Journal of the Operational Research Society (2008) 59, 1281–1287

www.palgrave-journals.com/jors

Decision-making and the newsvendor problem: an experimental study

U Benzion[1,2], Y Cohen[3], R Peled[1] and T Shavit[3]*

[1]Ben-Gurion University of the Negev, Beer-Sheva, Israel; [2]The Max Stern Academic College of Emek Yezreel, Israel; and [3]The Open University of Israel, Raanana, Israel

This paper investigates repetitive purchase decisions of perishable items in the face of uncertain demand (the newsvendor problem). The experimental design includes: high, or low profit levels; and uniform, or normal demand distributions. The results show that in all cases both learning and convergence occur and are effected by: (1) the mean demand; (2) the order-size of the maximal expected profit; and (3) the demand level of the immediately preceding round. In all cases of the experimental design, the purchase order converges to a value between the mean demand and the quantity for maximizing the expected profit.

Journal of the Operational Research Society (2008) **59,** 1281–1287. doi:10.1057/palgrave.jors.2602470
Published online 11 July 2007

Keywords: inventory; learning; behaviour; management; optimization

1. Introduction

In the newsvendor problem, the decision-maker, facing uncertain demand distribution, has to decide how many units to buy each day. Since Whitin (1955) first presented the newsvendor problem, it has become one of the classic models in inventory management. Interest in the newsvendor problem and its various versions remains unabated and many extensions to it have been proposed in the last decade (Lau and Lau, 1997; Khouja, 1999; Shore, 2004). The newsvendor problem focuses on the purchase of perishable products. The mathematical model maximizes the expected profit by determining the optimal order-size. For the sake of convenience, the order-size of the maximal expected profit is abbreviated as 'Optimal order'. Optimal order and expected profit are functions of: (1) the item cost and the marginal profit, and (2) the demand distribution (Nahmias, 1994).

In this paper, we present an experiment in which participants play the role of newspaper storeowners and decide on how many papers to order, given known demand distribution. We use the results in order to answer the following questions:

1. Do decision-makers act according to the theoretical prediction?
2. Do the orders of the decision-makers converge throughout the experiment?
3. What is the effect of alternative parameters on the participants' orders? (Different demand distributions and costs, marginal profit levels etc.)

We used computerized learning experiments and each individual was assigned a single combination of different conditions (uniform or normal demand distribution, and low or high marginal profit). The participants were asked to decide on their order quantity in the course of 100 periods. While other experimental studies focused on uniform distribution demand, we also used the normal distribution demand.

The paper is organized as follows. First, we present a short review of the literature. Second, we define the hypotheses of our study and in the third section we describe the experimental procedure. Next, we present the primary results and provide some possible explanations. Finally, we summarize the conclusions.

2. Literature review

The classical newsvendor problem (Whitin, 1955) deals with a single-period inventory. Unless it is sold, it will lose part or all of its value. The newsvendor (the decision-maker), facing uncertain demand D from a known distribution function $F(D)$ with a probability density function $f(D)$, has to decide on the order quantity Q. The newsvendor problem is extremely popular and it has been extensively reviewed by Gallego and Moon (1993), Silver *et al* (1998), Khouja (1999), Petruzzi and Dada (1999) to mention a few.

Since the cost of each unit is C and the selling price for the customer is P, the marginal profit C_u equals

Correspondence: T Shavit, Department of Management and Economics, The Open University of Israel, 108 Rabutzki, PO Box 808, Raanana 43104, Israel.
E-mail: shavittal@gmail.com

$P - C$. The marginal loss C_0 equals C (or if a salvage value s is returned $C_0 = C - s$). The newsvendor model finds the optimal order quantity (Q^*) by maximizing the expected profit $\pi(Q)$.

To compute the expected profit of a given order Q, the profit is divided into two cases:

(a) for demand exceeding the order quantity—$Q < D$: $\pi(Q) = (P - C)Q = C_u Q$
(b) for demand lower than order quantity—$Q > D$: $\pi(Q) = (P - C)D - C_0(Q - D) = C_u D - C_0(Q - D)$

In the mathematical development below, (b) is divided into $(b_1) = C_u D$ and $(b_2) = C_0(Q - D)$. Computing the expected profit of an order of Q items, based on (a) and (b) yields,

$$E[\pi(Q)] = \overbrace{C_u \int_Q^{\infty} Qf(D)dD}^{(a)} + \overbrace{C_u \int_0^Q Df(D)dD}^{(b_1)} - \overbrace{C_0 \int_0^Q (Q - D)f(D)dD}^{(b_2)} \quad (1)$$

The well-known formula for optimality conditions of (1) is:

$$F(Q^*) = \frac{c_u}{c_u + c_0} \quad (2)$$

Carlson and O'keefe (1969) were the first to report an experiment with the newsvendor problem. In this instance, the newsvendor problem was part of a much larger experiment in scheduling decision-making. The authors reported participants as making erratic decisions so that no conclusion could be made except that 'participants made almost every kind of mistake'. Fisher and Raman (1996) provided evidence from a firm engaged in manufacturing fashion apparel to indicate that order-purchase decisions do not correspond to the optimal order. In other studies, Sterman (1989) and Diehl and Sterman (1995) discussed the anchoring phenomenon and insufficient adjustment bias in an inventory distribution system experiment with multiple actors, time periods, feedback and time delay. However, these studies were not designed to disentangle biases in the newsvendor context.

Schweitzer and Cachon (2000) conducted an important experimental test of the newsvendor problem model. In their study, they analysed 15 decision periods of ordering for each subject with known uniform distribution. They show that participants systematically deviate from the optimal order and that when marginal profit is larger (smaller) than the cost, participants tend to order less (more) than the optimal order. Bolton and Katok (2004) extended their work using 100 decision rounds. They found that enhanced experience improves newsvendor performance, although this improvement is, on average, rather slow.

Both Fisher and Raman (1996) and Schweitzer and Cachon (2000) claim that there are behavioural factors that lead to deviation from the optimal order, such as risk and loss aversion, underestimation of opportunity cost or waste aversion.

3. Hypotheses

This study examines three hypotheses. One is based on theoretical model but the other two refer to behaviour-based learning theories.

First, we assume that the participant, who follows the optimal order calculated by the newsvendor problem model, is also biased towards the demand distribution mean. We base this assumption on the 'central tendency bias' as discussed, for example, by Hollingworth (1910), Helson (1964) and Crawford *et al*, (2000).

Formally, we assume that the participants order (Q) is a weighted average of the optimal order Q^*, and the distribution mean, $E(D)$.

$$Q = \alpha_t^* E(D) + (1 - \alpha_t)^*(Q^*) \qquad \text{for } 0 \leqslant \alpha_t \leqslant 1 \quad (3)$$

The mean coefficient, α_t, is the strength of the 'central tendency bias' for each subject.

3.1. Hypotheses

H$_1$: Participants' order quantity
The participants' order quantity is a weighted average of the optimal order and the demand distribution mean. For initial stages $\alpha_t \approx 1$ (ie the order is close to the mean demand). In a classic learning process, the effect of recent outcomes declines with experience, and the average marginal increase in profit declines with experience. As a result, the decision-maker's order converges to a subjective level.

H$_2$: Learning
Individuals learn during 100 periods, and as a result:

(a) The coefficient of the mean declines over time and so $\alpha_t < 1$ in late periods.
(b) The average profit increases.
(c) The mean and the optimal order weights converge to a subjective level.

We assume that participants are affected by recent outcomes (see Erev and Barron, 2001). Johnson *et al* (2005) found that in the context of trading stocks, consumers strongly prefer to buy winning stocks and sell losing stocks. We use the same effect for the consumer of newspapers. If the difference between previous round demand and previous round order is positive (negative), the subjects increase (decrease) the order as they would with a winning (losing) stock. The effect of feedback on inventory decision-making and the learning process was tested in different tasks (Atkins *et al*, 2002; Diehl and Sterman, 1995).

H_3: Effect of previous round results
The current order is higher/lower than the previous order if the difference between previous demand and previous order is positive/negative. Over time (ie in later stages) the influence of the previous round declines.

4. The experiments

The experiments included 60 management students, sophomores and juniors, who had taken a basic course in statistics. The experiments took place at a computer laboratory (The experiment was programmed using Visual Basic and Excel.), and lasted approximately 1 h. Each subject was free to progress at his or her own pace independently of the other participants in the experiment.

The participants were divided into four groups before the experiment to examine the combinations of two profit levels and the two variance levels (using different distributions). Two of the groups (one for each cost level) were assigned the same cost, selling price, and demand distribution as described by Schweitzer and Cachon (2000): a uniform demand range of 1–300 products. The other two groups (one for each cost level) were assigned a normal demand distribution with the same mean ($\mu = 150$) and a SD ($\sigma = 50$) that ensures that 99.7% of the demand distribution is within the range of 1–300.

For the low profit level, the values of the optimal order for the uniform and normal distributions respectively are 75 and 116. For the high profit levels, the values of the optimal order are 225 and 184 for the uniform and normal distributions, respectively.

We tested the normal distribution since the demand in real-life situations may have normal distribution. We also wanted to prevent distribution effect by using two different distributions to test whether the results depended on the distribution. By using the normal distribution we extend the previous analysis by comparing the different distributions separately.

Throughout the experiment, the participants made 100 inventory purchase decision rounds, following 10 rounds of practice. In each round, participants were informed of the cost and price of the product. Each round was followed by: a presentation of the actual demand; the total cost of the order; the total revenue; the demand/supply surplus; the forfeited profits due to inventory shortage; and the profit. The data was presented in a table format. (For discussion on the effects of feedback format, see Atkins *et al*, 2002.)

Before the experiment, participants were handed written instructions (see Appendix A), including examples. The demand's distribution was given to the participants as follows:

(1) For uniform distribution, participants were told that each value from 1 to 300 has the same likelihood of being chosen.
(2) For normal distribution, participants were given a table with demand results of 100 simulated days. This represented the normal distribution in a palpable manner.

To provide concrete incentives, at the end of the experiment, one of the rounds was randomly selected and the participants were paid proportionally (The average payment was 20 NIS, or about \$5.) to the profit in the selected round (in cash).

5. Results

Table 1 presents the average weight (a_t) of the mean in the first and last 20 periods according to Equation (3). To validate the effect of learning on the order decisions, we used a paired *t*-test to compare the average weight in the first 20 periods and in the last 20 periods. The average order in the five blocks of 20 periods each is presented in Figures B.1 and B.2.

First, we see that the average coefficient (α_t) declines over time, meaning that the tendency to move towards the distribution's mean declines while the subjects move closer towards the optimal solution of the newsboy problem. This is also shown in Figures B.1 and B.2.

A total of 76.6% of the subjects move toward the direction of the optimum; 1.6% stay at the mean; and the rest, 21.6% of the subjects move away from the optimum. This indicates that participants change their quantity toward the optimal order.

Next, we used *t*-tests to examine the hypothesis that the average coefficient is not different from one, meaning that the order is equal to the mean. The results show that in most of the treatments the average weight is not

Table 1 Average 'mean coefficient' ** in the first and last 20 periods—normal and uniform

Distribution	*Margin profit*	*Average coefficient of mean* (α) *First 20 periods*	*Average coefficient of mean* (α) *Last 20 periods*	*First to last 20 periods t-value (paired t-test), p-value*
Uniform	Low	0.98	0.63^{++}	$t = 3.22,\ p < 0.01$
	High	0.79^{++}	0.47^{++}	$t = -3.96,\ p < 0.01$
Normal*	Low	1.14	0.7^{+}	$t = 2.6,\ p = 0.01$
	High	1.07	0.5^{++}	$t = -4.3,\ p < 0.01$

*$\alpha > 1$ signifies average order on the side of the mean that is opposite to the optimal level.
**$Q = \alpha_t E(D) + (1 - \alpha_t)(Q^*)$ [(Q) is a weighted average of the optimal order Q^*, and the distribution mean, $E(D)$].
$^{+}$Significance of 10% level for the hypothesis that the average weight $\neq 1$.
$^{++}$Significance of 5% level for the hypothesis that the average weight $\neq 1$.

significantly different from one in the first 20 periods, while in the last 20 periods, the average order is significantly different from one.

The results are consistent with hypotheses H_1 and H_{2a}. Participants' order quantity is a weighted average of the optimal order and the demand distribution. In the initial rounds, $\alpha_t = 1$.

Next, we calculated for each subject the absolute change in the order (in percentage) from one period to the next period. The absolute change is used as a measure of convergence. Next, we calculated the average change for each block of 20 periods.

Sixty-seven per cent of the subjects show a decline in the average change between the first 20-period block and the last 20-period block.

In Table 2, we show the percentage of subjects that show an average absolute change in each range in the last 20-period block.

Table 2 Percentage of subjects in different ranges of absolute change in the first and last 20 periods

Average absolute change	*First 20 periods (%)*	*Last 20 periods (%)*
Exactly 0%	3.2	16.4
0% < change < 5%	16.3	37.7
5% < change < 10%	18	21.3
10% < change	62.5	24.6

Table 2 shows that only 24.6% of the subjects show an average change higher than 10% in the last 20 periods compared to 62.5% in the first 20 periods. This shows significant convergence into a stable order over the experiment. Note, however, that this stable order is not the optimal order from the mathematical model.

A total of 54.1% of the subjects show an average change lower than 5% in the last 20 periods, compared to 19.5% in the first 20 periods.

Table 3 shows the average profit in the first and last blocks of 20 rounds.

For each case we present the optimal order's average profit (in brackets). This profit was calculated by using the optimal order in each period instead of the subject's order.

Table 3 shows that the average profit in the last 20 rounds is higher than the average profit in the first 20 rounds, meaning that the profit is increasing between the first rounds and the last rounds, consistent with Hypothesis H_{2b}.

Table 4 presents the average rate between the actual profit and the profit calculated by using the optimal order as follows:

$$100 \times \frac{\text{actual profit} - \text{profit using mean}}{\text{optimal profit} - \text{profit using mean}} \quad (4)$$

Table 4 shows that the rate between the actual profit and the profit calculated by using the optimal order is improving between the first rounds and the last rounds. The results are

Table 3 Average profit in the first and last 20 periods—normal and uniform

		Average profit in the first and last 20 periods		
Distribution	*Margin profit*	*First 20 periods*	*Last 20 periods*	*First to last 20 periods t-value (paired t-test), p-value*
Uniform	Low	−225 (1)	−124 (32)	$t = 1.9,\ p = 0.04$
	High	621 (640)	705 (754)	$t = 9.16,\ p < 0.01$
Normal	Low	−24 (147)	99 (183)	$t = 3.96,\ p < 0.01$
	High	919 (931)	994 (1006)	$t = 5.73,\ p < 0.01$

In brackets, we present the optimal order average profit. This profit was calculated by using the optimal order in each period instead of the subject's order.

Table 4 Average rate*, between the actual profit and the profit calculated by using the optimal order in the first and last 20 periods—normal and uniform

		Average profit in the first and last 20 periods		
Distribution	*Margin profit*	*First 20 periods (%)*	*Last 20 periods (%)*	*First to last 20 periods t-value (paired t-test), p-value*
Uniform	Low	−16	40	$t = 5.571,\ p < 0.01$
	High	36	71	$t = 2.48,\ p = 0.01$
Normal	Low	−48	43	$t = 3.68,\ p < 0.01$
	High	13	80	$t = 3.5,\ p < 0.01$

*We calculated the rate as follows: 100 × (Actual profit − Profit using mean)/(Optimal profit − Profit using mean).

Table 5 Number of changes towards and away from last demand

Distribution	Margin profit		Average changes towards and away from previous demand: First 20 rounds	Last 20 rounds	First to last 20 rounds t-value (paired t-test), p-value	Over all
Uniform	Low	Towards prior demand	6.7	3.3	$t = 3.14,\ p < 0.01$	22.9
		Away from prior demand	2.6	0.9	$t = 3.16,\ p < 0.01$	7.3
		Paired t-test	$t = 3.69,\ p < 0.01$	$t = 3.32,\ p < 0.01$		$t = 5.23,\ p < 0.01$
	High	Towards prior demand	10.1	5.2	$t = 4.9,\ p < 0.01$	35.2
		Away from prior demand	2.1	1.4	$t = 2.9,\ p < 0.01$	9.7
		Paired t-test	$t = 7.15,\ p < 0.01$	$t = 3.4,\ p < 0.01$		$t = 6.0,\ p < 0.01$
Normal	Low	Towards prior demand	7.5	5	$t = 2.8,\ p < 0.01$	32.9
		Away from prior demand	2.6	0.8	$t = 3.5,\ p < 0.01$	8.2
		Paired t-test	$t = 5.3,\ p < 0.01$	$t = 4.1,\ p < 0.01$		$t = 6.1,\ p < 0.01$
	High	Towards prior demand	7.3	3.3	$t = 5.6,\ p < 0.01$	29
		Away from prior demand	2.6	1.2	$t = 1.7,\ p = 0.05$	8.5
		Paired t-test	$t = 4.8,\ p < 0.01$	$t = 4.3,\ p < 0.01$		$t = 7.7,\ p < 0.01$

consistent with the finding that subjects move towards the optimum and away from the mean.

The negative rate at the first 20 rounds in the low profit groups is a result of an average loss since participants order above the optimal order, and the cost of an unsold product (9 NIS) is three times more than the profit from a sold product (3 NIS).

In Table 5, we present the number of times each subject changed his or her order quantity from one round to another in the first and last 20 rounds. We distinguish between changes toward previous demand (Change toward last demand is: (1) if the immediate previous demand is larger than the immediate previous order—an order increase from the previous order and (2) if the immediate previous demand is lower than the immediate previous order—an order decrease from previous order.) and away from previous demand.

Table 5 shows that participants change their order towards the demand of the previous periods more frequently than away from this demand in all the treatments. This indicates that participants are affected by the prior round demand. This is consistent with hypothesis (3). The effect of prior rounds becomes weaker in the last 20 periods indicating that the subject learns throughout the experiment that past information is not relevant to current decision-making.

Overall, the number of changes towards and away from prior demand in the first 20 periods is higher than in the last 20 periods, indicating that participants converge to a subjective order level and a subjective mean weight. This is consistent with hypothesis H_{2c}.

6. Conclusions

In general, there is a convergence to a stationary order quantity and stationary mean coefficient throughout the experiments. This convergence is reflected by a declining number of changes throughout the 100 rounds and an increase in the participants' profits. However, we also demonstrated that subjects converge away from the level of stocking that optimizes expected profit.

We found that in the first purchase decision rounds, participants tend to be more biased toward the mean demand than in the last rounds. This bias persists, since, despite the general convergence to a stationary order, we found a significantly positive mean coefficient in the last rounds too. While we cannot explain why subjects do not converge to the expected value of optimal order in the newsvendor problem, the existence of bias towards the mean, which we found can partly explain the way subjects move from the optimal order.

The results are also consistent with the hypothesis that subjects are affected from previous experience. If past demand is higher than past order (demand surplus), participants tend to increase their order. Participants are affected by the demand or supply surplus. If past demand is lower than past orders (supply surplus), participants tend to reduce their orders. However, participants learn throughout the experiment to reduce this effect.

Clearly, one should be very careful in generalizing from simple experiments to behaviour and prices in real inventory problems. The experiment was conducted in a laboratory with students and a virtual product. However, in real-life situations the inventory managers may deal with many different products. Moreover, inventory managers may use their experience and not theoretical results when deciding on the order quantity. We hope, however, that the intriguing results of this study will motivate further research of the interaction between individual behaviour biases and the inventory problems.

Acknowledgements—The authors would like to express their gratitude to the Open University of Israel for supporting this research.

References

Atkins PWB, Wood RE and Rutgers PJ (2002). The effects of feedback format on dynamic decision-making. *Org Behav Human Decis Process* **88**: 587–604.

Bolton G and Katok E (2004). *Learning-by-doing in the newsvendor problem: A laboratory investigation of the role of experience and feedback*. Working Paper, Smeal College of Business, Penn State University, http://lema.smeal.psu.edu/katok/BoltonKatok%2007-13-2004.pdf.

Carlson J and O'Keefe TB (1969). Buffer stocks and reaction coefficients: An experiment with decision making under risk. *Rev Econ Stud* **36**: 467–484.

Crawford LE, Huttenlocher J and Engebretson PH (2000). Category effects on estimates of stimuli: Perception or reconstruction? *Psychol Sci* **11**: 280–284.

Diehl E and Sterman J (1995). Effects of feedback complexity on dynamic decision-making. *Org Behav Human Decis Process* **62**: 198–215.

Erev I and Barron G (2001). *On adaptation, maximization, and reinforcement learning among cognitive strategies*. Working Paper, Columbia University Business School.

Fisher M and Raman A (1996). Reducing the cost of demand uncertainty through accurate response to early sales. *Opns Res* **44**(4): 933–946.

Gallego G and Moon I (1993). The distribution free newsboy problem: Review and extensions. *J Opl Res Soc* **44**(8): 825–834.

Helson H (1964). *Adaptation-Level Theory*. Harper & Row: New York,

Hollingworth HL (1910). The central tendency of judgment. *J Phil Psychol Scient Meth* **7**: 461–469.

Johnson J, Tellis GJ and Macinnis DJ (2005). Losers, winners, and biased trades. *J Consum Res* **32**(2): 324–329.

Khouja M (1999). The single-period (news-vendor) problem: Literature review and suggestions for further research. *Omega* **27**: 537–553.

Lau H-S and Lau AH-L (1997). A semi-analytical solution for a newsboy problem with mid-period replenishment. *J Opl Res Soc* **48**(12): 1245–1259.

Nahmias S (1994). Demand estimation in lost sales inventory systems. *Nav Res Logist* **41**: 739–757.

Petruzzi NC and Dada M (1999). Pricing and the newsvendor problem: A review with extensions. *Opns Res* **47**(2): 183–194.

Schweitzer ME and Cachon GP (2000). Decision bias in the newsvendor problem with a known demand distribution: Experimental evidence. *Mngt Sci* **46**(3): 404–420.

Shore H (2004). A general solution for the newsboy model with random order size and possibly a cutoff transaction size. *J Opl Res Soc* **55**(11): 1218–1228.

Silver EA, Pyke DF and Peterson RP (1998). *Inventory Management and Production Planning and Scheduling*, 3rd edn. John Wiley: New York.

Sterman JD (1989). Modeling managerial behavior: Misperceptions of feedback in a dynamic decision making experiment. *Mngt Sci* **35**: 321–329.

Whitin TM (1955). Inventory control and price theory. *Mngt Sci* **2**: 61–80.

Appendix A. Experiment instructions

- This is a computerized experiment in decision-making. You will function as a retailer of a single product.
- The experiment is composed of a large number of rounds in which you will be asked to make inventory decisions.
- In each round you are able to order the product from your supplier at a wholesale cost. You will then sell the product to consumers at a higher price.
- Consumer demand in each round is randomly selected from known distribution.
- The prices and profits in every round will be in experiment tokens.

Possible scenarios

- *Overage*—If fewer products are demanded than the quantity you ordered, you will have to dispose of some inventory (ie you cannot keep unused inventory for future periods).
- *Shortage*—If more products are demanded than the quantity you ordered, you will have to forgo some sales.

Data after each round

After ordering the quantity from the supplier in each round, the realized demand and the profit will be presented to you.

Theoretical example

The decision screen. The decision screen will not change during the experiment.

Data		
Round:	1	Your order quantity: ________
Price:	15	Confirm
Cost:	5	

You then decide your order quantity. And then press the confirm button. Assume that your order decision is: 380 units and the realized demand was: 136

The results screen.

In product units:		In experiment tokens:	
Your order quantity:	380	The order cost:	1400
The realized demand:	136		
The quantity purchased:	136	The total revenue:	2040
Overage of product:	244		
Shortage of product:	—	The forgone sales value:	—
		Total profit:	640

Payment for experiment. Part of your payment will be fixed (10 NIS) and the other part depends on your profit/loss level.

Following the completion of the experiment, one of the rounds will be randomly picked and will determine the payment for the experiment. This means that the payment is dependent on the quality of your decision. The profit/loss of the picked round will be divided by 50 and added to a fixed sum of 10 NIS.

Assume that the profit in the chosen round was: 704

Your payment for the experiment will be: $10+(704)/50=$ 24.08.

Appendix B. Average Order in 20-round Blocks

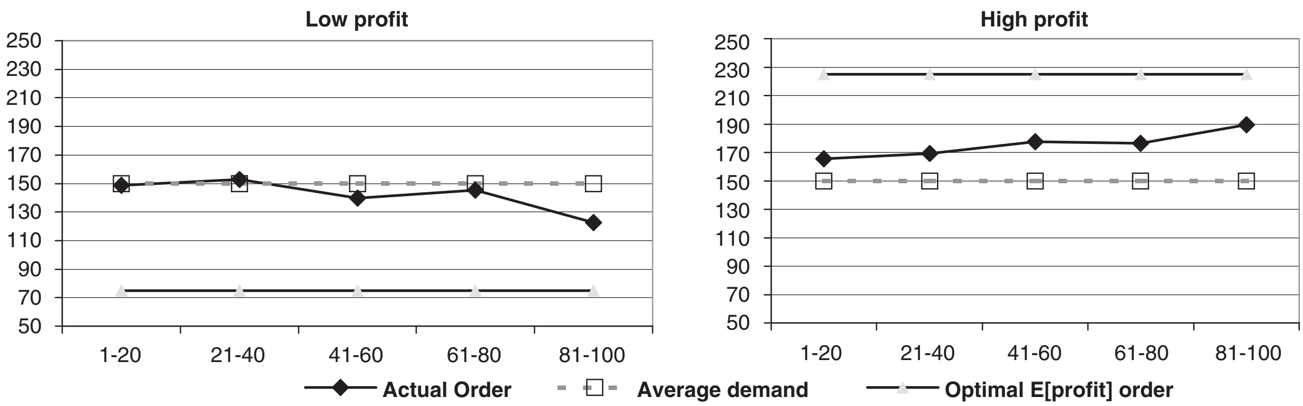

Figure B.1 Average orders along 100 rounds for ***Uniform distribution*** low/high profit

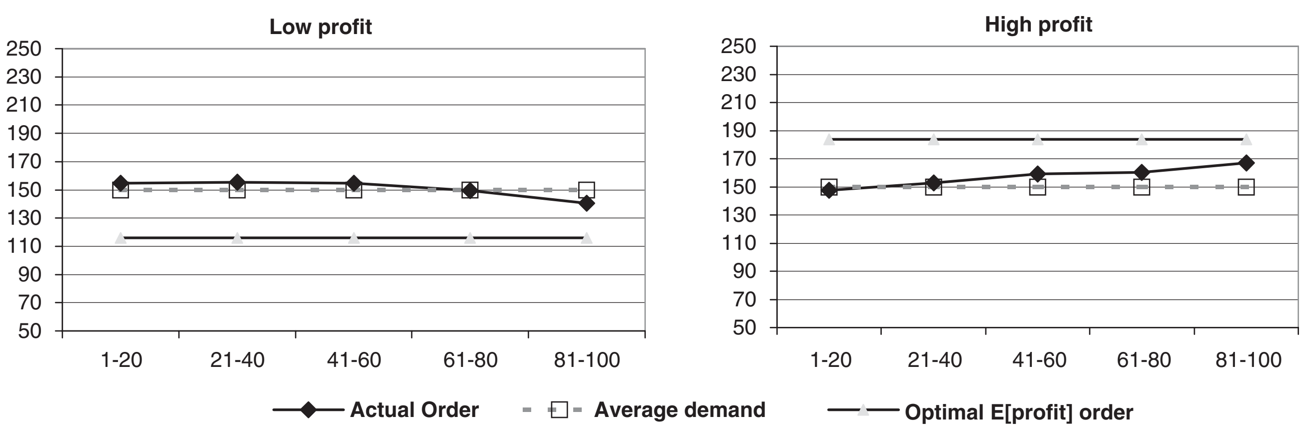

Figure B.2 Average orders along 100 rounds for ***Normal distribution*** low/high

Received July 2006;
accepted April 2007 after one revision