. The aggregate average in the low safety stock condition is 88 and in the high safety stock condition it is 61, both significantly different from 75 (both Wilcoxon, two-tailed $p < 0.001$). Eeckhoudt et al. (1995) demonstrate that risk aversion, preferring the expected value of a gamble to taking the gamble, implies that newsvendors optimally order less than the expected profit maximizing choice. So risk aversion would pull orders below 75 in *both* conditions, contrary to what we see in Figure 2.

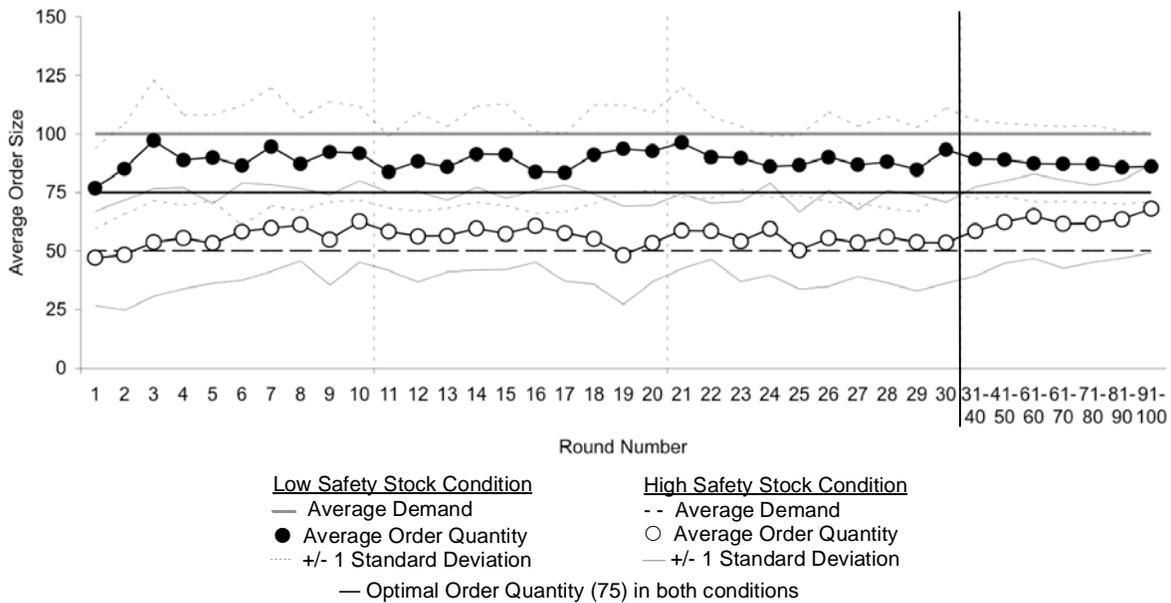


Figure 2: Order quantities by round. For rounds 1-30 the orders are averaged over subjects. For the remaining rounds they are also averaged over 10-round blocks.

Fitting a simple trend line to the data in Figure 2 (where each data point corresponds to the average order per round) we find a trend in the direction of the optimal order for both high and low safety stock conditions (OLS, two-tailed $p < 0.001$ in both cases), consistent with the extended experience hypothesis. The trends, however, emerge slowly: The overall average order increase in the High Safety Stock condition is 0.126 units per round (standard error = 0.012), and the overall average order decrease in the Low Safety Stock condition is 0.038 units per round (standard error = 0.011). Like in Schweitzer and Cachon, little trend is apparent when attention is restricted to the first 30 rounds (OLS two-tailed $p > 0.500$ in both safety stock conditions).

When we compare the magnitude of the trend between the low and the high safety stock conditions, we find that the trend variable is somewhat more pronounced in the high safety stock condition (two-tailed $p = 0.0605$). Even in the high safety stock condition, however, the improvement is gradual: The average order for the final 10 rounds is weakly significantly different from the optimum order of 75 for the high safety stock condition, while strongly so for the low safety stock condition (Wilcoxon, $n =$ number of subjects, two-tailed $p = 0.089$ and 0.002 , respectively). Finally, the standard deviations of orders are quite high for both conditions (Figure 2), indicating a good deal of variation between newsvendors even with extended experience. So while extended experience helps, there is a good deal of room for improvement. (Schweitzer and Cachon (2000) also observe somewhat better performance in their high safety stock condition.⁸)

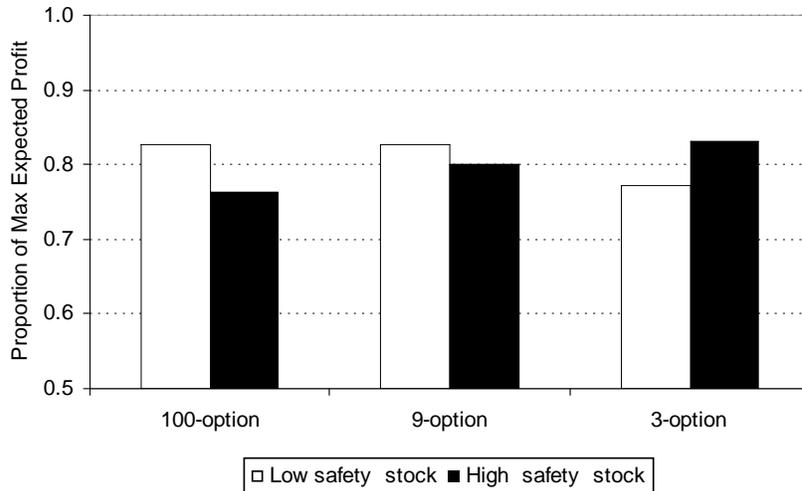
3.4 Results: Flat Maximum Hypothesis and the Reduced Options Treatments

We next move to the tests of the flat maximum hypothesis, having to do with the treatments with the reduced number of options (9 and 3). Figure 3 shows the proportion of maximum expected profit achieved (defined in section 2.3). These range from 0.75 to 0.85. Figure 3 also displays test results for the flat maximum hypothesis, that a thinner set of options improves performance. Looking first at the aggregate results (over all rounds): For the low safety stock condition, the order of performance is the opposite of that predicted. The order is correct for the high safety stock condition but not statistically so. So there is no significant evidence in the aggregate data for the flat maximum hypothesis.

We might expect a flat maximum effect to become more evident with experience. Figure 4 shows how expected profit evolves over time for each of the three treatments. Estimating a simple trend line for the data in Figure 4, there is a significantly positive experience effect on performance in all cases except 9-option high safety stock (OLS, 9-options high safety stock treatment, two-tailed $p = 0.920$; $p < 0.001$ all others). But experience tends to reduce differences across treatments: By the last 10 rounds, most differences in the proportion of the maximum expected profit achieved go in the opposite direction of that implied by the flat maximum hypothesis (see the bottom portion of the right panel of Figure 3). Hence, decreasing

⁸ This tendency to do better in high safety stock conditions is observed in all our experiments (see section 6).

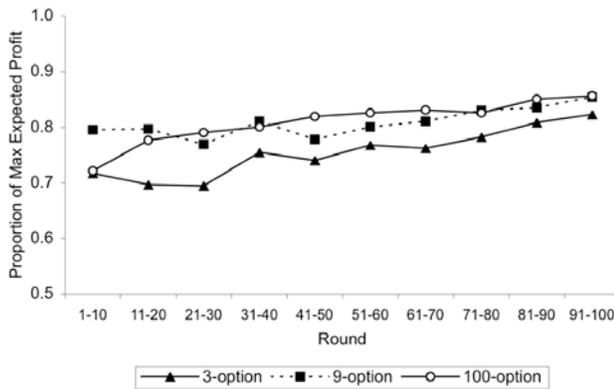
the number of ordering options has no systematic, positive effect on performance, even with experience.



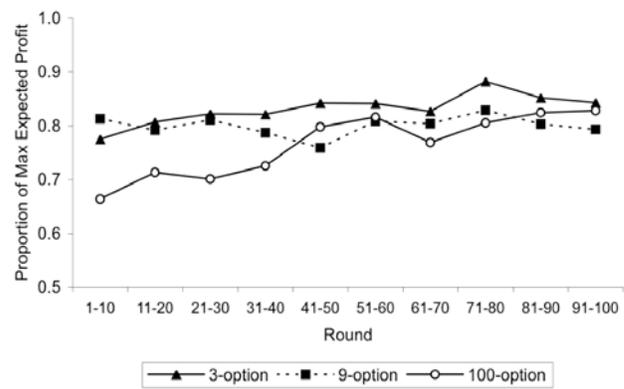
Mann Whitney tests of the flat maximum hypothesis (one-tailed, sample size = number of subjects)

Comparison	p-value
Low, all rounds	
H ₀ : 3-option ≤ 9-option	0.997
H ₀ : 3-option ≤ 100-opt	0.918
H ₀ : 9-option ≤ 100-opt	0.663
High, all rounds	
H ₀ : 3-option ≤ 9-option	0.155
H ₀ : 3-option ≤ 100-opt	0.102
H ₀ : 9-option ≤ 100-opt	0.179
Low, last 10 rounds	
H ₀ : 3-option ≤ 9-option	0.883
H ₀ : 3-option ≤ 100-opt	0.953
H ₀ : 9-option ≤ 100-opt	0.814
High, last 10 rounds	
H ₀ : 3-option ≤ 9-option	0.117
H ₀ : 3-option ≤ 100-opt	0.454
H ₀ : 9-option ≤ 100-opt	0.873

Figure 3: Proportion of maximum expected profit achieved in 3- 9- and 100-option treatments. The figure displays averages for all 100 rounds.



Low safety stock



High safety stock

Figure 4: Proportion of maximum expected profit achieved as experience accumulates for 100-9- and 3-options treatments of study 1.

It is in many respects remarkable that decreasing the number of options from 100 to 3 yields so little improvement in the expected profitability of decisions. The thinning of the option space eases the cognitive processing requirements. Particularly notable is the 3-option, low

safety stock treatment in which the option with the highest expected profit is also the option with the least variance in profit, making it the best choice for risk averse newsvendors (Eeckhoudt et al. 1995). So again we see that risk aversion cannot explain the pattern of deviation from optimal ordering.

We now move to examining individual patterns of behavior. Table 2 shows the categorization for the 3-option treatments. While somewhat ad hoc in nature, the categorization is tied to known behavioral biases. We first separate out newsvendors whose orders are statistically correlated with the previous demand draw. This behavior is consistent with the gambler’s fallacy (Kahneman and Tversky 1972), based on a fallacious belief that independent draws are positively correlated (*Positive*; ex., ‘hot hand’ fallacy in basketball) or negatively correlated (*Negative*; ex., believing a number on the roulette wheel is ‘due’). We then separate out any newsvendors whose choices are not statistically different from random. Remaining newsvendors are classified by modal behavior: *Optimum* if the most prevalent action is the optimum order; *AvgD* if it is average demand; and *LowestR* if it is the strategy with lowest expected return.

Table 2: Ad hoc classification of newsvendors by ordering behavior.

<i>Category</i>	<i>Pattern of behavior</i>
<i>Positive</i>	Orders are positively correlated with the previous round demand; specifically, correlation coefficient is positive at the 0.05 level of significance.
<i>Negative</i>	Orders are negatively correlated with the previous round demand; specifically, correlation coefficient is negative at the 0.05 level of significance.
<i>Random</i>	Orders cannot be distinguished from a strategy of selecting each option with equal probability; chi-squared test at the 0.05 level of significance.
<i>Optimum</i>	Newsvendors who do not fit into the first three categories and whose modal order was the optimum (75).
<i>AvgD</i>	Newsvendors who do not fit into the first three categories and whose modal order is the average demand (100 in the low, 50 in the high safety stock condition).
<i>LowestR</i>	Newsvendors who do not fit into the first three categories and whose modal order is the one with lowest return (115 in the low, 35 in the high safety stock condition).

Figure 5 shows the proportion of newsvendors in each of the six categories for the last 50 rounds of play (results for the first 50 rounds are similar). About two-thirds either correspond to the gambler’s fallacy (about 40%) or have a modal order of the average demand (25%). About 30% have a modal order that is the optimal order. Choices of about 5% are not distinguishable from random.

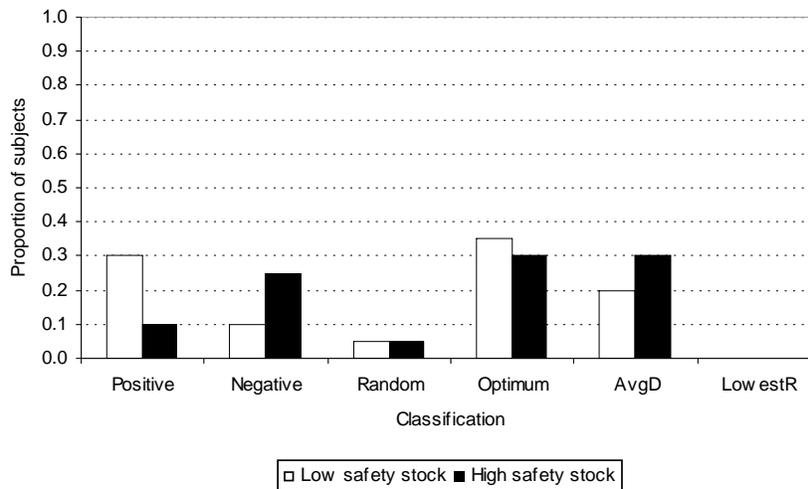


Figure 5: Newsvendors classified for the last 50 rounds, 3-option treatments of study 1.

To summarize the principal findings from study 1: The results from the 100-option treatments are consistent with Schweitzer and Cachon’s. Extended experience improves newsvendor performance, but slowly. There is little support for the flat maximum hypothesis. Aggregate performance is little if at all improved as we decrease the number of ordering options from 100 to 3; and what difference there is tends to diminish with experience. Strikingly, even with only 3 options to choose from, many newsvendors fail to adjust to the optimum, either exhibiting some form of the gambler’s fallacy or staying anchored on ordering average demand.

4. Study 2: Tracking Performance of Foregone Options

4.1 Hypothesis

In this study, we investigate the effect of giving newsvendors feedback about the payoffs associated with options *not* taken (as well as those taken). Learning from forgone options is the basis of fictitious play learning models; these, under certain circumstances, converge to optimal behavior (Brown 1951). A number of empirical studies show that learning from forgone payoffs, combined with adaptive learning, fits observed behavior from a variety of games (Camerer and Ho 1999). Forgone payoff information might enable newsvendors to move off the average demand anchor more quickly, and avoid the gambler’s fallacy behavior we saw in study 1 (Figure 5). We label this treatment manipulation FORE (see Table 1). In a separate treatment, we give newsvendors additional actual and forgone options payoff feedback in the form of 10-

round moving averages. Feedback variability can slow learning (Tversky and Kahneman 1986). The moving averages smooth the variability of single round feedback, potentially improving the effectiveness of the forgone payoff information. We label this treatment manipulation MAVG (see Table 1).

H3: Tracking hypothesis. *Providing payoff information for forgone options leads to newsvendor decisions that achieve a higher proportion of maximum expected profit.*

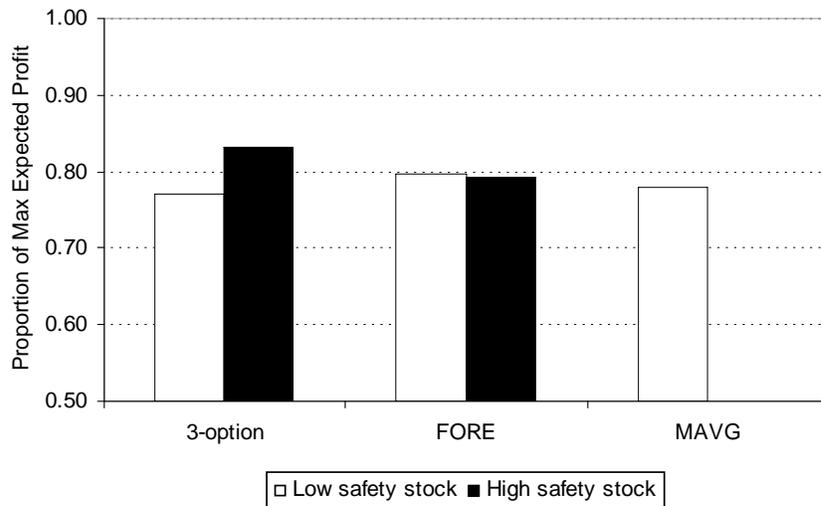
4.2 Method

For study 2, we take the 3-option treatments of study 1 as baseline (see Table 1). The thinned option space of the 3-option game makes it easier to process the information we introduce. For FORE and MAVG treatments, the only changes made to the 3-option games described in section 3.1 concern displays for the additional information. For the FORE treatments, information concerning the payoffs to forgone options for the most recent round was added to the pop-up results box (on-line Appendix A3.1), while for previous rounds this information was displayed in the history box (on-line Appendix A2). The MAVG treatment was conducted in the low safety stock condition. It includes information on the 10-round moving average of the profit for the three ordering options (on-line Appendix A2).⁹ Each FORE treatment had 20 subjects. The MAVG treatment had 18 subjects.

4.3 Results

Figure 6 displays the proportion of maximum expected profit achieved (as defined in section 2.3) in each treatment of study 2. The 3-option data is from study 1. Looking first at the aggregate results (over all rounds), the additional information on forgone payoffs in FORE treatments does not significantly improve profitability; information about moving average in the MAVG treatment has no positive effect.

⁹ Schweitzer and Cachon gave similar information to their subjects. Specifically, they provided tables showing all possible profit outcomes for each inventory order (but not expected values or other statistical analysis). So in principle, subjects had access to the same forgone profit information that we provide. However, to get this information required an action by their subjects, whereas we presented it to all subjects with no action required. Also, the information might arguably be more effective when attention is restricted to 3 options.



Mann Whitney tests of the tracking hypothesis (one-tailed, sample size = number of subjects)

Comparison	p-value
Low, all decisions	
H ₀ : FORE ≤ 3-option	0.342
H ₀ : MAVG ≤ 3-option	0.362
H ₀ : MAVG ≤ FORE	0.663
High, all decisions	
H ₀ : FORE ≤ 3-option	0.832
Low, last 10 decisions	
H ₀ : FORE ≤ 3-option	0.016
H ₀ : MAVG ≤ 3-option	0.617
H ₀ : MAVG ≤ FORE	0.965
High, last 10 decisions	
H ₀ : FORE ≤ 3-option	0.952

Figure 6: Proportion of maximum expected profit achieved in baseline (3-options) treatments, with forgone option feedback (FORE), and with 10-round moving average added (MAVG). The figure displays averages for all 100 rounds.

Figure 7 shows how the proportion of maximum expected profit achieved changes over time.

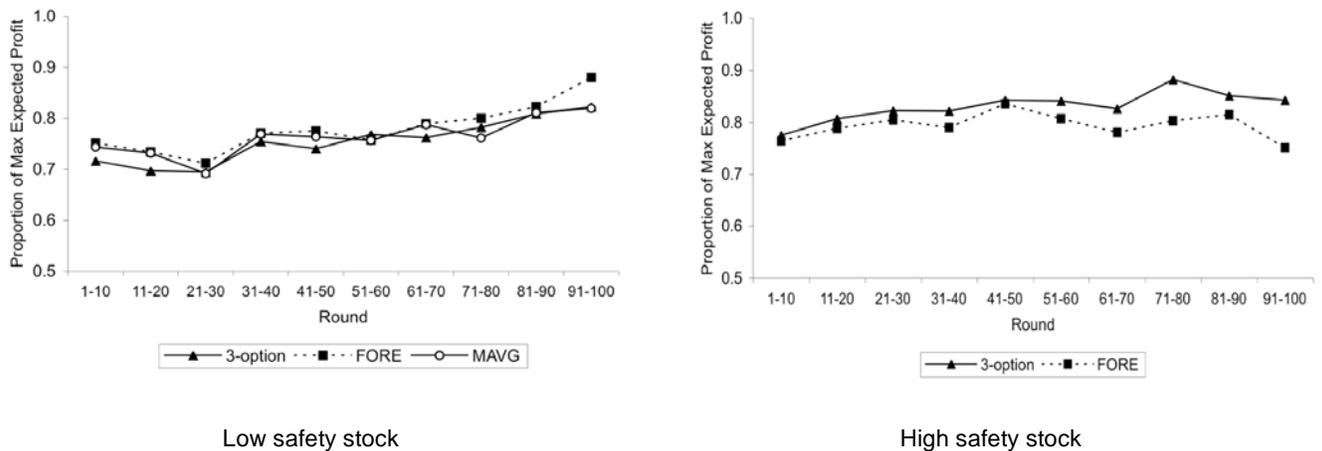


Figure 7: Proportion of maximum expected profit achieved with and without additional feedback as experience accumulates. Baseline (3-options) treatment compared with forgone option feedback (FORE) and with 10-round moving average added (MAVG).

As in study 1, an experience effect is evident in most treatments. However, there is little difference in performance across the treatments. For the last 10 rounds, only the low safety stock comparison between the 3-option and forgone payoffs treatments is consistent with the tracking hypothesis (see test results in Figure 6)—but adding moving average information leads to

virtually no difference with the 3-option data. There is little systematic evidence in the data that forgone payoff information improves newsvendor performance.

Figure 8 shows the break-out of newsvendors using the same classification method introduced in section 3.4. A comparison with Figure 5 finds the introduction of foregone payoff and moving average information has little effect on the pattern of behavior (chi-squared test across the three treatments, high and low safety stock conditions pooled, $p = 0.405$).

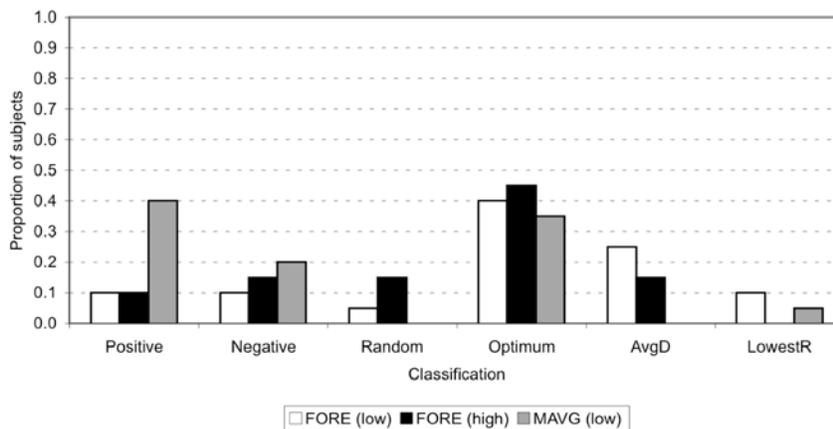


Figure 8: Newsvendors classified for the last 50 rounds, treatments from study 2.

We conclude that tracking information does little to remedy suboptimal behavior.

5. Study 3: Standing Orders

5.1 Hypothesis

In this study, we take a more invasive approach to improving performance: we restrict newsvendors to ordering a standing (fixed) quantity for a sequence of 10 demand periods. We label this manipulation 10P (see Table 1). The impetus for our approach is the ‘law of small numbers’, a tendency to believe that statistically (too) small samples are representative (Tversky and Kahneman 1971). In fact, our data suggests that many newsvendors in our sample jump too quickly to draw conclusions about the optimum order’s expected profitability. For example, for newsvendors not classified as Optimum in Figure 8, the average sample run for the optimum order is 2.4 consecutive orders, with a median and mode of just 1. The uninformative nature of this kind of cursory sampling might explain why some tend to stick to the expected demand anchor, while others persist in the gambler’s fallacy.

H4: Law of small numbers hypothesis. *Restricting newsvendor decisions to longer term standing orders leads to newsvendor decisions that achieve a higher proportion of maximum expected profit.*

We will see that this restriction improves performance. As a point of comparison, we also run a treatment in which newsvendors order for one demand period at a time but receive, prior to ordering, a statistical analysis of order profitability including the expected profitability (on-line Appendix A4). We label this manipulation UPFRONT (see Table 1). This manipulation permits a test of whether it is the restriction on ordering behavior in 10P that is critical to behavior or whether the additional information the subjects gain from the extended sampling is an adequate explanation.

5.2 Method

All of the treatments in study 3 provide the same information about payoffs to foregone options included in the FORE treatments of study 2, and so we use these FORE treatments as the baseline (see Table 1).

The 10P treatment restricts newsvendors to a standing order for 10 demand periods at a time. To get a sense of how much more informative 10 consecutive draws are than one draw, observe that the standard error of the average of 10 independent observations is smaller by a factor of the square root of 10 than the standard deviation of a single observation. We collect data for 100 decisions, so 10P newsvendors participate in 1000 demand periods. (The token-to-dollar exchange rate was adjusted to make payoffs per decision comparable to payoffs per decision in other treatments; see section 2.2.) The first 100 demand draws were the same as in the comparable low and high safety stock conditions. After ordering, 10P newsvendors observe the outcomes for each of the ten individual demand periods covered by the standing order (see the screen shot in on-line Appendix A3.2). Additionally we compute and display to them the average performance over the 10 demand periods of the option chosen as well as the two options not chosen. This summary information about averages is shown for all past decisions, and it is located on the same screen subjects use to make the next decision (see on-line Appendix A2). Information on past individual demand-period outcomes for each 10 demand-period round may be reviewed by using a Details button.

The UPFRONT treatment was conducted in the low safety stock condition. Subjects placed orders for one demand period at a time (total of 100 orders), but were given, at the beginning of the session, a sheet with the expected profit and range of profit associated with each of the three order options (see on-line Appendix A4). This fuller accounting of profitability (as opposed to simply stating the expected profit) explicitly stated the variability of profit associated with each option, and was intended to avoid leading the subjects.¹⁰ The 10P high safety stock treatment had 20 subjects. The 10P low safety stock treatment and the UPFRONT treatment had 18 subjects.

5.3 Results

Figure 9 displays the proportion of maximum expected profit achieved (as defined in section 2.3) in each treatment of study 3. The FORE data is from study 2. Looking first at the aggregate results (over all rounds), the 10-demand-period constrained treatments (10P) significantly improve performance relative to FORE and UPFRONT. UPFRONT performance is not significantly better than performance in FORE.

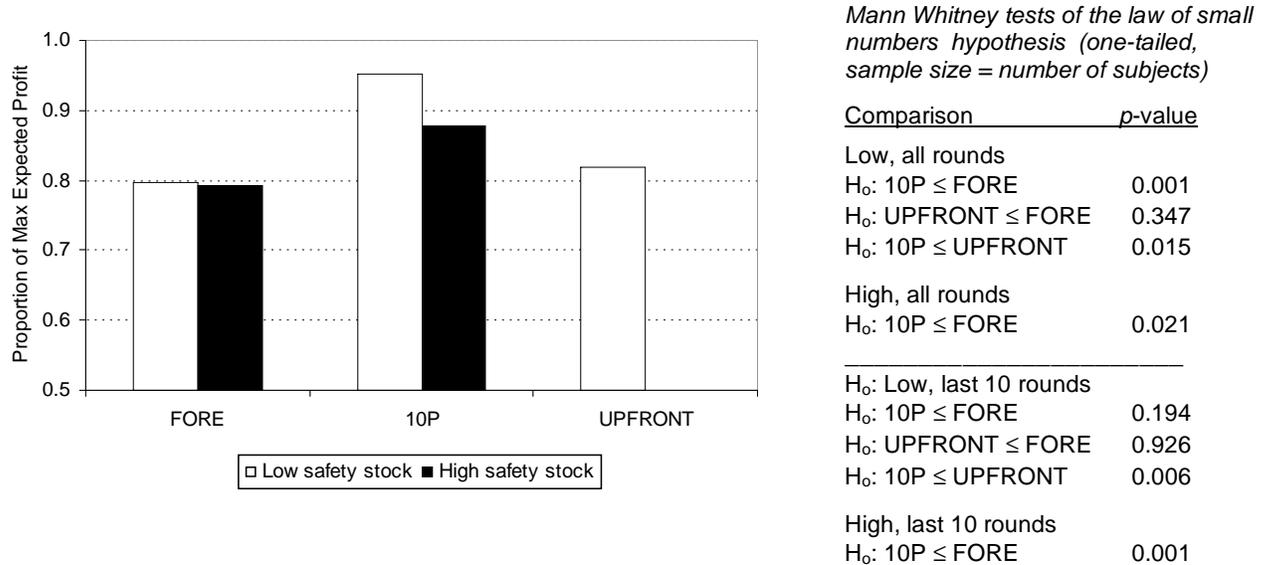


Figure 9: Proportion of maximum expected profit achieved in forgone option feedback (FORE) treatment compared with standing order for 10 demand periods (10P) and with upfront expected profit information (UPFRONT). The figure displays averages for all 100 rounds.

¹⁰ The information provided in this sheet is a condensation of information Schweitzer and Cachon (2000) made available to their subjects. See also footnote 5.

Figure 10 shows how the proportion of maximum expected profit achieved changes over time. An experience effect is evident in most treatments. For the last 10 rounds, 10P significantly outperforms FORE in the high safety stock condition, but not significantly so in the low safety stock condition (test results in Figure 9). However, the performance difference between 10P and FORE low safety stock treatments is stronger than looking at the last 10 rounds implies, because 10P significantly outperforms FORE for every 10 round block shown in Figure 9 save the last (Mann Whitney one-tailed p -values are below 0.01 for all comparisons save $p = 0.0759$ for rounds 81-90 and 0.194 for rounds 91-100). In the last 10 rounds in both high and low safety stock conditions, constrained newsvendors collectively achieve over 90% of expected profit potential.

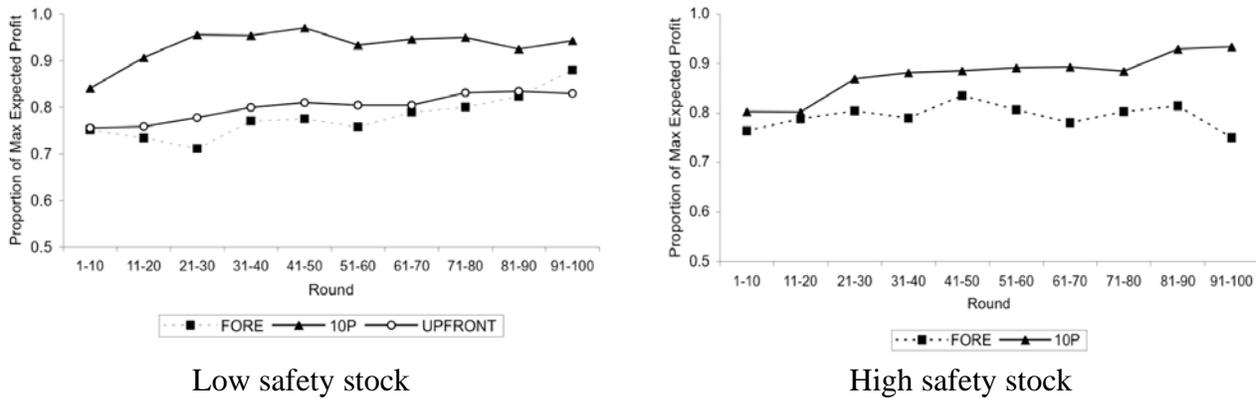


Figure 10: Proportion of maximum expected profit achieved as experience accumulates. Forgone option feedback (FORE) treatment compared with standing order for 10 demand-periods (10P) and with upfront expected profit information (UPFRONT). In all treatments, each round corresponds to a single decision.

One might posit that the better 10P performance is simply due to the fact that newsvendors observe more demand draws in this treatment, rather than to the standing order constraint. However, as has already been noted, giving newsvendors UPFRONT information (on-line Appendix A4)—all the information that in theory they need to make the right decision—does not match the performance in the 10P treatments, either if we compare overall performance or performance in the last 10 rounds (see Figures 9 and 10).

Of course, it is possible that UPFRONT subjects did not comprehend the information provided to them. But consider another way to see that the number of demand draws cannot be the entire story: Figure 11 compares FORE and 10P newsvendor performance, controlling for the number of demand draws observed. Specifically, the figure displays the proportion of maximum

expected profit achieved during each of the first ten rounds in the 10P treatments, and compares this with the entire 100 round sequence in the FORE treatments, grouped in blocks of ten demand periods. Hence the total number of observed demand draws (100) is the same across treatments (note that the x-axis in Figure 11 is in terms of demand-periods). For the second half of the rounds (on the graph, demand periods 51-100), the performance in the 10P treatments is better, on average, than the performance in the FORE treatments for both high and low safety stock conditions, and significantly so for the low safety stock conditions (t-test comparing the proportion of maximum expected profit achieved aggregated by subject over demand periods 51-100, n = number of subjects, two-sided $p = 0.009$ and $p = 0.242$ for, respectively, low and high safety stock conditions). It appears that newsvendors experiment for the first 4 decisions (4 rounds, 40 demand-periods) and then settle on near-optimal behavior, albeit the statistical evidence for this is clear only in the low safety stock condition.

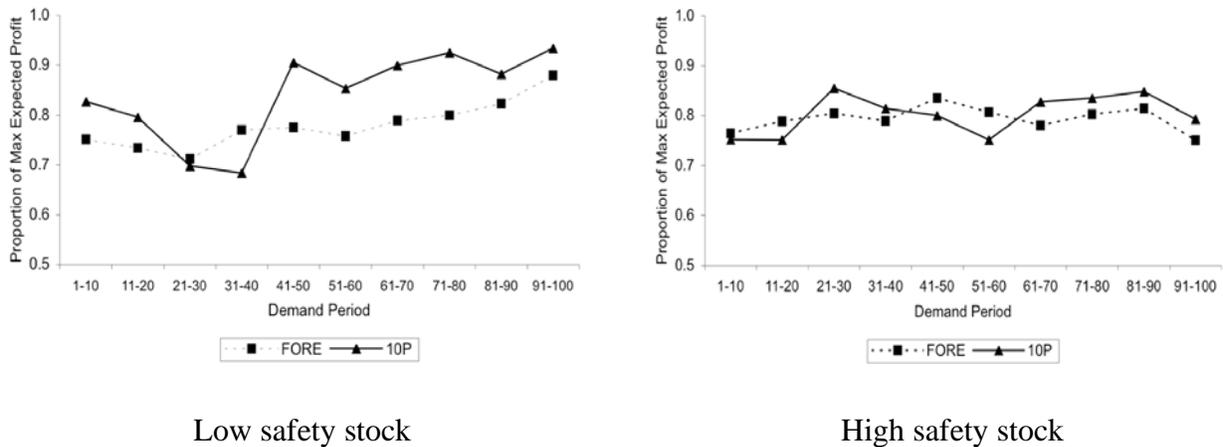


Figure 11: Proportion of expected profit achieved for the first 100 demand draws. Forgone option feedback (FORE) treatment compared with standing order for 10 demand-periods (10P).

Figure 12 compares specifically the choices made in the 10th round in the 10P treatments to the choices made in the 100th round in the FORE treatments. This comparison again controls for the number of demand periods observed. The optimal order is more frequent in the 10P treatments; chi-square tests confirm that the difference is strongly significant in the low safety stock condition (one-tailed $p = 0.019$) and is weakly significant in the high safety stock condition (one-tailed $p = 0.099$). The test on the pooled data is highly significant (one-tailed $p = 0.018$).

Overall, we conclude that, whatever effect the additional demand observations have, the standing order restriction on behavior is a critical factor in the observed better 10P performance.

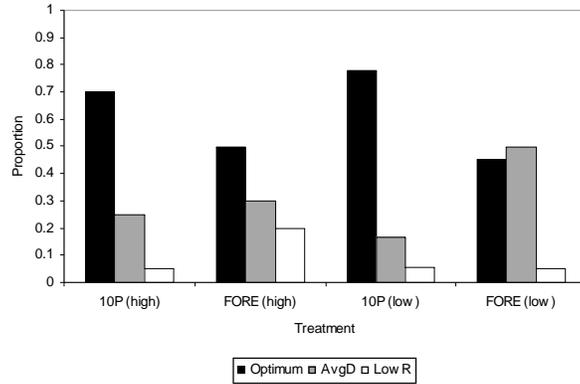


Figure 12: Proportion of newsvendors selecting the optimal order, the average demand order, and the lowest expected profit order in the 10th round of the 10P treatments and in the 100th round of the FORE treatments.

Figure 13 shows that the 10-demand-period restriction has a strong effect on the pattern of individual ordering, effectively wiping out negative correlation (the roulette wheel version of the gambler’s fallacy) as well as anchoring to average demand. The restriction also greatly increases the proportion of newsvendors who place the optimal order.¹¹ A chi-squared test comparing the FORE treatments (data in Figure 8) to the 10P treatments (Figure 13) yields a p -value of 0.002 (high and low data pooled). A chi-squared test comparing the low FORE treatment to UPFRONT is insignificant (p -value = 0.985).

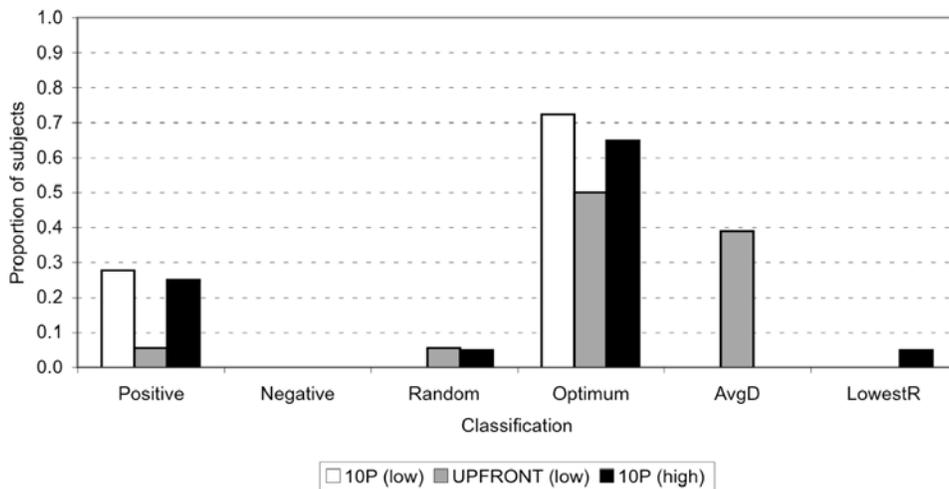


Figure 13: Newsvendors classified for the last 50 rounds, treatments from study 3.

¹¹ In the 10P treatment we classify participants as “Positive” or “Negative” based on the correlation of their decisions with the *average* of the 10 demand draws they observed following the previous decision.

As an aside, note that the superior performance in the 10P treatments relative to the performance in the FORE treatments is evidence that our results are not due in any major way to the fact that the optimal order of 75 is an extreme choice. The order of 75 is an extreme choice in both treatments and so cannot explain the better performance of 10P.

Why do standing orders improve performance? One explanation is to note that restricting behavior to standing orders makes the demand distribution look as if it has a lower variance. This promotes learning in that each data point is now a more reliable indicator of expected profitability of the chosen order quantity. In contrast, the upfront information did not significantly improve performance; it appears that subjects either ignored, disregarded or did not comprehend this information.

6. Summary of the Main Findings

In this section, we present regression estimates culled from the data on the 3-option treatments. The results pull together all of our main findings save for those regarding the flat maximum hypothesis (since the 9- and 100-option treatments are not included), and yield some further insights as well.

Table 3 presents two estimation approaches. In the least squares dummy variable model (LSDV), the dependent variable is the expected profit associated with the round decision (for 10P treatments, the value used is the expected profit averaged over the 10 demand periods affected by the round decision). The 3-option low safety stock treatment of study 1 is taken as the baseline. The independent variables then reflect all the factors we manipulated in the three option studies: a *Round* variable (decision number 1 to 100); dummy variables for *High* (=1 if high safety stock), *Forgone* (=1 if forgone payoffs), *Moving average* (=1 if moving average information), *Standing order* (=1 if the 10-demand-period restriction), and *Upfront info* (=1 if upfront information); and an interaction variable $Round \times Standing\ order$ to check for differences in the trend due to the 10P treatment restriction. See Table 1 for which independent variables apply to which treatments. LSDV estimates are obtained using fixed effects for decision makers.¹²

¹² We also estimated the LSDV model using AR(1) to correct for potential autocorrelation. The resulting estimates are virtually identical to the LSDV estimates presented in Table 3, and none of the conclusions change.

In the logit model, the dependent variable is equal to 1 if the newsvendor placed the optimal order, and 0 otherwise. Thus the logit model allows us to examine how different factors affect the likelihood of placing the optimal order. The independent variables are the same as for the LSDV model, but estimates are from logit regression with random effects for decision makers. (See Greene 1993 for a general discussion of both the LSDV and logit models.)

Table 3: Two models of factor influence on newsvendor performance in the 3-option treatments. The baseline treatment is the 3-option, low safety stock treatment from study 1.*

Dependent variable	Estimates [two-tailed <i>p</i> -value]	
	LSDV	Logit
	<i>Expected profit**</i>	<i>Optimal order=1</i>
Independent variables (factors manipulated)		
<i>Constant</i>	102.99*** [0.000]	-0.362 [0.000]
<i>High safety stock</i>	4.33 [0.000]	0.016 [0.631]
<i>Round</i>	0.13 [0.000]	0.004 [0.000]
<i>Round × Standing order</i>	0.01 [0.516]	0.0006 [0.252]
<i>Foregone</i>	-0.94 [0.182]	0.157 [0.671]
<i>Moving average</i>	-2.22 [0.142]	-0.105 [0.087]
<i>Standing order</i>	15.64 [0.000]	0.163 [0.003]
<i>Upfront info</i>	2.93 [0.002]	-0.017 [0.828]
<i>Random effects</i>	–	0.092 [0.004]
<i>Observations</i>	15,600	15,600
<i>R-sq</i>	0.307	–
<i>Likelihood</i>	–	-7828.2

* See Table 1 for factors manipulated in each study.

** Expected profit per round (decision) averaged over number of demand draws per round, measured in laboratory tokens.

*** Average of the estimated fixed effects.

The main conclusions we draw from Table 3 are:

All other things equal, newsvendors do better in the high safety stock condition than in the low safety stock condition. We see from the LSDV model that, on average, high safety stock condition newsvendors capture more of the expected profit potential (the *High* coefficient in the LSDV model is significant). Interestingly, this is not because high safety stock newsvendors

order the optimal amount significantly more often (the *High* coefficient in the logit model is insignificant). From a tabulation of the data, newsvendors in the low safety stock, 3-option treatment of study 1 are 50% more likely to place the lowest expected profit order (here, the above average demand order quantity) than are newsvendors in the high safety stock, 3-option treatment of study 1. (Forgone payoff information reduces the gap to a still substantial 20%.) Schweitzer and Cachon also found that newsvendors do better in the high safety stock condition. They speculated that this might be because the optimal order in the low safety stock condition (less than average demand) is less intuitive to newsvendors than the optimal order in the high safety stock condition (more than average demand). The popularity of above average demand orders in our low safety stock conditions is consistent with this explanation.

Experience improves performance. Holding all other factors fixed, experience significantly improves profit performance and increases the probability that newsvendors will place the optimal order (see *Round* coefficients). From Table 3, the effect of experience in the 10P treatments (see *Round* \times *Standing order* coefficients) is not significantly different from its effect in other treatments, consistent with our finding that much of the learning in the 10P treatments of study 3 happen within the first few decisions 10P subjects make (see Figures 11 and 12 and accompanying discussions).

Performance tracking information alone does not improve performance. The regressions confirm our finding that there is no significant, positive effect from adding information about *Forgone* option payoffs. Also, the estimates of the *Moving average* coefficients are negative.

Restricting ordering to standing orders for 10 demand periods improves performance both in terms of expected profit potential exploited and in terms of the probability of making the optimal order. We can see in the regressions that this constraint has the biggest impact of any factor studied. In contrast, upfront information has a smaller (but significant) positive effect on expected profitability, and shows no significant effect with regard to the probability of choosing the optimal order. An important caveat, reflected in the nesting of the variables in the regression analyses, is that we tested the standing orders restriction in conjunction with both the 3-option restriction and feedback on forgone payoffs (see Table 1); while neither feature alone was found to have any substantial impact on performance, either or both may be important to the improved performance we saw from adding the standing orders restriction.

7. Conclusions

The results of our study, summarized in the previous section, imply that how experience and feedback are organized for the decision maker may have an important influence on whether inventory is stocked optimally. The insight from behavioral theory that led to the biggest improvement in performance was the law of small numbers bias, the observation that people tend to draw conclusions from inappropriately small samples. The nature of these too-quick conclusions, however, appears to vary widely across individuals, as evidenced by the diversity in search patterns on display in Figures 5, 8 and 13. We note this to highlight the likely importance of developing a theory of optimal inventory institutions that is robust with respect to multiple kinds of misjudgments.

Our study identifies several institutional factors that may promote optimal stocking. First, inhibiting inappropriate responses to short-term information may be critical to keeping people from over-reacting to short-term fluctuations. This task may take on added importance given the advent of technology such as ERP and other tools for supply chain coordination and information sharing that are capable of generating voluminous data. Second, in our study, knowledge gained through personal experience led to the biggest improvement in performance. Gaining experience takes time, which can be expensive. So a high performing inventory control system need deliver appropriate experience, efficiently. One possibility is a carefully designed employee training program. Sterman's (1989) proposed "management flight simulator" is along these lines.

The standing order treatment had two other features, neither of which had significant impact on their own, but may nevertheless have contributed to improved performance. The first feature is the restriction of the number of ordering options. While limiting the number of options by itself did little, it nevertheless permitted us, in the standing order treatments, to provide more focused feedback on the performance of both chosen and unchosen options. In reality, the question of how to limit options in a way that does not eliminate good decisions is a potentially important one that requires further pursuit. The second feature is information about the performance of forgone stocking options. We conjecture that the elimination of this information from the standing order treatments would not substantially affect the learning we see in the standing order treatments. While the small sample bias motivated the standing order treatments,

there may be other explanations for why standing orders mitigate the gambler's fallacy and anchoring behaviors. We have made no attempt, however, to identify or test them here.

Acknowledgements

Katok gratefully acknowledges the support of the National Science Foundation, award number SES 0214337. Bolton gratefully acknowledges the support of the National Science Foundation, award number SES 0351408. We both thank Axel Ockenfels and the Deutsche Forschungsgemeinschaft for financial support through the Leibniz-Program. We thank Gérard Cachon and Maurice Schweitzer for sharing information and materials about their study, and seminar participants at Cornell, the University of Illinois at Urbana-Champaign, Harvard Business School and Penn State for insightful comments.

References

- Benzion, U., Cohen, Y, Peled, R. and Shavit, T. (2005), Decision-making and the newsvendor problem—an experimental study, Working paper.
- Bolton, Gary E. and Anthony M. Kwasnica, eds. (2002), *Special Issue: Experimental Economics in Practice, Interfaces*, 32.
- Boudreau, J.W, Hopp, W., McClain, J.O., and Thomas, L.J. (2003), Commissioned paper on the interface between operations and human resources management. *Manufacturing & Service Operations Management*, Vol 5, No. 3, pp. 179-202.
- Bowman, E.H. (1963), Consistency and optimality in managerial decision making. *Management Science*, Vol. 9, No. 2, pp. 310-321.
- Brown, G. (1951), Iterative solution to games by fictitious play, in T. Koopsmans (ed.) *Activity analysis of production and allocation*, Wiley: New York.
- Cachon, G.P. (2002), Supply chain coordination with contracts. S. Graves and Ton de Kok (Eds.). *Handbook in OR & MS, Supply Chain Management*. Elsevier, North-Holland.
- Camerer, C.F. and Teck-Hua Ho. (1999), Experience-weighted Attraction Learning in Normal Form Games, *Econometrica*, Vol. 67 (4) pp. 827-874.
- Eeckhoudt, L., C. Gollier and H. Schlesinger (1995), The risk-averse (and prudent) newsboy, *Management Science*, 41, 786-794.

- Erev, I. and Roth A.E. (1998), Predicting how people play games: Reinforcement learning in experimental games with unique mixed strategy equilibria," *American Economic Review*, 88, 848-881.
- Gilovich, Thomas, Dale Griffin and Daniel Kahneman, Eds. (2002), *Heuristics and Biases: The Psychology of Intuitive Judgment*, Cambridge University Press: Cambridge.
- Greene, William H. (1993), *Econometric Analysis*, 2nd edition, Macmillian: New York.
- Harrison, G.W. (1989), Theory and misbehavior in first-price auctions, *American Economic Review*, 79, 749-762.
- Hogarth, Robin (1987), *Judgement and Choice*, John Wiley and Sons: New York, 2nd edition.
- Holt, C.A. (1992), ISO Probability matching. University of Virginia Working Paper.
- Holt, C.C., Modigliani, F., Muth, J.F. and Simon, H.A. (1960), *Planning Production, Inventories, and Work Force*, Prentice-Hall, Inc., Englewood Cliffs, NJ.
- Juran, D.C. and Schruben, L.W. (2004), Using worker personality and demographic information to improve system performance prediction, *Journal of Operations Management*, 22, pp. 355-367.
- Kagel, John H. and Alvin E. Roth, eds. (1995), *Handbook of Experimental Economics*, Princeton: Princeton University Press.
- Kahneman, D. and Tversky, A. (1972), Subjective probability: A judgment of representativeness, *Cognitive Psychology*, 3, 430-454.
- Katok, E., Lethrop, A., Tarantino, W. and Xu, S.H. (2001), Using dynamic programming-based DSS to manage inventory at Jeppesen. *Interfaces*, Vol. 31, No. 6, 54-68.
- Lurie, N.H and Swaminathan, J.M (2005) Is timely information always better? The effect of feedback frequency on performance and knowledge acquisition, Working paper.
- Lee, H., P. Padmanabhan and S. Whang, (1997) Information distortion in a supply chain: the bullwhip effect. *Management Science*, 43, 546-558.
- Porteus, E.L. (1990). Stochastic inventory theory. D.P. Heyman and M.J. Sobel (Eds.). *Handbook in OR & MS, Vol 2*. Elsevier, North-Holland, 605-652.
- Prasnikar, Vesna and Roth, Alvin E. (1992). Considerations of fairness and strategy: Experimental data from sequential games. *Quarterly Journal of Economics*, Vol 107, No. 3, pp. 865-888.

- Rapoport, A. (1967), Variables affecting decisions in a multistage inventory task. *Behavioral Sciences*, Vol. 12, pp. 194-204.
- Rapoport, A. (1966), A study of human control in a stochastic multistage decision task. *Behavioral Sciences*, Vol. 11, pp. 18-32.
- Roth, Alvin E. and Erev, Ido. (1995), Learning in extensive-Form Games: Experimental data and simple dynamic models in the intermediate term. *Games and Economic Behavior*, Vol 8, No. 1, pp. 164-212.
- Schultz, K.L., Juran, D.C., Boudreau, J.W., McClain, J.O., and Thomas, L.J. (1998), Modeling and worker motivation in JIT production systems. *Management Science*, Vol 44, No. 12, pp. 1595-1607.
- Schweitzer, M.E. and Cachon, G.P (2000), Decision bias in the newsvendor problem with known demand distribution: experimental evidence. *Management Science*, Vol. 46, pp. 404-420.
- Shanks, D.R., Tunney, R.J, and McCarthy, J.D. (2002), A re-examination of probability matching and rational choice. *Journal of Behavioral Decision Making*, Vol 15, pp. 233-250.
- Siegel, Sidney (1964), *Choice, Strategy, and Utility*, McGraw-Hill: New York.
- Siegel, Sidney and Fouraker, Lawrence E. (1960), *Bargaining and Group Decision Making*, McGraw-Hill: New York.
- Sterman, J., (1989), Modeling managerial behavior: misperceptions of feedback in a dynamic decision making experiment. *Management Science*, 35, 321-339.
- Sterman, J., (2000) *Business dynamics: systems thinking and modeling for a complex world*. Irwin McGraw-Hill.
- Tversky, A. and Kahneman, D. (1986), Rational choice and the framing of decisions, *Journal of Business*, 59, 251-284.
- Tversky, A. and Kahneman, D. (1971), The belief in the law of small numbers, *Psychological Bulletin*, 76, 105-110.