Newsvendor Problem Experiments: Riskiness of the Decisions and Learning by Experience

Arshavskiy V.,[1] Okulov V.,[2] Smirnova A.[3]

ARTICLE INFO

Available Online May 2014

Key words:
Newsvendor Problem;
Competition;
Risk;
Rationality;
Learning by experience.

ABSTRACT

We describe the results of three controlled (multi–rounds) experiments in which the players made decisions on purchasing products by fixed price and further selling under demand uncertainty. In our experiments we have introduced the competition; if aplayer has a lot of product ordered, he could win over buyers away from his competitor. We study the rationality of decisions and the learning by experience. Analysis shows that in each round the average decision of players was more aggressive (risky) than the optimal decision obtained under the assumption that the players maximize the expected profit. But by the end of the experiments the variance of the individual decisions decreased, so the choice of decision was more confident and conscious.

Introduction

The experiments, which we describe in this article, are based on the well known problem of operational management, that is, Newsvendor problem or Newsboy problem (Qin, Wang, Vakharia, Chen & Seref,2011). The classic version of the problem is the following: "How many newspapers should the reseller purchase in the morning, if it is not certain how many newspapers he can sell during the day?" If the seller purchases a few newspapers (almost a risk–free decision), he will sell all newspapers but will get a little profit. If the seller makes a risky decision and purchases a lot of newspapers, he can either make much more profit, or remain with unsold newspapers and suffer losses. The simplest version of the problem suggests that the purchase price of newspapers and the retail price are fixed; each buyer can buy only one newspaper, and yesterday's newspaper worth nothing. The demand for newspapers is random, but theseller knows the probability distribution of the number of buyers that can be evaluated according to his past experience.

Despite its simplicity, the problem is of practical importance. Many companies from the food shops up to large power companies are forced to do stocks of seasonal goods and faced the problem of random demand. There are numerous examples when the company suffered losses due to excess inventory or lost profits due to underestimating demand. In (Fisher & Raman, 1996) it is suggested that some companies systematically loose profits, because their decisions on stock amount are not optimal for random demand.

The Newsvendor problem is often used as a basis for controlled experiments, see (Benzion, Cohen, Peled & Shavit, 2008; Bolton & Katok, 2008; Bostian, Holt & Smith, 2008) and many more. Usually the goal of such experiments is to compare the decisions on stock amount of the participants of the experiment with model theoretical decisionobtained for some ideal rational individual. The results of many of the experiments allow one to conclude that these decisions differ.

The theoretical solution of the classical Newsvendor problem can be obtained, only if we know the probability distribution of the number of possible buyers and target function of the seller. It is assumed that the seller is rational; he makes decisions that maximize some utility function. If this utility function is in a direct proportion to monetary wealth, so the seller is risk–neutral and the best criterion for decision making is to maximize the expected profit (Eeckhaudt & Gollier, 1995).

[1]Research Scientist,Center for Technology and Economics Development, NYU Abu Dhabi and Graduate School of Management, Saint–Petersburg State University, Email: victor.archavski@gmail.com

[2]Associate Professor, Department of Finance and Accounting, Graduate School of Management, Saint–Petersburg State University, Email: okulov@gsom.pu.ru

[3]Master in Management,Graduate School of Management,Saint–Petersburg State University, Email: alenav.smirnova@gmail.com

But there are some "but". Firstly, it is questionable that all people (or most people) are neutral to risk. Secondly, the information about the probability distribution of potential buyers doesn't have any sense for ordinary people who participate in experiments. Ordinary person doesn't know how to use this information, and he/she acts intuitively. Thirdly, aparticipant of a repeated multi–round experiment observes the immediate (current) profit and based on that he adopts the decisions in next rounds. He can estimate the expected profit (it's the abstract value) only post factum by analyzing accumulated profit which was got as the result of previous decisions. So it is hard to assume that the real decisions of people and model theoretical decision for ideal rational individual[4] will be the same because neither presuppositions nor the available information content nor the criteria for decision making agree with each other in both cases (the reality and the model).

We took as a basis the classical Newsvendor problem and organized three different experiments. In every experiment there were some number of periods (20–40 rounds of the game), during which the players made their decisions under the same conditions (the same product, prices and the same probability distribution of buyers). In the first experiment the optimal decision is defined by only the uncertainty of the number of buyers (we call it "a game with nature"). In the second experiment the optimal decision depends on the uncertainty of demand and the decision a competitor made in a current period ("static game"). In the third experiment the optimal decision depends on demand uncertainty, the decision of a competitor in a current period, and the decisions of a participant and his/her competitor during previous periods ("dynamic game").[5] Important point is that the parameters in each experiment remained unchanged, but the optimal decision differed from experiment to experiment, it has become increasingly risky.

Our goal was to compare the participants' decisions with the model theoretical decision. We have analyzed the average decision (averaged over all the participants' decisions in a particular round) and the variance of individual decisions. We were interested whether the average decision and the variance of individual decisions differ at the beginning and at the end of the each experiment, whether it was close to the optimal decision if the same problem was solved repeatedly. By comparing the results of different experiments we wanted to understand whether the participants "feel" that the optimal decision has changed and become more risky.

Literature Review

The Newsvendor problem is well known for a long time. A lot of researches were devoted to this problem, but till now this problem is of the interest for scientists and experts (see reviews (Khouja, 1999; Petruzzi& Dada, 1999; Qin, Wang, Vakharia, Chen &Seref, 2011)). Most often, this problem is considered from the point of view of the operational management (Alfares& Elmorra, 2005; Mostard& Teunter, 2006; Zheng& Liu, 2011). We will focus mainly on the review of papers which describe the results of controlled experiments based on Newsvendor problem.

First we briefly discuss the theoretical solution of the classical Newsvendor problem. There is a seller, who every day orders a product at price c per unit and sells it at price p per unit. Assume, every buyer can buy one unit of a product; the number of buyers $\tilde{X}$ is random, the cumulative distribution function $Q(X)$ is known. If a seller ordered Y units of a product, then his/her (random) profit $\tilde{\pi}$ will be equal to (Eeckhoudt& Gollier, 1995):

$$\begin{aligned} \tilde{\pi}(Y,\tilde{X}) &= p \cdot \tilde{X} - c \cdot Y \quad if\ Y > \tilde{X} \\ \tilde{\pi}(Y,\tilde{X}) &= p \cdot Y - c \cdot Y \quad if\ Y \le \tilde{X} \end{aligned} \quad (1)$$

Equations (1) work under assumption that unclaimed goods cannot be sold at any price, a seller has no fixed costs, and he/she has no direct or indirect costs because of unsatisfied demand.

The key assumption of the model is the rationality of a seller. If one assumes that a seller acts to maximize expected profit, the optimal solution is found based on thecondition: $E[\tilde{\pi}] \to max(Y)$(here $E[\dots]$ is the mathematical expectation). Then the optimal decision Y_{opt} satisfies the following equation:

[4] We call it as "optimal decision".
[5] In the second and third experiments all participants were divided into pairs, and two participants played against each other. More information is provided further in the paper.

$$\frac{Q(Y_{opt})}{1 - Q(Y_{opt})} = \frac{p - c}{c} \qquad (2)$$

From (2) it is clear, that the optimal decision of a rational seller depends on product profitability and on probability distribution of number of buyers. The optimal decision Y_{opt} should be less than the expected number of buyers if product profitability is low, that is, $\frac{p-c}{c} < 1$. If sales profitability is high $\left(\frac{p-c}{c} > 1\right)$the optimal decision should be higher than the expected number of buyers.[6]

The results of controlled experiments, during which participants solved classical Newsvend or problem for different c and p and different types of probability distribution of buyers, are described in several articles. In (Schweitzer & Cachon, 2000) the players made decisions on the purchase of low–profitability goods($p - c \cong 0{,}2c$) and high–profitability goods ($p - c \cong 2c$) during 15 rounds;it was studied the adequacy of average decision of players to model one. These experiments were repeated with small variations by (Benzion, Cohen, Peled & Shavit, 2008; Bolton & Katok, 2008; Bostian, Holt & Smith, 2008; Véricourt, Bearden & Filipowicz, 2011).

The main result of all these experiments is the following: the average decision of the players differs from the model decision (2). It is riskier if one deals with low–profitability product (average decision is higher than the model one) and less risky if one deals with high–profitability product (average decision isless than optimal one)—"a pull–to–center effect". These results cannot be explained by the assumption that the players on average are always risk averse or risk loving. Maybe the players' behavior is irrational. The authors of (Schweitzer& Cachon, 2000) analyze how the optimal decision would change under different target function of the seller, but after some experimentation with different price and demand function they conclude that real decisions do not correspond to optimization of any target function. To explain these results, they consider a number of decision making heuristics, including demand chasing.

In (Bostian, Holt & Smith, 2008; Benzion, Cohen, Peled& Shavit, 2008), the authors examined how the number of buyers in previous round affects the player's decisionin the next round. The results show that if the number of buyers in a previous period was more (less) than the quantity of product ordered, the playermore likely will increase (decrease) his/her order quantity for the next period. The same effect was mentioned in (Atkins, Wood & Rutgers, 2002). In (Ho, Lim & Cui, 2010), the authors constructed the player's behavioral model with two parameters related to the asymmetric reaction to overstocking and understocking. In (Bostian, Holt & Smith, 2008; Gavirneni& Xia, 2009), the authors studied the "anchoring effect" — the possible connection between players' decisions and the available information (the average number of buyers) or some additional information (the decision of other player or expert's advice).

It is worth to mention that almost in all experiments the average decision of players (especially during first rounds) is close to the expected number of buyers. For most of participants it is the only understandable information about future demand (indeed, not everyone knows what the variance or quadratic deviationis). But, one can ask, does a player's decision become closer to the optimal decision if this player solves the problem repeatedly? By the end of 100–round experiments it was observed, that the average decision of the players becomes notably closer to the optimal decision (Benzion, Cohen, Peled & Shavit, 2008). Thatis true for both low– and high–profitability of product. In (Bolton& Katok, 2008) the importance of "learning by doing" is also shown; own experience of players is more important than additional information.

The influence of gender and nationality of players on their decision making is of main interestin (Véricourt, Bearden, & Filipowicz, 2011; Feng, Keller & Zheng, 2011). It is interesting to note that on average the female decisions are more anchored to the initial information about expected number of buyers, and during the experiment the female players are less likely to adjust their initial decision.

In a number of papers the main attention was paid to the differences in average decisions which are made individually or in group. The results of different experiments on this question show that the decision made after the group discussion is closer to optimal decision in comparison with the decisions made individually. In repeated experiments the decisions made in first rounds do not show any noticeable difference between

[6] Strictly speaking, if $p - c = c$ then Y_{opt} is equal to the median of the distribution. However, in most practical cases the distribution median and the mathematical expectation (expected value) are close to each other.

individual and group decisions, but with the course of experimentthe group starts to win individuals in the amount of profit they get (Kocher & Sutter, 2005; Cooper & Kagel, 2005; Charness & Karni, 2007). However in (Cox, 2002) any significant difference between group and individual decisions was not found.

It should be noted that in almost all experiments the participants solved the classical (isolated) Newsvendor problem, where the optimal decision of a player is not related with the decisions of any other players. More sophisticated variants of competitive environment where the decision maker has to take into account the current decision of other player are considered mainly theoretically. Particularly, the single period Newsvendor problem with 2 and 3 sellers was first studied in (Parlar, 1988) and (Wang & Parlar, 1994). The authors of (Netessine & Rudi, 2003) summarized the results of the studies for the problem with arbitrary number of sellers with interdependent demand functions. In (Nagarajan & Rajagopalan, 2009) the theoretical solution of Newsvendor problem was considered for the case of two competitors who sell substitute products. It is shown that starting from some particular ratio between costs (costs of unsatisfied demand and storage of unsold goods were also taken into account) and price of the product for one of sellers, his/her optimal decision on number of product to order does not depend any more on decisions of a competitor. A controlled experiment that documents the behavioral ordering regularities in competitive newsvendor environments is described in (Ovchinnikov, Katok, Moritz & Quiroga, 2013). The authors propose an analytical model that extends the standard theory of newsvendor competition by deriving the optimal (best reply) policy for competing with a behavioral newsvendor. The paper (Moritz, Hill & Donohue, 2013) investigates the relationship between cognitive reflection and newsvendor decision making.

To sum up our brief review, we want to point out, that the question of players' rationality in competitive environment and riskiness of intuitive decisions is still open. By rationality we mean making the decisions, which are close to optimal ones (which maximize the expected profit); by riskiness we mean the readiness to make the decisions, which can bring a higher profitand incur big losses.In contrast to the previously described experiments, we wanted to examine not only the average of players' decisions, but also the scatter of their individual decisions. We believe that a small spread of decisions is an important sign that the participants understand their actions and intuitively follow some strategy.Perhaps the participants by trial and error will find this strategy during the experiment.

That's why we decided to organize a series of multi–round experiments, when the participants have deal with the same product and prices, face the same uncertainty of demand, but the optimal decisions in these experiments are different. During the experiments we aimed to check the following hypothesis:

H1: hypothesis about risk–neutrality. The average decision of participants corresponds to model decision of Newsvendor problem, which is got under assumption, that participants are risk–neutral and maximize expected profit.

H2: hypothesis about learning. In repeated rounds participants are looking for an optimal decision for themselves by trial and error; that is why the average decision of participants in last rounds differs from average decision in first rounds.

H3: hypothesis about learning. In repeated rounds participants are looking for an optimal decision for themselves by trial and error; that is why the variance of participants' decisions in last rounds is less than the dispersion of participants' decisions in first rounds.

Experiments

To organize the experiments we used «Z–Tree» platform (Zurich Toolbox for Ready–made Economic Experiments) (Fischbacher, 2007). The experiments took place in computer classes of Higher School of Economics (HSE Campus in St.–Petersburg). About 40 students of economics and finance participated in each experiment. The participation was voluntary; there was no prior selection of participants. All participants knew about money reward based on the results of experiments and the way this reward is calculated. The reward was calculated as a cumulated profit, which a participant "earned" during the experiment. At the same time there was a minimal guaranteed reward for participation (200 rubles). The typical reward was approximately 500 rubles (about $16). All the participants were guaranteed anonymity of their answers and the reward sum.

More details about the organization of the experiment, including instruction for participants, one can find in (Archavsky & Okulov, 2012). The design of the experiment we used is typical and corresponds to the standards of controlled experiments management. We suppose that the experiments as they were designed by us have all conditions to state that the participants acted rationally and on their own. Every player made the decisions based on his/her own feeling of what is “right”, that is, of common sense and life experience. It is important to note that all experiments took place on the same day, and each student participated in all three experiments. Therefore, the participants could learn from their own experience not only during a single experiment; they could carry the experience over the following experiments.

The first experiment lasted 20 rounds during which the participants solved classical Newsvendor problem individually and independently from each other. The purchase price was $c = 80$, the selling price was $p = 100$. The randomness of demand was described in the following way: “*You have 20 potential buyers and there is a probability of ½ that a buyer will come and buy one unit of a product independently on whether this buyer did it in the previous round*.” Such description sets the discrete random number of buyers $\tilde{X}$ with binomial distribution:

$$q(X) \equiv \mathbb{P}\{\tilde{X} = X\} = \frac{N!}{X! \cdot (N-X)!} \cdot \theta^N; \qquad X = 0, 1, \dots, N; \quad 0 < \theta < 1 \qquad (3)$$

where $q(X)$ is the probability that the number of buyers will be exactly equal X. The parameters of distribution are: $N = 20$ and $\theta = 0.5$.

The cumulative distribution function $Q(X)$, mathematical expectation $E[\tilde{X}]$, and variance $D[\tilde{X}]$ for random value being described in such a way are:

$$Q(X) \equiv \mathbb{P}\{\tilde{X} < X\} = \sum_{x=0}^{X} q(x)\,; \quad E[\tilde{X}] = N \cdot \theta = 10; \quad D[\tilde{X}] = N \cdot \theta \cdot (1-\theta) = 5 \qquad (4)$$

In the experiments based on Newsvendor problem the normal or uniform distributions are usually used. To deal with normal distribution of the number of buyers the participants should be given the mathematical expectation and standard deviation, but the participants may have no idea what these figures mean. It would be even more difficult for the participants to imagine the possible number of buyers based on random sample given to them, as it was done in (Benzion, Cohen, Peled & Shavit, 2008). The uniform distribution of the number of buyers is easy to present, but the usage of uniform distribution in the experiment often leads to large scattering of the number of buyers; it can contradict with life experience of the participants, according to which random values tend to grouping around some average value. The huge dispersion of number of buyers during repeated experiment can disorient a player and his/her decisions can become chaotic. The description of random value proposed by us is easy to understand for the participants, because it corresponds to their life experience (in everyday life a participant can meet a group of 20 people where every group member can do something with a probability of ½).

Having distribution function of the number of buyers (3, 4) and using formula (1) one can easily find the expected profit (mathematical expectation of $\tilde{\pi}$) for different player’s decisions (Fig.1). The black dot corresponds to the maximum expected profit. The optimal decision Y_{opt} which maximizes the expected profit equals to 8 units of a product.

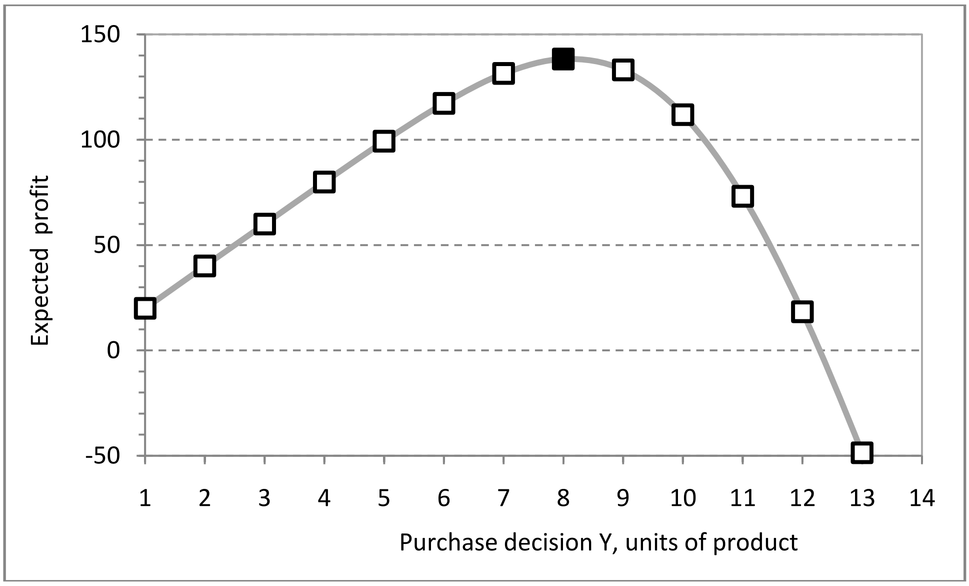

Fig.1. Expected profit for different decisions on product being purchased by a player. The purchase price is $c = 80$, the selling price is $p = 100$; the distribution of number of buyers follows binomial law with parameters $N = 20, \theta = 0.5$. The optimal decision is $Y_{opt} = 8$.

The second experiment lasted 20 rounds; the participants were divided into pairs (one player acted against another); the instruction was the same, but with one additional condition: "*If your buyer was not able to buy a product from you, because you've ordered too less, he will go to your competitor to buy a product there. However, this buyer will stay loyal to you and next time he will firstly go to you. The buyer which is loyal to your competitor can also come and buy a product from you, if your competitor is not able to sell him a product.*"

In this case the profit of the seller 1, who made a decision Y_1, depends on a decision Y_2 of a competing seller 2:

$$\tilde{\pi}_1(Y_1, Y_2, \tilde{X}_1, \tilde{X}_2) = p \cdot min(\tilde{X}_1 + max(\tilde{X}_2 - Y_2, 0), Y_1) - c \cdot Y_1 \qquad (5)$$

where $\tilde{X}_1$, $\tilde{X}_2$ are random numbers of buyers, who came to buy a product fromseller1 and seller 2 accordingly.

On Fig.2 one can see the best decisions for player 1 if he/she would know exactly the decision his/her competitor will make. Obviously, in reality the player 1 doesn't know the competitor's decision. That's why the best decision the player 1 can make is to think that his/her competitor is an ideal rational individual, who maximizes the expected profit too. The player 2 should think the same about the player 1. The equilibrium in this situation is reached, when both competing players make the identical decisions: $Y_{1;opt} = Y_{2;opt} = 9$. (Note that it follows from the symmetry of the problem). It is easy to prove that this set of decisions maximizes total expected profit of a pair of players: $E[\tilde{\pi}_1 + \tilde{\pi}_2] \rightarrow max$Comparing with the first experiment, the optimal decision in the second experiment has changed. We think it is interesting to find out, whether the participants understand through insight that they should make more risky decisions, that is, they should order more units of a product during the second experiment.

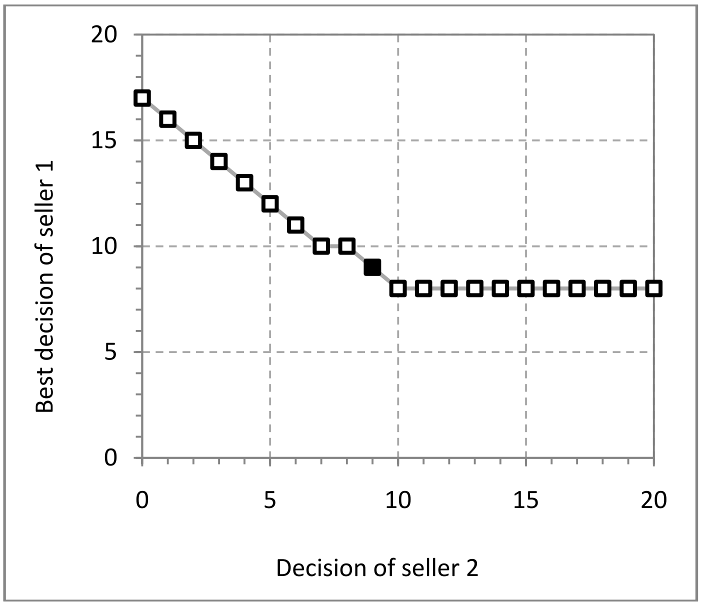

Fig.2. Best decision of a player depends on his/her competitor's decision. The purchase price, the selling price and the distribution of buyers are the same as on Fig.1. If both players are ideal rational individuals, they will make the identical decisions $Y_{opt} = 9$ (black dot).

The third experiment lasted 40 rounds. All the participants were divided into pairs (one player acted against another); the instruction was the same as in second experiment, but with additional condition: "*If your buyer was not able to buy a product from you, because you've ordered too less, he will go to your competitor to buy a product there. If this buyer can buy from your competitor, next time they will firstly go to your competitor. If your unsatisfied buyer cannot buy from your competitor, he will stay loyal to you. The same is applied to the buyers of your competitor: his potential buyers can become loyal to you if they cannot buy the product from your competitor and then they will come to you and buy.*"

So the decision making in current round can directly influence the results of the next rounds. Indeed, if a player does not make risky decisions, he/she can lose own loyal buyers and the profits in next rounds will be less. In limit a player can lose almost all buyers. Obviously in the first 10–20 rounds it makes sense to risk to capture the buyers of competitor.[7] But if a player makes risky decisions, he/she can face big losses in the current round. By the end of the experiment there is no any more reason to make risky decisions.

Finding the optimal decision in such dynamic experiment is a difficult task which requires additional assumptions about the utility for a player of added future profit comparing with the utility of current profit. We assumed that the utility of profit in any round (either current or future) is the same. Let us give a rough estimate of the optimal decision in the first round. If the players make the decisions $Y_{1;t=1} = Y_{2;t=1} = 10$, the loss of expected profit for each player in the first round will be 17.8; if $Y_{1;t=1} = Y_{2;t=1} = 11$, then the loss will be 60.0; if $Y_{1;t=1} = Y_{2;t=1} = 12$, then the loss will be 122.6, and so on. Let the players make the decisions $Y_{1;t=1} = Y_{2;t=1} = 10$, then easy to calculate that a player can capture one buyer of competitor with probability 17.0%, two buyers with probability 6.3%, and so on (Tab.1).

Suppose the player 1 has captured two buyers of the competitor, and then this player simply makes the best decision as he/she did in the second experiment. Since the number of loyal buyers of player 1 has grown to 22, and the number of loyal buyers of the player 2 has fallen to 18, the best pair of decisions in the future will be $Y_{1;t} = 10$; $Y_{2;t} = 8$. Obviously, under these decisions the expected profits of player 1 will increase in

[7]Case study of Urban Outfitters described in (Ovchinnikov, Katok, Moritz & Quiroga, 2013) is a telling illustration of the practical importance of such problem. Urban Outfitters, a major apparel retailer, anticipated that competitors, somewhat depressed by the crisis, would decrease their purchase of the goods, and therefore Urban Outfitters increased its inventory levels. In the end, Urban Outfitters was able to capture a larger share of demand and achieved higher sales and profitability.

next rounds. Approximate estimates of the expected profits in the future and the losses in the current round are presented in Tab.1.

Table 1. Changes of expected profit under different decisions in the first round (compared with the decisions $Y_{1;1} = Y_{2;1} = 9$).

Decisions of players in 1st round	Change of expected profit in 1st round	Probability of capture …				Change of expected profit in next 39 rounds
		1 buyer	2 buyers	3 buyers	4 buyers	
$Y_{1;t=1} = Y_{2;t=1} = 10$	–17.8	0.170	0.063	0.017	0.003	86.7
$Y_{1;t=1} = Y_{2;t=1} = 11$	–60.0	0.148	0.054	0.015	0.003	74.2
$Y_{1;t=1} = Y_{2;t=1} = 12$	–122.6	0.098	0.034	0.009	0.001	46.7

From Tab.1 one can see that in first round the best decision of a player is close to $Y_{1;t=1} = 11$. This decision is risky and leads to a decrease in the expected profit in the first round (because the expected number of buyers is 10), but it gives hope to recoup losses by increasing the expected profits in the next rounds (after all, the number of loyal customers can grow up). The decision $Y_{1;t=1} = 12$ is too risky and future profits cannot fully compensate for the loss of expected profit in the first round. The decision $Y_{1;t=1} = 10$ is insufficiently risky, because the competitor's decision $Y_{2;t=1} = 11$ may result in a decrease of loyal customers the first player has.

In last rounds, if two sellers end with 20 loyal buyers, the best decision is $Y_{1;t=40} = 9$. The best decisions' trajectory depends on the implementation of a random number of buyers for both sellers. If the number of loyal buyers for seller 1 has grown, he must continue to order a lot of the goods (more than 9). If the number of loyal buyers has decreased, the seller must reduce the order of the goods in the following rounds. Therefore, we can speak only about the best average decision. By simulation method we calculated the trajectory of average best decision during the 40 rounds of experiment (Fig.6).

It is interesting to find out do the participants understand that the optimal decision has changed? Are they ready to take on additional risk for the sake of future profits? How will the participants change their decisions during experiment?

Results and Discussion

During the first experiment all participants worked individually and made 780 decisions. Scattering of their decisions was pretty high: from $Y = 3$ to $Y = 20$. But the most frequent decisions were $Y = 8$ and $Y = 10$ (see Fig.3).

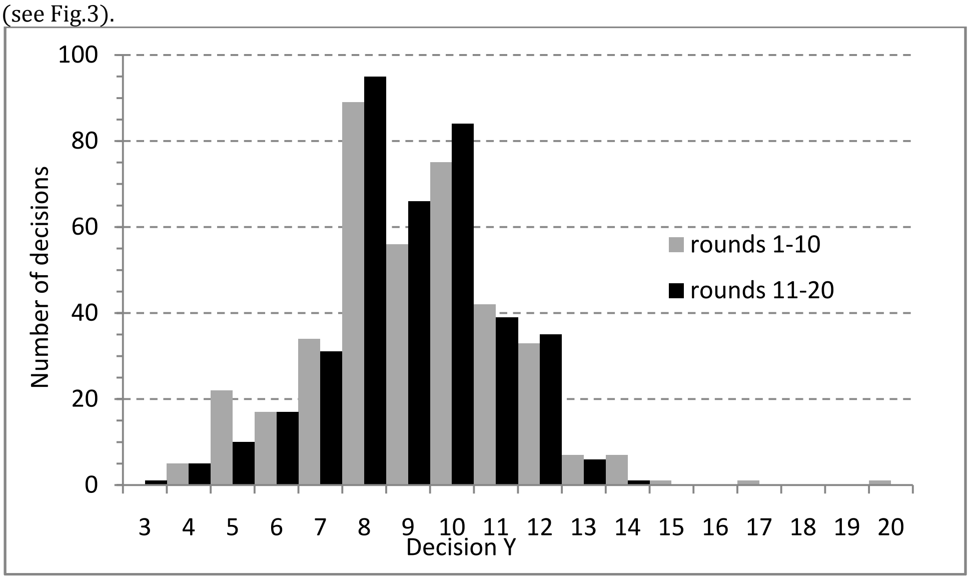

Fig.3. The histogram of the decisions made by participants in first experiment during first 10 rounds and last 10 rounds. The optimal decision is the purchase of 8 units of the product.

We should mention that there were few risky decisions, that is, $Y \geq 13$, which correspond to negative expected profit (Fig.1). There were many decisions $Y = 5;\ 6;\ 7$; they are not risky, but bring not bad profit. This can be seen as the evidence of participants' rationality in decision making. Nevertheless, the players make on average riskier decisions in comparison with the optimal decision $Y = 8$.

Using well–known statistics formulae we calculated for every round and for the whole experiment the mean observation Y_{aver}, sampling variance, and standard deviation for decisions made. Although the decision $Y = 8$ (the optimal decision under assumption about participants' risk–neutrality) was the most frequent decision in first experiment, the average decision $Y_{aver} = 9.05$ turned out to be significantly different from model decision $Y_{opt} = 8$, that is why hypothesis H1 about risk–neutral behavior of participants should be rejected. Statistic hypothesis we checked was formulated in the following way: the mathematical expectation of a participant's decision in the first experiment equals 8.0. Because the value of t–statistic (=14.36) exceeds the critical value on 95% confidence level, we can reject the null hypothesis.

We should admit that decision to buy 9 units of a product leads to only insignificant decrease in profit (Fig.1) and, probably, participants were not able to feel the difference between the decisions $Y = 8$ and $Y = 9$ during 20 rounds. However, the decision $Y = 7$ also leads to slight reduction in profit, but this decision was not as "popular" among participants as the decisions $Y = 8, 9, 10$; it was equally as popular as the risky decisions $Y = 11$ and $Y = 12$, which lead to significant decrease in profit. All this allows us to reject the hypothesis about participants' risk–neutrality and makes us admit that the decision made by participants do not correspond with model assumption about profit maximization.

In some experiments (Benzion, Cohen, Peled & Shavit, 2008) it was mentioned that the average decision of participants was constantly changing and by the end of the experiment it became closer to optimal one. In our experiments we did not observe it (Fig.4). Testing the hypothesis H2 about the equality of mathematical expectations of decisions in first 10 rounds and in last 10 rounds showed that the hypothesis cannot be rejected (t–statistic is 0.187, critical value is 1.96).

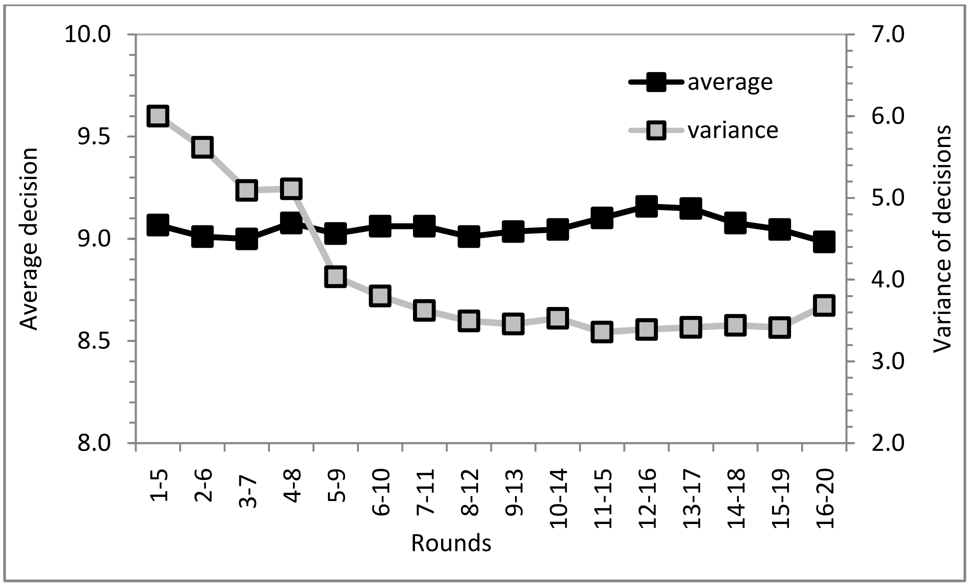

Fig.4. Average decision and variance of decisions (sampling of all participants' decisions during 5 rounds) in first experiment. Optimal decisions is Y=8. Variance of the number of buyers is 5.

While considering the hypothesis about learning, we assumed that in the beginning of experiment the participants can act unconsciously and try to find the best solution by trial and error. Some of players started with very risky decisions, some of them — vice versa — almost risk–free; decisions of players were intuitional and they looked like a mess.[8] But at the end of the experiment, if learning by doing really takes

[8]After the experiment many participants noted that they were trying "to guess" demand in first rounds.

place, participants will make decisions consciously, taking into account the experience they got, and it will be appeared as decreasing of scattering of individual decisions — during last rounds variance of decisions should be smaller.

Fig.4 shows the sampling variance of participants' decisions during rounds 1–5, 2–6 and so on till the last rounds 16–20. It can be seen that in the beginning of experiment (during first 5 rounds) the variance of decisions exceeded the variance of the number of buyers but during last rounds the variance of decisions decreased noticeable and became lower than the variance of the number of buyers. We tested the hypothesis H3 about the equality of variances in first 5 rounds and last 5 rounds. F–statistic is 1.63, critical value is 1.27, so this hypothesis should be rejected.

So we can say that learning by doing took place and it manifested itself not as an approximation to optimal model decision, but as more conscious decisions which were more stable from round to round. Unfortunately, we cannot compare these results with other experiments, because, apparently, such a question was not under research previously.

During the second experiment the average decision of players was also riskier than the optimal equilibrium decision $Y = 9$ (Fig.5). Hypothesis H1 is rejected. Starting from the first rounds the participants of the second experiment understood that the decision should be riskier comparing with decisions in the first experiment (the average decision during first 5 rounds in the second experiment was 9.75; the average decision during last 5 rounds in the first experiment was 9.0). Low sampling variance of the decisions during first rounds also tells us about consciously made decisions; the decisions were not chaotic while players were trying to understand the task of the second experiment, which was changed comparing to the first experiment.

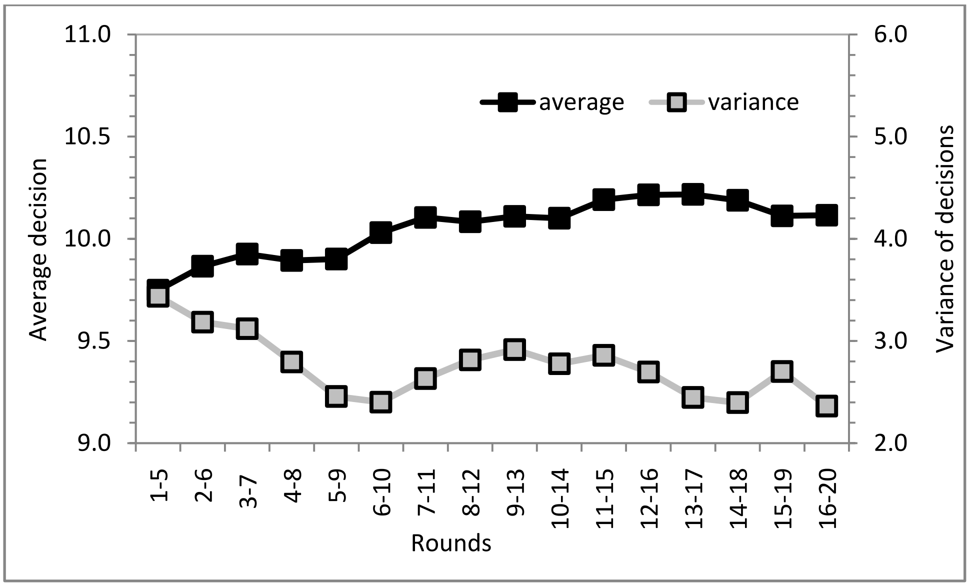

Fig.5. Average decision and variance of decisions (sampling of all participants' decisions during 5 rounds) in the second experiment. The best decision is $Y = 9$.

Moreover, by the end of the experiment the players' decisions became riskier (average decision grew to $Y_{aver} = 10.1$) and confident (sampling variance of decisions felt to 2.5). In the first experiment the variance of the number of buyers was 5, in the second experiment the variance became higher — approximately 7.5. Formally H2 should be rejected based on the results of the second experiment, and H3 should not be rejected.

The results of the third experiment are difficult to interpret and here we present only preliminary attempt to analyze the players' decisions. The evolution of the average decision (averaged during 5 rounds) in comparison with the evolution of the "optimal" decision is shown on Fig.6. Firstly, starting from the first round players made much riskier decisions ($Y_{aver} = 11.0$) comparing with the decisions in the second experiment ($Y_{aver} = 10.1$). Approximately till the 10th round the average decision was close to the optimal one.

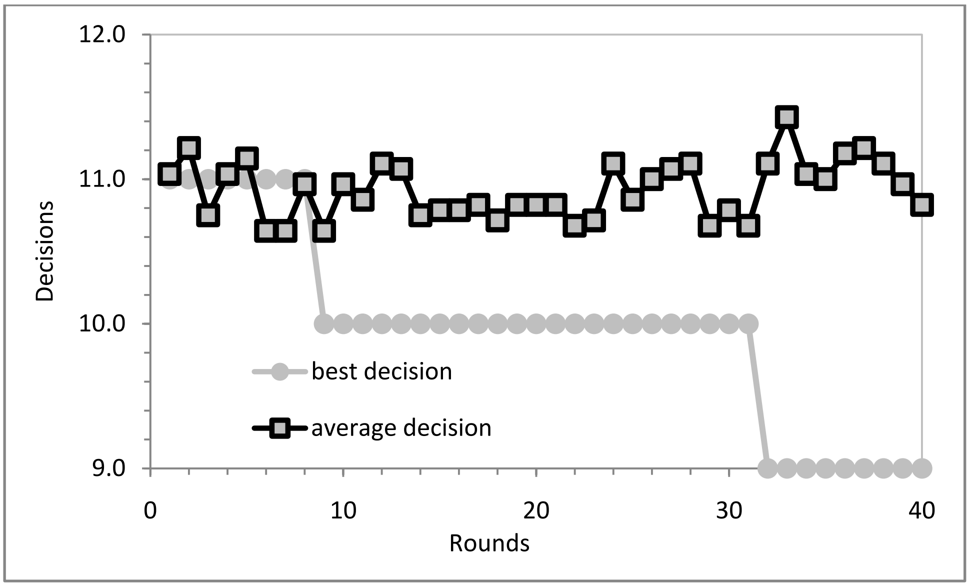

Fig.6. Average decision of participants in every round of the third experiment. The "best" decision was obtained under the condition that both players have the same number of loyal buyers in the round under consideration.

But further the average decision remained too risky, and such a behavior of participants differs a lot from the behavior of ideally rational agent. The rational player should not make risky decisions closer to the end of the experiment, because by the end of the experiment the additional profit from new captured buyers becomes very low. So it is difficult to explain the behavior of real participants. May be the players got very excited and during the experiment put a new goal for themselves – not to maximize the profit, but to capture new buyers as much as possible .

It is difficult to check statistically hypothesis H1, because optimal decision is changing from round to round. However based on the results of the experiment we can say that the average decision in first 10 rounds do not differ from the average decision in last 10 rounds (H2 should be rejected).

To check the hypothesis H3 about learning in this experiment we cannot use the variance of players' decisions. That is because by the end of the experiment some players had less than 10 loyal buyers and for such participants optimal decision was 4 or even lower. Some players increased the number of their loyal buyers till 30 or even more, and their optimal decision in such case was 13 or higher. Of course, by the end of the experiment sampling variance of all the decision had increased. To eliminate the influence of different number of loyal buyers, we introduced a new variable of "normalized decision": $Z_t = Y_t/N_t$, where Y_t— a player's decision in round t, N_t— number of loyal buyers by the beginning of round $t.$ To have a chance to compare variances in different experiments we calculated sampling variance of random variable $Z_t \cdot N$, where $N = 20$— the number of loyal buyer at the beginning of experiment.

Fig.7 illustrates the changes in sampling variance of "normalized decisions" during the third experiment. Participants got used to new conditions and, taking them into account, they probably were looking for a new strategy for decision making. By the end of the experiment their understanding of the problem had become better; that was reflected in the decrease of the variance of "normalized decisions".

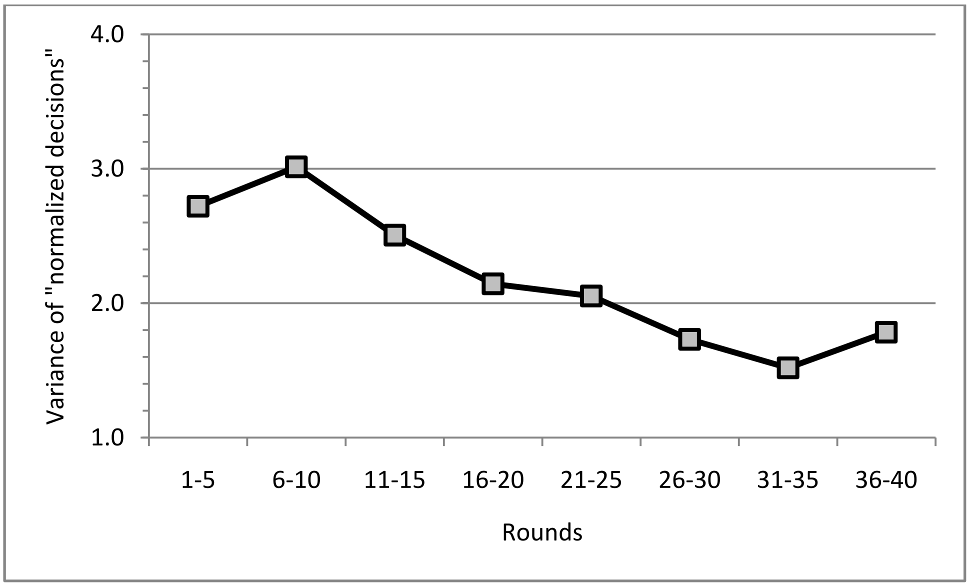

Fig.7. Variance of "normalized decisions" (sampling of all players' decisions during 5 rounds) in the third experiment.

We can try to estimate the effectiveness of players' decisions by comparing profits the players got with profit the fully informed and rational risk–neutral agent would get. Tab.2 shows profits per round, which participants got on average in each experiment, and the expected profits of the ideally rational player. Effectiveness of the decisions was calculated as the ratio between average profit got by participants and the expected profit.

Table 2. The average profit of players and the expected profit at optimal decision–making.

	Average profit (per round)	Expected profit (per round)	Effectiveness of decisions
1st experiment	103	138	75%
2nd experiment	123	155	79%
3rd experiment	86	133*	65%

*estimate provided that the number of loyal customers is always 20.

Compare the data from Tab.2, we can assume, that low effectiveness of the decisions in the first experiment is due to the search of a strategy during first rounds. In the second experiment the participants were able to find a strategy faster, that's why the effectiveness of a decision is higher. Third experiment showed a "failure" of effectiveness, probably, because of riskier decisions in the second half of the experiment. We can assume it is because of 1) high excitement during the game and substitution of goal for a game; 2) misunderstanding of the terms of the game; 3) participants were not able to compare properly the value of future profits and current losses. May be, there are another reasons; it is a field for future research.

Conclusion

Series of experiments we organized, where participants solved different variations of Newsvendor problem, and the results of similar experiments organized by other researches, show that decisions made under uncertainty do not correspond to model theoretical decisions. In all our experiment the average decisions of participants were riskier than the decision of a rational player who maximizesthe expected profit. This is not news. But surprisingly, the participants felt how the optimal decision was changing. It became riskier from experiment to experiment, and average player's decision followed this change. Perhaps the participants were not able to assess the risk accurately, but felt how much the risk has changed.

Analysis of the results illustrates that the variance of participants' decisions was decreasing significantly after first 5–10 rounds of each experiment. We interpret it as learning by doing. Observing the current profit the players gained more experience in each new round of the game and made their decisions more deliberately and consciously. Perhaps a player acted in accordance with a certain own strategy of decision making under uncertainty. This strategy was individual and not the best on average, but it was self–intuitive for them. Seems that during third experiment the players replaced the mysterious aim “to maximize the expected profit” by clear aim “to capture the competitor's buyers”.

Acknowledgments

The authors acknowledge Saint–Petersburg State University for a research grant 16.23.976.2012. We want to thank Zurich University for providing us with Z–Tree programto organize the experiments. We also want to thank the administration and staff of National Research University “Higher School of Economics” (HSE Campus in St.–Petersburg) for their help in organization of experiments. Special thanks to Professor Zenkevich N. from Graduate School of Management of St. Petersburg State University for useful discussion.

References

Alfares, H., &Elmorra, H., (2005). The distribution–free newsboy problem: Extensions to the shortage penalty case. International Journal of Production Economics, 93/94(1), 465-477.

Archavsky, V., &Okulov, V., (2012). An Experiment on Decision Making under Risk and Uncertainty. Working Paper 13(R)–2012. Graduate School of Management, St. Petersburg State University: SPb.

Atkins,P., Wood,R., &Rutgers, P. (2012). The effect of feed back formation dynamic decision–making. Organizational Behavior & Human Decision Processes, 88(2), 587-604.

Benzion,U., Cohen,Y.,Peled,R.,&Shavit,T. (2008). Decision–making and the news vendor problem: An experimental study. Journal of the Operational Research Society, 59(9), 1281-1287.

Bolton,G., &Katok, E., (2008). Learning–by–doing in the news vendor problem: A laboratory Investigation of the role of experience and feedback. Manufacturing & Service Operations Management, 10(3), 519-538.

Bostian,A., Holt,C., &Smith,A. (2008). Newsvendor “Pull–to–Center” Effect: Adaptive Learning in a Laboratory Experiment. Manufacturing & Service Operations Management, 10(4), 590-608.

Charness, G., Karni, E.,& Levin, D. (2007). Individual and group decision making under risk: An experimental study of Bayesian updating and violations of first–order stochastic dominance. Journal of Risk &Uncertainty, 35(2), 129-148.

Cooper, D., &Kagel, J. (2005). Are Two Heads Better Than One? Team versus Individual Play in Signaling Games. American Economic Review, 95(3), 477-509.

Cox, J. C. (2010). Trust, reciprocity, and other–regarding preferences: groups vs individuals and males vs females. In R. Zwick R. & A. Rapoport (Eds.), *Experimental Business Research.* Kluwer Academic Publishers: Boston, Mass.

Eeckhoudt,L., &Gollier,C. (1995). *Risk evaluation, management and sharing*. Harvester–Wheatsheaf:NY.

Feng,T., Keller,L., &ZhengX. (2011). Decision making in the news vendor problem: Across–national laboratory study. Omega, 39(1), 41-50.

Fischbacher,U.(2007). Z–Tree: Zurich Tool box for Ready–made Economic Experiments. Experimental Economics, 10(2), 171-178.

Fisher,M., &Raman,A. (1996). Reducing the cost of demand uncertainty through accurate response to early sales. Operations Research, 44(1), 87-99.

Gavirneni,S., &Xia,Y. (2009). Anchor Selection and Group Dynamics in Newsvendor Decisions – A Note. Decision Analysis, 6(2), 87-97.

Ho, T., Lim, N., &Cui, T. (2010). Reference Dependence in Multilocation Newsvendor Models: A Structural Analysis. Management Science, 56(11), 1891-1910.

Khouja,M. (1999). The single–period (news–vendor) problem: Literature review and suggestion for further research. Omega, 27(5), 537-553.

Kocher, M., &Sutter, M. (2005). The Decision Maker Matters: Individual Versus Group Behavior in Experimental Beauty–Contest Games. The Economic Journal, 115(500), 200-223.

Moritz, B., Hill, A., &Donohue, K. (2013). Individual differences in the newsvendor problem: Behavior and cognitive reflection. Journal of Operations Management, 31(1/2), 72-85.

Mostard,J., & Teunter,R. (2006). The news boy problem with resalable returns: A single period model and case study. European Journal of Operational Research, 169(1), 81-96.

Nagarajan, M., &Rajagopalan, S. (2009). A Multiperiod Model of Inventory Competition.Operations Research, 57(3), 785-790.

Netessine, S., &Rudi, N. (2003). Centralized and Competitive Inventory Models with Demand Substitution. Operations Research, 51(2), 329-335.

Ovchinnikov, A., Katok, E., Moritz, B., & Quiroga, B. (2013). How to Compete Against a Behavioral Newsvendor. Darden Business School Working Paper No. 2236503. Available at SSRN:http:// ssrn.com /abstract=2236503

Parlar, M. (1988). Game Theoretic Analysis of The Substitutable Product Inventory Problem with Random Demands. Naval Research Logistics, 35(3),397-409.

Petruzzi,N., & Dada, M. (1999). Pricing and the newsvendor problem: A review with extensions. Operations Research, 47(2), 183-194.

Qin,Y., Wang,R., Vakharia,A., Chen,Y., &Seref,M. (2011). The newsvendor problem: Review and directions for future research. European Journal of Operational Research, 213(2), 361-374.

Schweitzer,M., &Cachon,G. (2000). Decision Bias in the Newsvendor Problem with a Known Demand Distribution: Experimental Evidence. Management Science, 46(3), 404-420.

Véricourt,F., Bearden,J., & Filipowicz,A. (2013). Sex, Risk and the Newsvendor. Journal of Operations Management, 31(1/2), 86-92.

Wang, Q.,&Parlar, M.(1994). A three–person game theory model arising in stochastic inventory control theory. European Journal of Operational Research, 76(1), 83–97.

Zheng,H., &Liu,J. (2011). Fuzzy Newsboy Problem with Random Variables in a Supply Chain Environment. International Journal of Information and Management Sciences, 22(1), 27-42.