An Experimental Study of Posted Prices

Seungbeom Kim, Sriram Dasu

Marshall School of Business, University of Southern California, Los Angeles, CA 90089, United States

[seungbek@usc.edu, dasu@marshall.usc.edu]

In industries with perishable goods dynamic pricing schemes are often used and many models have been proposed to maximize seller's revenues. These models are based on a range of assumptions about how customers make buying decisions. Some assume that customers are myopic and are not forward looking while others assume that customers are strategic and anticipate future prices. In the second stream, a standard assumption is that strategic customers maximize their expected utility. All these models predict that customers will purchase in the current period if and only if their valuation exceeds some threshold. We use laboratory experiments to gain insights into how customers make purchase decisions when they have the option of buying at a higher price or waiting for a lower price but incur the risk of the product being out of stock. We find that the quantal response model (QRM), a quasi-rational model, provides a more accurate description of customers' decisions. We also explore how customers learn in these types of settings.

[Key Words: Behavioral Operations, Markdown, Bounded Rationality, Strategic Customers, Revenue Management, Posted Pricing, Quantal Response Equilibrium, Learning]

1. Introduction

Marcy finds a dress priced at $150 that she likes. From her past experience she knows that two weeks later she could buy the dress for a 10% discount provided it is still available. Should Marcy buy the dress or should she wait? Assumptions about how customers make these decisions lie at the core of optimal pricing and quantity decisions for firms selling perishable goods. A number of authors have developed pricing strategies and stocking decisions employing game theoretic approaches (Aviv and Pazgal 2008; Elmaghraby et al. 2008; Yin et al. 2009; and Shen and Su 2007). In these papers the seller is a Stackelberg leader and customers are followers. Customers have private valuations and try to maximize their expected surplus by anticipating future prices, timing of price changes, and stock-out probabilities. Models differ by types of pricing policies that are considered by firms and factors that influence customers' decision. However, there are some core assumptions. Customers discover the equilibrium strategies of other players, which enables them to estimate stock out probabilities and price paths. They purchase if the expected utility of buying now is greater than that obtained by waiting.

There is now growing experimental evidence, in a range of decision problems, that customers do not employ strategies predicted by game theoretic models. While the rational utility maximization models (RUMM) are elegant they require each customer to either compute or discover through careful observation the strategies of other customers and the firm and to act optimally. In the context of Newsvendor problems researchers (Bolton and Katok 2008; Lim and Ho 2007; Ho et al. 2010; Wu and Chen, 2013; and Chen et al. 2012) have found that bounded rationality models provide a more accurate description of the actual buying behavior.

In this paper we employ experiments to investigate buying behavior in a two-period posted pricing scheme (Aviv and Pazgal 2008; and Dasu and Tong 2010). Subjects are aware of their private valuation, the distribution of the valuations of all other customers, the number of other customers, and the number of units for sale. They can buy at a higher price or wait for a discount but face the risk of a stock-out. This model was analyzed by Dasu and Tong (2010) assuming that customers are fully rational and risk neutral. We are interested in the following questions:

1. How do customers evaluate the option of a certain current payoff against an uncertain future payoff?
    a) Do they conceptualize the problem as a utility maximization problem with known (or accurately computed) probabilities?
    b) Is their decision making consistent with boundedly rational models that assume that individual customers' utility estimates include a random error term.
2. How do customers discover stock-out probabilities?
    a) Do they compute stock-out probabilities or do they learn about stock-out probabilities through observation?
    b) Are there biases in how they estimate probabilities?
3. How do customers learn to make decisions in these types of environments?

In the following section we present the theoretical background and related literature for this study. In section 4, we summarize the theoretical predictions under alternative models that were identified in section 3. Design of the experiments is presented in section 4. In sections 5, 6, 7, and 8 we present our experimental findings. We conclude in section 9.

2. Theoretical Background and Related Literature

Dynamic pricing modelers assume that customers are either myopic and purchase whenever their valuation exceeds the current price of the product (Gallego and Van Ryzin 1994) or customers are strategic and anticipate future opportunities to purchase the product at a lower price (Aviv and Pazgal 2008) and time their purchases so as to maximize expected utility. Cachon and Swinney (2009) identify a third type of customer that they refer to as the bargain hunter, who always waits for the lowest price. Talluri and Van Ryzin (2005), Elmaghraby and Keskinocak (2003), and Shen and Su (2007) provide comprehensive reviews of dynamic pricing models.

Researchers have considered different types of pricing schemes. In the posted pricing scheme the firm announces the prices at the beginning of the season and in the contingent pricing scheme (Cachon and Swinney 2009) the firm's pricing strategy evolves dynamically based on observed sales. Aviv and Pazgal (2008) find that neither approach dominates the other. Researchers (Liu and Van Ryzin 2008; and Dasu and Tong 2010) have also found that bulk of the benefits of dynamic pricing is derived by a single price change. Based on these findings, in this paper we restrict ourselves to a two-period posted pricing scheme. We further assume that the first period price (P_1) is higher than the second period price (P_2). Although, a declining price path is not always optimal (Su 2007), it is optimal in our setting (Dasu and Tong 2010). This is also a common practice in retailing.

Consistent with prior literature, we assume that each customer has a private valuation v, that corresponds to his or her willingness to pay. We are interested in how customers make the following decision:

Buy at high price and derive a value of $v - P_1$*, or wait and incur the risk of earning zero if there is a stock out or earning a value of* $v - P_2$ *otherwise.*

This type of decision lies at the core of dynamic pricing models. To make this decision subjects have to determine (i) the probability of getting the product in the second period (π) and (ii) compare the option of getting ($v - P_1$) for sure or getting with probability $\hat{\pi}(v - P_2)$ and 0 with probability 1-$\hat{\pi}$ where $\hat{\pi}$ is their estimated in-stock probability.

Liu and Van Ryzin (2008) provide a review of how different researchers have modeled decision making by strategic customers. They identify two broad approaches. In the first approach, strategic customers explicitly compute the equilibrium strategies of other players in the game and compute the stock-out probabilities. In the second approach, customers do not compute the equilibrium strategies. Instead, it is assumed that the properties of the equilibrium such as stock out probabilities and price paths are observable characteristics of the market and customers learn this through observation. Liu and Van Ryzin (2008) argue that the first approach is more applicable in small markets and the second in large markets where the impact of a single customer on the overall market is infinitesimal. Anderson and Wilson (2003) suggested that customers may be able to learn about the probabilities through third parties.

Researchers have long recognized the limitations of RUMM (Mackelvey and Palfrey 1995; and Kahn and Baron 1995). These models do not accurately capture either the decision rules employed by the customers or the factors that influence these decisions. Kahn and Baron (1995) find that even when the stakes are high subjects fail to evaluate trade-offs in a manner proposed by utility theory. In recent years there have been a number of experimental studies examining how decisions are made by subjects faced with Newsvendor type problems. Experimental data has provided little support for rational decision models (Ho and Zhang 2008; Ho et. al 2010; and Wu and Chen 2013). Wu and Chen (2013) find that the quantal response frame-work, quasi-rational models, provides a more accurate description of the decisions.

In RUMM, it is assumed that each subject accurately estimates the utility of each of the decisions; i.e. $U(v - P_1)$ if they choose to buy in period 1 and $\pi U(v - p_2)$ if they choose to wait. In quantal response models (QRM) (McFadden 1976; and McKelvey and Palfrey 1995), it is assumed that each of these utilities is estimated with some error.

$$\widehat{U}(v - p_1) \ = \ U(v - p_1) + \ \varepsilon_{1,k} \tag{1}$$

$$\pi\widehat{U}(v - p_2) \ = \pi \ U(v - p_2) + \ \varepsilon_{1,k} \tag{2}$$

Due to the error terms $\varepsilon_{i,k}$, the likelihood of picking the dominant decision depends on the difference between the two options. If the error terms ($\varepsilon_{i,k}$) in equations (1) and (2) are unbiased and have extreme value distributions, then the probability of purchasing in the first period is given by a Logit model (McKelvey and Palfrey 1995).

If subjects are boundedly rational then they have to learn the in-stock probability and to make good decision. Camerer and Ho (1999) present an overview of different approaches that have been proposed for how subjects develop their strategies. They also propose a general model called the experienced weighted attraction (EWA) model that combines belief based models and reinforcement models. In reinforcement models the likelihood of a subject choosing a strategy in a game depends on the historical pay-offs due to that strategy. Subjects, however, do not keep track of strategies employed by other players. In the belief based models subjects keep track of strategies employed by other players. A subject selects his or her strategy based on what (s)he believes will be employed by others. The EWA model combines the beliefs and reinforcement models. EWA model also incorporates the possibility of fictitious play; i.e. when subjects carry-out what-if analysis for strategies that they did not employ.

In our experiment subjects are not informed about other players' strategies but only observe whether or not the product is in stock. We employ an exponential learning model to measure how rapidly subjects discover the in-stock probability. We explore if waiting and experiencing a stock-out stings more than when they buy in the first period and find the product was available in the second period. Finally, we employ the EWA model to learn about the rate at which subjects learn to make decisions and the role of fictitious play.

Experimental methods have been used to gain some insights into how customers make these complex trade-offs (Güth et al. 1995; Güth et al. 2004; Bearden et al. 2008; and Bendoly 2012). Much of the early experimental work on dynamic pricing was concerned with durable goods and focused on the discount rate (Güth et al. 1995; and Güth et al. 2004).

Bearden et al. (2008) use experiments to study the behavior of sellers. In their experiments financially motivated subjects have to sell a small number of objects over a finite time horizon. At various points in time a customer arrives and offers to buy one unit at some price. The subjects have to decide whether or not to accept the bid. The objective is to maximize the sales revenue over the time horizon. They find that subjects use policies that are sub-optimal but are qualitatively similar in structure to optimal policies. In our experiments subjects are customers and have to decide whether to purchase or wait for a lower price.

Osadchiy and Bendoly (2011) study how customers respond to a two-period posted pricing mechanism that is similar to the problem we study. We assume that all customers are present from the start, while they inform their subjects that customers arrive according to a Poisson process. In each trial in their experiment, subjects are informed about their time of arrival and their private valuation. In half the trials subjects are also informed about the likelihood of getting the product. Based on subject's purchasing decisions the authors try to determine whether or not each subject is myopic, strategic, or a bargain hunter. They find that in aggregate, information about the stock-out probability influences customer's behavior**.**

Although our context is similar to that of Osadchiy and Bendoly (2011), the objectives of our experiments and how they are carried out are vastly different from theirs. In our experiment the stock-out probability remains the same but the private valuations change from one trial to the next. The same valuations, however, are repeated a few times. As a result we are able to gain very different insights. Because subjects have to make repeated decisions under the same parameter settings we have more information about whether or not they are myopic, bargain hunters, or strategic. We are also able to detect biases in how they estimate stock-out probabilities and shed some light on how they learn. We are able to characterize the aggregate demand curve.

3. Decision Making Under Alternative Assumptions

We implemented the two-period posted pricing problem analyzed by Dasu and Tong (2010) and Aviv and Pazgal (2008). The seller who is a monopolist announces two prices, P_1 and P_2 for the first period and the second period, respectively, with $P_1 > P_2$. There are two types of customers: high type and low type. Low type customer's valuation lies between P_2 and P_1, while the high type customer's valuation exceeds P_1. The number of high type customers (N_1), the number of low type customers (N_2), and the number of units for sale (K) are common information. The number of high type customers is less than the number of units for sale ($N_1 < K$); however, the total number of customers is greater than the number of units for sale ($N_1 + N_2 > K$); therefore, there is a possibility of stock out in the second period. High type customers and low type customers valuations are drawn from Uniform[P_1,200]and Uniform[P_2,P_1], respectively. Customers' valuations are private but the distribution of the valuations is common information. We assume that valuations do not change from one period to the next and discounting is not relevant because the time gap between the first and second period is too small. High type subject can buy in the first period or buy in the second period provided there is supply. The expected payoff for a risk neutral customer with valuation v, in the first period and the second period is $(v - P_1)$ and $\pi_2(v - P_2)$, respectively. Value of π_2, the probability of getting the product in the second period, depends on how consumers make their decisions. If consumers are risk neutral rational utility maximizers, then theory predicts that there is a unique Bayesian Nash equilibrium for equation (3) which is given by the smallest solution τ in the range $[P_1, 200]$. All high type customers with valuations $v \in [\tau, 200]$ will bid in period 1 and the customers with valuation $v \in [P_2, \tau]$ will bid in period 2. We call this a threshold policy with a threshold at τ.

$$(\tau - P_1) = \pi_2(\tau)(\tau - P_2) \tag{3}$$

, where $\pi_2(\tau)$: the probability of getting the product in period i when all the customers use the threshold policy with a threshold at τ.

In our experiments subjects do not interact with each other. They are informed that the decisions for other buyers are made by the computer using a pre-determined rule and that this rule is not based on either the private valuation of the subject or her decision. The computerized rule tries to maximize the expected profit of each of the other buyers. Further this optimization is done without considering the specific valuation or decision of all the other simulated buyers. In actuality the computer merely determines whether or not product is in-stock based on the probability given by equation (3).

Different assumptions about customer's decision making processes result in different predictions about observed buying behavior. Below we list the alternatives in the form of conjectures and comment on what we expect to observe.

C1. Some subjects are myopic and some others are bargain hunters: Myopic customers will always purchase in the first period. Bargain hunters will always wait.

C2a. Subjects are risk neutral and employ RUMM: They will always buy in the first period if their valuation exceeds the threshold predicted by equation (3) and wait otherwise.

C2b. Subjects are risk averse and employ RUMM: These subjects purchase in period 1 if and only if the utility of buying in the first period exceeds the expected utility of buying in the second period:

$$U_1(v - P_1) \geq E[U_2(v - P_2)] \quad (4)$$

Each subject has his or her own threshold level that depends on his or her utility function. If the subject's valuation exceeds this threshold then then the subject will always buy in period 1 and will wait otherwise.

C3. Subjects are boundedly rational: Subjects will buy in the first period with some probability that depends on their private valuation. The likelihood of buying in the first period increases with the valuation. Subjects' buying decisions will not be consistent with a threshold policy.

If subjects are boundedly rational then they have to learn the in-stock probability and learn how to make decisions.

C4. Losses loom larger than gains: The pay-off in the first period acts as a reference point for determining losses and gains. Trials in which a subject waits and discovers that a product was not in stock will be waited more heavily. As a result, *subjects will learn faster when the in-stock probability is lower*.

C5. The belief model and the reinforcement model are employed by subjects: If subjects learn rapidly then we will observe that each subject converges quickly to one of the alternatives and as a result the percentage that buys in the first period will stabilize.

4. Design of the Experiments

265 financially motivated subjects participated in four sets of experiments termed: **SIM, SIM-I, CONV,** and **LOT**. **SIM** is a simulation of the two-period posted pricing problem that is in accordance with the model described in section 3. **SIM** consisted of 3 treatments and each treatment in turn consisted of a number of trials. In each trial in a treatment private valuations changed. Each treatment in **SIM** differed only in terms of the size of the market (the number of low type customers). The number of goods for sale, and the purchase prices (P_1 and P_2) were not altered. As a result, the in-stock probability in the second period was different in each treatment. The in-stock probabilities were 50%, 83%, and 16% in treatment 1, 2, and 3. We will refer to these as **SIM1**, **SIM2**, and **SIM3**. In **SIM-I** subjects were informed of the in-stock probability in the second period. **SIM** and **SIM-I** were identical otherwise. We refer to the 3 treatments in **SIM-I** as **SIM1-I**, **SIM2-I** and **SIM3-I**.

The focus of our study is on **SIM.** We want to identify whether or not subjects employ RUMM or whether their decisions are better described by quasi-rational models such as quantal response models. By comparing **SIM** and **SIM-I,** we sought to understand the impact of in-stock probability information. **SIM** and **SIM-I** also allow us to determine if some subjects are consistently myopic or bargain hunters. Each subject participated in only one of the treatments to eliminate any learning effect that carry over from one treatment to the next.

LOT consisted of 5 surveys. We sent an email to the participants of the **SIM** and **SIM-I** experiment two or more day after they participated in the study asking them to fill out a survey. During the **SIM** and **SIM-I** experiments we informed subjects that they will be asked to participate in a different type of experiment and collected their email address. 142 out of 179 subjects who participated in **SIM** or **SIM-I** (79% of the subjects) participated in the **LOT** experiment.

LOT was modeled after the surveys proposed by Holt and Laury (2002) to measure risk attitude. Decisions in the survey corresponded to the decisions arising in the **SIM** experiment. In each survey subjects had to make 16 decisions. While making decision j subjects had to choose between Option B that guaranteed $5j or a risky Option A that paid $40 + $5j with some probability and 0 otherwise. Thus decision j corresponded to a resale price (v) of $120 +$5j in **SIM.** The 5 surveys have the same set of payoffs but with different chance of success for Option A: 16%, 33%, 49%, 66%, and 83% (Table 1). To reduce ordering effect we randomized the order of the 5 surveys. The full surveys are in the appendix.

Table 1: Parameters for the **LOT** experiment

SURVEY	Probability of success in the risky option A	Corresponding Treatments in **SIM**
LOT16%	16%	**SIM3, SIM3-I**
LOT33%	33%	**N/A**
LOT49%	49%	**SIM1, SIM1-I**
LOT66%	66%	**N/A**
LOT83%	83%	**SIM2, SIM2-I**

By comparing **LOT**, **SIM** and **SIM-I** experiments we can determine whether subjects purchasing behavior is consistent with RUMM or with QRM. **LOT** gives us information about individual utility functions. Subjects were recruited from Amazon Mechanical Turk’s (AMT) online pool. Paolacci et al. (2010) and Buhrmester et al. (2011) provide an overview of this tool and its reliability. Readers are also referred to work by Archak et al. (2011), Kaufmann et al.(2013) and Toubia et al.(2013) for additional validation of AMT for experiments such as ours.

Subjects received a base payment of $2 and a bonus payment up to $2 that was based on their performance in **SIM (SIM-I)** experiment. To reduce portfolio effect, we informed the subjects that 10 trials would be randomly selected and that they will be paid based on their performance in these trials (Chen and Hogg, 2008; Katok, 2011). In each trial, during the experiments, subjects learnt whether or not the product was available in the second period. To further minimize portfolio effects we informed subjects the total number of trials is between 50 and 70. In the **LOT** experiment, subjects received a base payment of $0.5 and a bonus payment up to $1. To reduce portfolio effect, we informed that only one of the 80 decision questions would be selected. The lottery corresponding to that decision would be simulated on the computer and the outcome would determine the bonus.

The final set of experiment named **CONV** also consisted of 3 treatments. In all three treatments the in-stock probability was 49%. Unlike the other experiments, in each trial in a treatment the private valuation or resale value was the same. The three treatments in **CONV** differed in terms of the resale value. **CONV** was used to gain insights on how subjects learnt and whether or not they would converge to one of the two options – wait or buy in the first period.

All subjects were given an overview of the experiment and were asked to answer questions to ensure that they understood the objective of the experiment. Screen shots of the simulation tool used in **SIM, SIM-I,** and **CONV**, the survey used in **LOT**, and the detailed instructions given to the subjects are available in the online appendix. The average time the subjects spent in the **SIM (SIM-I**, **CONV**) experiment and on the **LOT** survey was approximately 21minutes and 7 minutes, respectively. Table 2 below summarizes the experiments and their objectives.

Table 2: Description of the experiments

Experiment Type	Number of Treatments	Description	Objective
SIM	3	Simulation of a two-period posted pricing problem. Valuations varied from one trial to the next. In-stock probability is not given.	To identify the decision rules employed by subjects when they have to learn the in-stock probability
SIM-I	3	Simulation of a two-period posted pricing problem. Valuations varied from one trial to the next. In-stock probability is given.	To identify the decision rules employed by subjects when exact in-stock probability is offered
CONV	3	Simulation of a two-period posted pricing problem. Valuations fixed from one trial to the next. In-stock probability is not given.	To investigate how subjects learn and whether they would converge to one strategy
LOT	5	Surveys in which the decision problem encountered in **SIM** is presented as a lottery	To explore whether the decisions in **SIM** are structurally similar to decisions made while evaluating lotteries

5. Analysis of LOT

Figure 1 shows the aggregate data for **LOT16%**, **LOT49%** and **LOT83%** treatments. For different decisions j, this figure shows the percentage choosing the certain option, Option B. The solid line is the percentage we would expect to observe under RUMM if all subjects are also risk neutral.

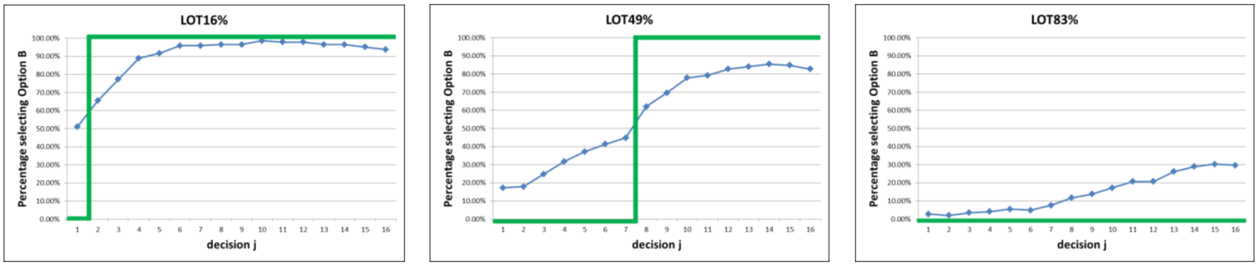

Figure 1: Aggregate results for **LOT** experiment

The chi-square goodness of fit test rules out the possibility that all subjects are rational, risk neutral and they employ the threshold policy predicted by theory. (All p-values = 0.00). Figure 1 does not, however, rule out the possibility that subjects may have different thresholds due to individual differences in risk aversion. Any subject who always chooses Option B if the decision number j exceeds some number, and always selects Option A below this number is considered to have employed a threshold policy. Subjects who always chose Option A or always selected Option B were also included in this set. The percentage employing threshold policies given in Table 3 below is similar to what is reported by Holt and Laury (2002). Nevertheless, it is surprising that 26.76% do not employ a threshold policy when you consider all the surveys together. The decision problem is one trivial step away from how the expected utility trade-off equation is evaluated. Yet RUMM is not consistently employed by at least 25% of the subjects.

Table 3: Percentage of subjects employing a threshold policy in the **LOT** experiment (total: 142)

Treatment	Percentage employing a threshold policy	Percentage always choosing option A	Percentage always choosing option B
LOT16%	129 (90.85%)	0 (0.00%)	61 (42.96%)
LOT33%	126 (88.73%)	2 (1.41%)	29 (20.42%)
LOT49%	116 (81.69%)	6 (4.23%)	10 (7.04%)
LOT66%	122 (85.92%)	22 (15.49%)	6 (4.23%)
LOT83%	130 (91.55%)	89 (62.68%)	0 (0.00%)
In ALL	104 (73.24%)	0 (0.00%)	0 (0.00%)

6. SIM: Simulation of the Two-Period Posted Pricing Problem

The parameters for the three treatments are given in Table 4. The threshold values (τ), and the in-stock probabilities (π) are based on the equilibrium for a game among risk neutral rational customers. In each trial using this probability the computer randomly determined whether or not the product was in stock. Resale prices (v) ranged between ($P_1 + 5$) and 200 in steps of 5.

Table 4: Parameters for experiment set **SIM**

Parameters	**SIM1**	**SIM2**	**SIM3**
Number of High Type Customers(N_1)	10	10	10
Number of Low Type Customers(N_2)	26	14	72
Number of units for sale(K)	20	20	20
First period price(P_1)	120	120	120
Second period price(P_2)	80	80	80
In-stock probability in the second period(π)	0.49	0.83	0.16
Threshold(τ)	160	>200	128
Number of trials	60	60	60

6.1 ANALYSIS OF THE SIM EXPERIMENT

Figure 2 shows the aggregate data for each of the 3 treatments. For different resale prices, this figure shows the percentage of decisions that resulted in the subjects buying in the first period at a price P_1. The solid line is the percentage that will be observed under RUMM if all subjects are risk neutral.

Figure 2: Aggregate results for **SIM** experiment

Observation 1: We reject the conjecture (C2a) that subjects employ any threshold policy at the aggregate level (All p-values = 0.00).

6.2 RUMM with Risk Aversion: Comparison of LOT and SIM

We next test whether buying behavior is consistent with RUMM with different levels of risk aversion (C2b). At the individual level we measured the percentage of subjects employing a threshold policy. Under 7% of the subjects employed threshold type policies. We also employed a more lenient criterion by counting subjects who deviated from a threshold policy at most twice.

Table 5: Percentage of subjects employing a threshold policy

Treatment	Threshold policy	
	For all resale prices (%)	Except for two resale prices (%)
SIM1	0 (0.0%)	5 (16.7%)
SIM2	2 (6.9%)	7 (24.1%)
SIM3	2 (6.7%)	11 (36.7%)

A significantly lower percentage of subjects adopted threshold policies in **SIM** relative to **LOT**. To the extent that **LOT** is more consistent with RUMM with risk aversion it is useful to compare for each subject the percentage of times they bought in the first period in **SIM** and the percentage of times they choose Option B in **LOT.** For additional justification for this metric readers are referred to Holts and Laury (2002). In Figure 3 each point corresponds to one subject, the x-axis is the percentage of decisions that the subject select option B, and the y-axis is the percentage of times the subject bought in the first period.

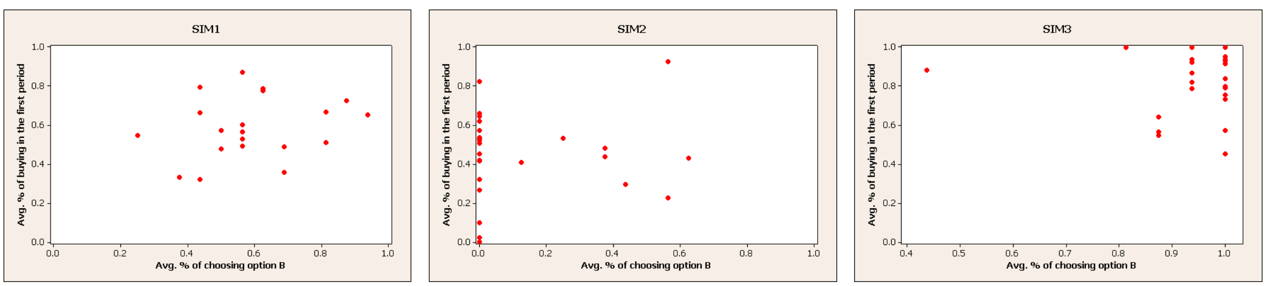

Figure 3: Comparison of buying behavior in **LOT** and **SIM** experiments

The coefficient of determinations (Adjusted R^2) of the linear regression for these points are 0.0%, 0.0%, and 0.0% for **SIM1, SIM2, and SIM3**, respectively. Therefore, we fail to find evidence that individuals employ the same decisions in **LOT** and **SIM**. There is a possibility that in the early trials the subjects focus on learning the in stock probability so we also conducted the same analysis with the last 30 trials. Here too we did not find strong linear relationships between the decisions in **SIM** and **LOT**. All of the coefficients of determinations were less than 15%.

Observation 2: Subjects do not employ the same decision strategy in LOT and SIM experiments. Even after accounting for differences in risk attitude, we reject conjecture C2(b) that subjects employ RUMM.

Although individual decisions in **LOT** and **SIM** are not the same, in aggregate **LOT** and **SIM** demand curves may be similar. Figure 4 below shows the percentage that buy in the first period in **SIM** for different resale values and the percentage that choose option B (the risk free option) in **LOT**.

Figure 4: Fitting **SIM** with **LOT**

The goodness of fit of the **LOT** curve to **SIM** observations is given in Table 6 below.

Table 6: Estimating **SIM** experiment with **LOT** result

	SIM1	**SIM2**	**SIM3**
Estimator	LOT 49%	LOT 83%	LOT 16%
Log-likelihood	-447.54	-696.94	-352.69
Chi Square Test P value	0.2853	0.0000	0.0001
Root Mean Square Errors	0.0681	0.3162	0.1306

Except for **SIM1**, Chi-Square Test P values are close to zero. In **LOT** subjects are informed of the probability of the risky option and hence the three **LOT** curves differ from each other. On the other hand in **SIM**, subjects have to discover or compute the in-stock probabilities. As a result we see less divergence among the three **SIM** curves relative to the **LOT** curves. Similarity between the **LOT49%** curve and **SIM1** curve suggests that subjects in **SIM** anchor on 50%.

6.3 Quantal Response Models (QRM)

Recall that payoffs of buying in the first period and waiting are $(v - P_1)$ and $(v - P_2)$, respectively, where v is resale price, P_1 and P_2 are the prices. We employ the exponential utility function to capture the risk attitude. Readers are referred to the work by Holt and Laury (2002) for a justification of this utility function. We let r denote the risk aversion parameter.

$$U(v-P_i)=\frac{1-exp\left(-r(v-P_i)\right)}{r} \tag{5}$$

, where $r>0, r=0$ and $r<0$ would imply risk aversion, risk neutrality, and risk preference, respectively and V is the payoff. The quantal response model assumes that the subject's estimate the utility of the two options with some error.

Let $R_1=U(v-P_1)$, $R_2=\pi U(v-P_2)$, and $R(v,P_1,P_{2,}\pi)=R_1-R_2=U(v-P_1)-\pi U(v-P_2)$. Under the QRM model the probability of purchasing in the first period when the resale price is v, is given by:

$$Pr(\ v,P_1,P_2,\pi,\gamma,r)=\frac{exp(\gamma R(v,P_1,P_2,\pi))}{1+\ exp(\gamma R(v,P_1,P_2,\pi)} \tag{6}$$

The rationality parameter (γ) and risk parameter (r) are estimated by maximizing the log likelihood function:

$$max_{\gamma,r}\sum_{v\in V}[O_1(v)Log(Pr(\ v,P_1,P_2,\pi,\gamma,r))+O_2(v)Log(1-Pr(\ v,P_1,P_2,\pi,\gamma,r))] \tag{7}$$

, where $O_i(v)$ is the number of observation in which subjects buy in period i, i = 1, 2, when the resale price is v.

The QRM model has a two parameters and the added flexibility should enable us to find a good fit. To control for this we use out of sample data to measure the quality of the fit. We randomly selected 50% of the subjects to fit the QRM model and tested the fit on the remaining 50% of the subjects. Table 7 summarizes the results for 5 different random partitions.

Table 7: Out of sample goodness of fit test for QRM models

Treatment	Partition No.	Rationality Parameter (γ)	Risk Aversion Parameter (r)	Log-Likelihood	Chi-square Test p value	Root Mean Square Deviation
SIM1	1	0.1029	0.0110	-206.15	0.9851	0.0622
	2	0.1113	0.0116	-212.36	0.8963	0.0765
	3	0.1092	0.0183	-195.83	0.9946	0.0559
	4	0.0913	0.0116	-195.82	0.5403	0.0872
	5	0.0852	0.0133	-193.34	0.1647	0.1046
SIM2	1	0.1226	0.0294	-212.16	0.0349	0.1096
	2	0.1784	0.0374	-215.05	0.5939	0.0842
	3	0.1107	0.0228	-223.44	0.0000	0.1612
	4	0.3161	0.0437	-252.81	0.0000	0.1682
	5	0.2791	0.0368	-249.79	0.0000	0.1346
SIM3	1	0.0594	0.0000	-152.36	0.5875	0.0963
	2	0.0652	0.0000	-160.34	0.7238	0.0905
	3	0.0697	0.0003	-163.81	0.9823	0.0651

4	0.0845	0.0192	-151.13	0.3454	0.1058
5	0.0783	0.0055	-170.98	0.8935	0.0739

In **SIM1** and **SIM3**, QRM model has a very strong estimation power. Chi-square tests have p values well above 0.05 and RMSDs are also below 0.10 for most of the partitions. In **SIM2**, Chi-square test p values are greater than 0.01 in only two of the five partitions and the maximum RMSD is 0.1682.

To take a deeper look at the **SIM2**, we compare QRM to RUMM and LOT by using Akaike Information Criterion (AIC) and Bayesian Information Criterion (BIC) for model comparison. AIC is evaluated as $\text{AIC} = (-2) \cdot LL + 2 \cdot k$ (Akaike 1974) and BIC is evaluated as $\text{BIC} = (-2) \cdot LL + k \cdot Log(n)$ (Swartz 1978; Liddle 2008), where LL denotes the maximum log likelihood, k is the number of parameters estimated and n is the number of datapoints used in the fit.

Table 8: Model comparison for **SIM2**: QRM, RUMM and LOT

Partition No.	Model	RMSD	Chi-square Test p value	Log-Likelihood	Number of Parameters	AIC	BIC	Rank
1	QRM	0.109574813	0.034903955	-212.16	2	428.33	426.73	1
	RUMM	0.508978457	0.0000	-1492.02	0	2984.04	2984.04	3
	LOT	0.340481916	0.0000	-317.09	0	634.18	634.18	2
2	QRM	0.084170186	0.593911877	-215.05	2	434.11	432.52	1
	RUMM	0.451081046	0.0000	-1352.02	0	2704.04	2704.04	3
	LOT	0.287227883	0.0000	-298.18	0	596.37	596.37	2
3	QRM	0.161204824	0.0000	-223.44	2	450.89	449.30	1
	RUMM	0.547405102	0.0000	-1672.02	0	3344.04	3344.04	3
	LOT	0.381176692	0.0000	-348.86	0	697.71	697.71	2
4	QRM	0.168215655	0.0000	-252.81	2	509.62	508.03	1
	RUMM	0.406035458	0.0000	-1296.02	0	2592.05	2592.05	3
	LOT	0.24782902	0.0000	-304.60	0	609.21	609.21	2
5	QRM	0.13463468	0.0000	-249.79	2	503.58	501.99	1
	RUMM	0.440086431	0.0000	-1404.02	0	2808.04	2808.04	3
	LOT	0.2876622	0.0000	-331.31	0	662.63	662.63	2

According to AIC and BIC, QRM model dominates RUMM and LOT estimations in all partitions. The QRM model we employed assumes that the subjects accurately estimating the in-stock probability and does not consider the dynamics of learning or the error in estimated probabilities. These assumptions may have been significantly violated in **SIM2**. In the next section we explore in greater detail how subjects estimate in-stock probabilities. Despite these limitations we find that the QRM model has a strong explanation power. The experimental data does not support the conjectures that (a) subjects accurately compute the in-stock probability and (b) they accurately evaluate the differences in the utility of the two options.

Observation 4: The observed data supports the conjecture (C3) that subjects' decision making is consistent with the assumptions underlying quantal response models.

This implies that subjects have to learn through observation the in-stock probabilities and learn to evaluate the two options. In the next few sections we explore each aspect of the decision problem separately.

7. Decision Making with Information about In-stock probability

In **SIM-I** subjects were told the in-stock probability. If we observe any variation from threshold policies in **SIM-I** it would be due to subjects not consistently evaluating the difference between $\mathrm{U}(v - P_1)$ and $\mathrm{U}(v - P_2)$; lending greater support to the QRM model. Table 9 provides the number of subjects that employed threshold policies in **SIM-I**.

Table 9**:** Percentage of subjects employing a threshold policy

Treatment	Threshold policy	
	For all valuation (%)	Except for two valuations (%)
SIM1-I	2 (6.7%)	5 (16.7%)
SIM2-I	4 (13.3%)	8 (27.7%)
SIM3-I	5 (16.7%)	13 (43.3%)

The percentages are marginally higher than those in **SIM** but considerably lower than those in **LOT**. Leading us to conclude that even in **SIM-I** the buying behavior does not support RUMM with risk aversion. Table 10 below shows the out-of-sample fit of the QRM model to **SIM-I**.

Table 10: **SIM-I**: Out of sample goodness of fit test for QRM

Treatment	Partition No.	Rationality Parameter (γ)	Risk Aversion Parameter (r)	Log-Likelihood	Chi-square Test p value	Root Mean Square Deviation
	1	0.1494	0.0431	-211.16	0.0012	0.1483
SIM1-I	2	0.0909	0.0219	-192.31	0.1623	0.1173
	3	0.1246	0.0249	-208.16	0.9129	0.0608

	4	0.1191	0.0267	-202.31	0.8852	0.0729
	5	0.1143	0.0231	-200.96	0.8543	0.0805
SIM2-I	1	0.1093	0.0302	-230.37	0.0002	0.1187
	2	0.1468	0.0327	-237.04	0.2931	0.0869
	3	0.1588	0.0425	-236.99	0.0000	0.1701
	4	0.1141	0.0253	-237.04	0.1813	0.0900
	5	0.1158	0.0313	-230.86	0.0011	0.1207
SIM3-I	1	0.1206	0.0194	-209.44	0.6988	0.1037
	2	0.1067	0.0229	-177.43	0.9967	0.065
	3	0.1187	0.0297	-175.07	0.9931	0.0639
	4	0.1263	0.0530	-143.05	0.1140	0.1300
	5	0.1023	0.0296	-156.61	0.9723	0.0670

Similar to the **SIM** experiment, QRM models provide good estimates of **SIM1-I** and **SIM3-I**. Chi-square tests have p values well above 0.05 except for one sampling. RMSDs are also below 0.15 in all samplings. In **SIM2-I**, Chi-square test p values are greater than 0.01 in the two samplings out of five samplings and the maximum RMSD is 0.1701. For **SIM2- I**, we compared QRM to RUMM and LOT by using Akaike Information Criterion (AIC) and Bayesian Information Criterion (BIC) for model comparison. Once again according to AIC and BIC, QRM model dominates RUMM and LOT in all partitions. The results are available in the on-line version of the paper.

SIM-I data strongly suggests that subjects' evaluation of the two options includes a random error term, a key component of the QRM model. In Table 11, decisions in **SIM-I** appear "closer" to decisions in **SIM**, than **LOT**. Below we compare the three experiments.

Table 11: Percentage purchased in the first period: **SIM**, **SIM-I** and **LOT**

In-stock Probability	SIM	SIM-I	LOT
16%	79.81%	80.96%	89.74%
49%	60.53%	65.48%	57.76%
83%	43.91%	39.66%	14.35%

We computed for each subject the percentage of purchases in the first period in **SIM** (or **SIM-I**) and the percentage of times the same subject selected option B (the certain option) in **LOT**. We use Bonferroni method which is non-parametric method for multiple comparisons. This method requires one of the treatments to be the control. Because **SIM-I** lies between **SIM** and **LOT**, it was chosen as the control level.

When in-stock probability is 49% we cannot reject the hypothesis that the decisions in **SIM1**, **SIM2** and **LOT** are similar. However, when in-stock probabilities are 83% and 16%, the decisions in **LOT** are significantly different from **SIM** and **SIM-I**. Further **SIM** treatments are not significantly different from their corresponding **SIM-I** treatments. The confidence levels for the tests were 95%. We also employed other parametric method such as Sidák method and Dunnett method, and they too generated the same results as Bonferroni method. Therefore we can assert that there is a structural difference between the decisions in **SIM-I** and **LOT**. The randomness in estimating the utility of the options in **SIM**, surprisingly, dominates the impact of the errors in in-stock probability estimates in **SIM**. As a result decisions in **SIM** and **SIM-I** are similar.

8. Learning of in-stock probabilities and utilities

We use an exponential smoothing model (Brown and Meyer, 1961, Pötzelberger and Sögner, 2003) to analyze the evolution of subjects' beliefs of the in-stock probability. We assume that subject initial estimate of the in-stock probability is 0.5. (Tversky and Kahneman, 1973) This assumption is also supported by the similarities among LOT 49%, SIM1 and SIM1-I.

Let $result_t$ denote the outcome in trial t:

$$result_t = \begin{cases} 0, & if\ product\ goes\ out\ of\ stock\ in\ trial\ t \\ 1, & if\ product\ is\ in\ stock\ in\ trial\ t \end{cases}$$

The belief of in-stock probability at trial t, π_t, is updated as,

$$\pi_t = \alpha \cdot result_{t-1} + (1-\alpha)\pi_{t-1}$$

, where $0 \leq \alpha \leq 1$ is the weight attached to the most recent outcome.

In the **SIM** experiment, the outcome in trial t $(result_t)$ is the same for all subjects. Because in the exponential model the belief of the in-stock probability at trial t, π_t, only depends on the outcome, we assume that π_t is the same for all subjects. The valuation observed by subject i at t, $v_{i,t}$, however, differs by subject.

Subject i's probability of choosing to "buy in the first period" at t is,

$$Pr_{i,t}(\, v_{i,t}, p_1, p_2, \pi_t, \gamma) = \frac{exp(\gamma R(v_{i,t}, p_1, p_2, \pi_t))}{1 +\ exp(\gamma R(v_{i,t}, p_1, p_2, \pi_t)}$$

, where $v_{i,t}$, p_1 and p_2 are the valuation to the subject i at trial t, the first period price and the second period price, respectively, and γ is the rationality parameter.

The weight parameter, α, and the rationality parameter, γ, are estimated by maximizing the log-likelihood function:

$$max_{\gamma,\alpha} \sum_{t=1}^{T} \sum_{i=1}^{N} \log\big(I(Action(i,t),0) - Pr_{i,t}(\, v_{i,t}, p_1, p_2, \pi_t, \gamma)\big)$$

, where $(i,t) = \begin{cases} 0, & if \;\; subject\ i\ decides\ to\ \text{"}wait\ for\ the\ second\ period\text{"} \\ 1, & if\ subject\ i\ decides\ to\ \text{"}buy\ in\ the\ first\ period\text{"} \end{cases}$,

and $I(x,y) = \begin{cases} 1, & if\ x = y \\ 0, & if\ x \neq y \end{cases}$.

The estimated parameters are in the following Table 12, and the evolution of the belief of in-stock probability is depicted in Figure 5.

Table 12: Estimation of the parameters by Exponential Smoothing Model

Treatment	In-Stock Probability	Rationality (γ)	Weight parameter (α)
SIM2	83%	0.0809	0.0048
SIM3	16%	0.0871	0.0925

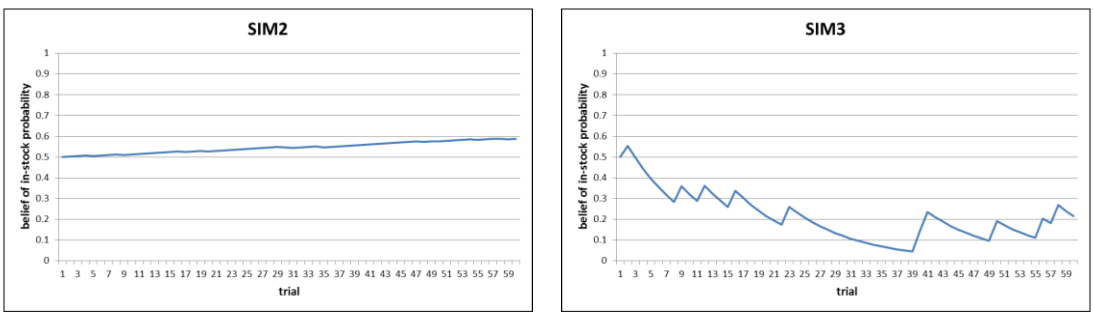

Figure 5: Belief of in-stock probability at trial, π_t.

The weight parameter α is smaller when the in-stock parameter is 83%, implying that in **SIM2** subjects do not react as strongly to the current outcome and take longer to migrate away from the initial belief of 0.5. In **SIM3** by the end of 30 trials the estimated in-stock probability is close to 0.16, the true in-stock probability. The estimate fluctuates between 0.0621 and 0.2305, with an average close to 0.16. On the other hand in **SIM2** at the end of 60 trials the estimated in-stock probability is close to 0.6. This lends support to our conjecture that losses loom larger than gains. When the in-stock probability is 16% if subjects wait they experience a stock out with 84% probability. This is perceived as a loss because they could have purchased in the first period. For a risk-neutral subject it is optimal to buy in the first period in **SIM3** if valuation is greater than 125. In **SIM2** it is optimal to always wait for a risk-neutral subject. Yet subjects do not react as aggressively to the potential lost revenue that they would have *gained* had they waited. They continue to buy in the first period, there by acting as if the in-stock probability is lower than 83%.

Observation 5: The experimental data is consistent with the conjecture (C4) that losses loom larger than gains. Subjects learn faster when the in-stock probability is lower.

The last set of treatments **CONV** focused on how subjects evaluated the utility of waiting relative to the utility of buying in the first period. In all three **CONV** treatments the in-stock probability was 49%, close to the initial estimate, to minimize the role of probability learning. In each treatment, which consisted of 20 trials, subjects had to repeatedly make a buy or wait decision for the same resale value. The resale values were different in each treatment. Here subjects did not know how many high or low type customers were present. They knew the number of units for sale, the number of buyers, and that other subjects' valuations ranged between 0 and 100. The data for the experiment are given below:

Table 13: Parameters for **CONV** experiment

Treatment	N	K	P_1	P_2	π_2	v	τ
CONV1	10	3	70	50	0.48	89	88.46
CONV2	10	4	60	30	0.49	80	88.82
CONV3	10	5	50	20	0.53	60	83.83

For a risk-neutral subject employing RUMM it is optimal to buy in the first period in CONV1 and to wait in CONV2 and CONV3. The threshold (τ) under RUMM is in Table 13. 86 subjects participated in the **CONV** Experiment. In Figure 6 we plot the percentage of times subjects choose to buy in the first period in **CONV1**, **CONV2**, and **CONV3**. If there is any learning that results in subjects converging to a risk neutral RUMM equilibrium then we should observe a trend.

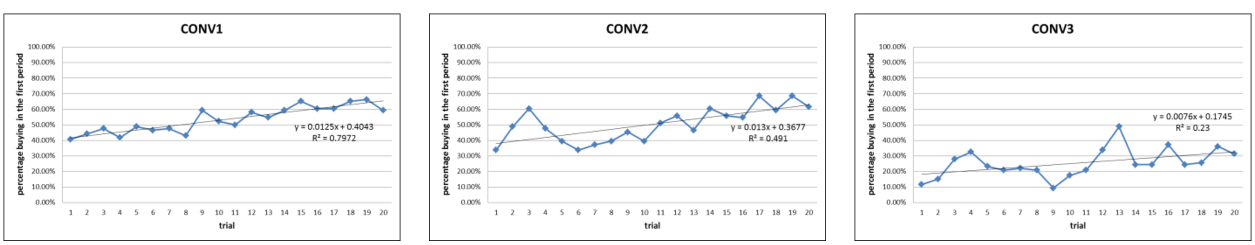

Figure 6: Percentage of buying now decisions at each trial in CONV

Table 14 below contains linear regression corresponding to these figures:

Table 14: Regression summary for **CONV**

Treatment	Slope	R-Sq	p
CONV1	0.0125	79.7%	0.000
CONV2	0.0130	49.1%	0.001
CONV3	0.0076	23.0%	0.032

The slopes are negligible in magnitude. Absence of a trend diminishes the possibility that all subjects will eventually behave like risk neutral rational utility maximizers. If subjects adopt RUMM with varying risk aversion levels then each subject should eventually converge to a decision. In that case, the percentage that buys in the first period should stabilize at some level and cease to oscillate. We counted the number of subjects that made the same decision in the last 5 and 7 periods.

Table 15: Percentage of subject that employed the same decision in the last few trials of **CONV**

	% who made the same decisions in the last 5 trials			% who made the same decisions in the last 7 trials		
	Buy in Period 1	Buy in Period 2	Total	Buy in Period 1	Buy in Period 2	Total
CONV1	23.26%	11.63%	34.88%	15.12%	6.98%	22.09%
CONV2	22.09%	8.14%	30.23%	16.28%	6.98%	23.26%
CONV3	4.65%	39.53%	44.19%	1.16%	37.21%	38.37%

Although in the aggregate data we failed to observe any significant trend, it appears that between 30 % and 40% of the subjects converged to a decision over 20 trials.

We employed the Experience Weighted Attraction model (EWA) (Camerer and Ho, 1999) to gain additional insights in to the learning process. The EWA computes for each subject i a parameter A_i^j that measures the attraction of strategy j for subject i. The likelihood of the subject employing strategy j is given by a logit model. In **CONV**, there are two strategies – buy in period j (j = 1 or 2). The probability of a subject adopting strategy j in trial t is given by:

$$\mathrm{P}_i^j(\mathrm{t}+1) = \frac{e^{\lambda \cdot A_i^j(t)}}{e^{\lambda \cdot A_i^1(t)} + e^{\lambda \cdot A_i^2(t)}}$$

, where λ is a measure of rationality.

The attraction weights for each subject are updates as follows:

$$A_i^j(t) = \frac{\phi \cdot N(t-1) \cdot A_i^j(t-1) + \left[\delta + (1-\delta) \cdot I\left(s^j, s_i(t)\right)\right] \cdot \pi(s^j)}{N(t)}.$$

$$N(t) = \rho \cdot N(t-1) + 1, t \geq 1.$$

The other parameters of the model are ρ $(0 \leq \rho \leq 1)$, ϕ $(0 \leq \phi \leq 1)$, and δ $(0 \leq \delta \leq 1)$. s^j denotes the j^th strategy, $s_i(t)$ is the strategy employed in trial *t*, *I(j,k) = 1* if *j = k* and 0 otherwise, and $\pi(s^j)$ is the payoff the subject would have realized in trial t if (s)he employed strategy s^j. $\pi(s^1)$ is always $(v - p_1)$ and $\pi(s^2) = (v - p_2)$ if the product is available and $\pi(s^2) = 0$ if the product goes out of stock. Readers are referred to Camerer and Ho (1999) for a detailed discussion of this model and the interpretation of the parameters.

We set $A_i^1(0)$ equal to $(v - p_1)$, $N(0)$ to 1, and constrained $A_i^2(0)$ to be between 0 and $(v - p_2)$. The parameters were determined by maximizing the log-likelihood function:

$$LL(A^2(0), \phi, \rho, \delta, \lambda) = \sum_{t=1}^{T} \sum_{i=1}^{N} \ln\left(\sum_{j=1}^{2} I\left(s_i^j, s_i(t)\right) \cdot P_i^j(t)\right)$$

We also computed the mean squared deviation:

$$MSD = \frac{\sum_{t=1}^{T} \sum_{i=1}^{N} \left[P_i^1(t) - I\left(s^1, s_i(t)\right)\right]^2}{N \cdot T}$$

, where T and N are the total number of trials and subjects, respectively.

The estimated parameters are given below in Table 16

Table 16: Estimated EWA Parameters for **CONV**.

Treatment	φ	ρ	δ	λ	$A^2(0)$	LL	MSD
CONV1	0.8695	0.5166	0.0000	0.0341	22.38	-445.38	0.2063
CONV2	0.8400	0.7733	0.1056	0.0661	29.97	-443.50	0.2039
CONV3	0.7921	0.6602	0.0005	0.0664	39.80	-359.61	0.1581

The following Figure 7 maps the relationship between the EWA model and the experimental results.

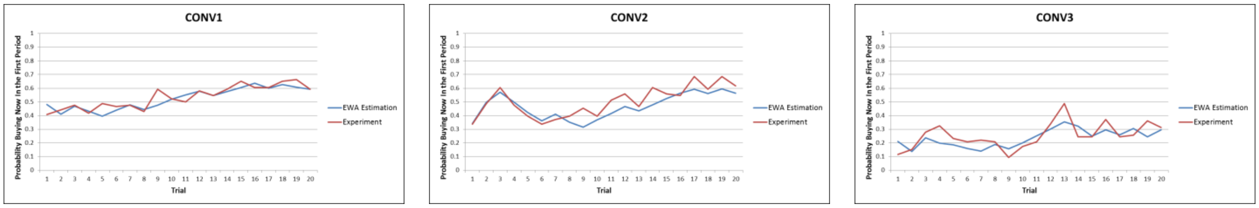

Figure 7: EWA Estimation

Value ϕ being close to 1 implies that subjects place a high value on past history and move gradually from the initial anchor. Values of δ being close to zero imply that subjects did not pay much attention to the strategy that they did not employ. A combination of high ϕ and low δ suggests that reinforcement model of learning is a better fit than the beliefs model. The rationality parameter λ is also close to zero. High values of λ imply that each subject converges rapidly to one of the two strategies. Low values of λ indicates that the percentage that buy in the first period is unlikely to stabilize, diminishing support for conjecture that subjects differ in risk aversion but employ RUMM.

Assuming that the EWA model depicts subjects' decision making with learning, we simulated decisions over 200 periods. In this simulation, in each period, the subject's decision is simulated based on the current attraction weight and the logit model, stock-out is generated, and the attraction weight is updated. We performed 1 million trials for each of the **CONV** treatments. Figure 8 graphs the histogram of the percentage of subjects buying in period 1 in the 200th period over the 1 million simulations. The X-axis is the percentage buying in the first period and the y-axis is the fraction of simulations in which this percentage was observed. Dispersion of the decisions lends further support to the assumptions underlying the QRM model. Even after 200 trials subjects do not converge to a decision but continue to randomize.

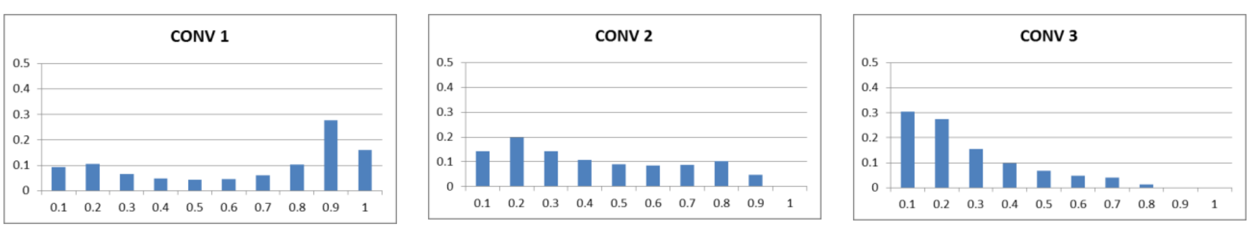

Figure 8: The percentage of subjects buying in period 1 in the 200th period

9. GENERAL DISCUSSION AND CONCLUSIONS

In this paper, we employed experiments to gain insights into how customers make purchase decisions for products that are dynamically priced. In particular we conducted experiments for a two-period posted pricing scheme. We investigated some of the common assumptions in theory and found that threshold policies, central to all rational utility maximizing models, are employed by very few subjects. Majority of the subjects randomize their decisions. As a result the aggregate buying behavior can be modeled accurately by quantal response models. Even when in-stock probability information is given, subjects continue to vacillate between the two decisions. When stock out probabilities and the pay-off are explicitly presented in the form of a survey, over 25% of the subjects did not employ a threshold policy. It is surprising that for this relatively simple decision problem, many subjects deviated from RUMM. We also found limited support for bargain hunting (Table 17).

Table 17: Percentage of subjects acting as bargain hunters or myopic customers

Treatment	Bargain Hunters		Myopic Customers	
	For all resale prices (%)	Except for two resale prices (%)	For all resale prices (%)	Except for two resale prices (%)
SIM1	0 (0.0%)	0 (0.0%)	0 (0.0%)	3 (10.0%)
SIM2	2 (6.9%)	3 (10.3%)	0 (0.0%)	0 (0.0%)
SIM3	1 (3.3%)	6 (20%)	0 (0.0%)	0 (0.0%)
SIM1-I	0 (0.0%)	0 (0.0%)	2 (6.7%)	3 (10.0%)
SIM2-I	1 (3.3%)	5 (16.7%)	1 (3.3%)	2 (6.7%)
SIM3-I	0 (0.0%)	0 (0.0%)	1 (3.3%)	10 (33.3%)

In **SIM2** and **SIM2-I**, for a risk neutral customer it is optimal to always wait for the second period act; i.e. act as a bargain hunter. As a result in these two treatments, bargain hunting is confounded with optimal decision making. In our experiment we draw attention to stock out risks which may cause subjects to occasionally purchase in the first period. By definition bargain hunters keep track of price reductions. This would also mean that they are very likely to be aware of stock outs. In that case our experiments imply that customers who are classified as bargain hunters because they always buy discounted products, are doing so not because it is a trait but because they have a lower willingness or ability to pay. Out of 179 subjects only 4 appeared to be bargain hunters.

In our study the percentage of buyers who acted myopically is also small. Our study, however, does not rule out the possibility that subjects outside the experimental setting may behave in a myopic manner.

The experiment explicitly draws subject's attention to the possibility of buying at a lower price. In absence of such framing, subjects may not have considered the option of waiting.

A surprising finding was the limited impact of probability information. We conducted two sets of experiments one in which stock-out information was provided and in another in which that information was provided. The decisions in these two experiments were surprisingly similar. This lends strong support to the assumption that subjects behave as if their assessment of the utility of each option includes a random error. Our **CONV** experiment suggests that this error term does not dissipate with repetition. We employed the EWA learning model to test whether subject's decisions would converge to a single strategy over time. We found that even after 200 trials the error term persists.

While subjects appear to have difficulty comparing the utility of the two options, they seem to learn in-stock probability levels quite rapidly. This can be inferred from the differences in the decisions made in the three treatments within **SIM**. Also exponential learning model shows that subjects discover the in-stock probability. Subjects learn faster when the in-stock probability is high because losses loom larger than gains.

9.1 Implications for the Firm and Existence of Equilibria

To optimize prices and stock levels dynamic pricing models need to compute expected revenues. Clearly RUMM with risk neutrality is a poor model. An improvement would be to base revenue estimates on risk aversion as measure by a conventional instrument such as the one we used in **LOT** (Holt and Laury, 2002). In our experiment we are guaranteed to sell all the units at the lower price P_2. The number of units for sale is less than the number of buyers. Incremental revenue from two prices is due to sales at the higher price P_1. From **LOT** data we computed the percentage that would buy in the first period and used that to estimate the error in the estimated incremental revenue (Table 18). Observe that the error can be as high as 67%. Ignoring bounded rationality can result in significant errors in revenue estimates, which in turn can lead to errors in optimal pricing and stocking decisions.

Table 18: Expected Additional Revenue at the High Price in **LOT** and **SIM** (% Error = (**LOT**-**SIM**)/**SIM**)

Experiment	**SIM1**	**SIM2**	**SIM3**
In-stock Probability	**49%**	**83%**	**16%**
SIM	242.11	175.65	319.24
LOT	231.03	57.41	358.97
% Error	-4.58%	-67.31%	12.44%

Table 18 also supports the conjecture that subjects' estimate of stock-outs are anchored around 50%. They under-estimate stock-out probability when it is below 50% and over-estimate stock-out probabilities when they are above 50%.

Our experiment employed a partial equilibrium model. Subjects were gaming against other fully rational customers simulated by the computer. In **SIM** we found that the simulated in-stock probabilities deviated from those that would have been observed if all subjects were quasi-rational.

Table 19: Comparison of calculated and actual in-stock probabilities in **SIM**

	SIM1	**SIM2**	**SIM3**
In-stock probability (π_2)	0.49	0.83	0.16
Re-calculated in-stock probability (γ)	0.46	0.80	0.16

In **SIM3**, in-stock probability based on the actual decisions of the subjects is identical to that which was simulated by the computer. Therefore, it is possible that the observed buying behavior is the same as the equilibrium behavior that would have been observed if subjects interacted with each other.

The recalculated in-stock probabilities in **SIM1** and **SIM2**, however, are different from the probabilities we employed in the treatment. To determine the equilibrium in-stock probability for these parameter settings, we ran an additional set of treatments by only changing the in-stock probability to 46% and 80%, respectively. We call these new treatments **SIM1-E** and **SIM2-E**, respectively. The parameters and results of the new treatments are summarized in Table 20.

Table 20: Parameters and results for experiment set **SIM1-E** and **SIM2-E**

Parameters	**SIM1-E**	**SIM2-E**
Number of High Type Customers(N_1)	10	10
Number of Low Type Customers(N_2)	26	14
Number of units for sale(K)	10	10
First period price(P_1)	120	120
Second period price(P_2)	80	80
In-stock probability in the second period(π)	0.46	0.80
Number of trials	60	60
Computed In-stock Probability	0.4640	0.7983

Surprisingly, in these treatments the in-stock probabilities calculated using the observed buying behavior are 46.40% and 79.83%. This result suggests the existence of equilibrium behavior when subjects interact with each other.

In summary we find strong support for QRM models in our experimental data. We find that ignoring bounded rationality can result in significant errors in revenue estimations. Next step would be to determine the equilibrium policies when all players in the exercise are subjects. It would also be useful to develop analytic pricing models in which customer decision making is based on quantal response models and thoroughly investigate the implication on the optimal policies for the firm.

References

Allon, G., A. Bassamboo, I. Gurvich. 2011. "We will be right with you": Managing customer expectations with vague promises and cheap talk. *Operations Research*, 59(6), 1382-1394

Anderson, C. K., J. G. Wilson. 2003. Wait or buy? The strategic consumer: Pricing and profit implications. *Journal of the Operational Research Society*, 54(3), 299-306

Archak, N., A. Ghose, P. G. Ipeirotis. 2011. Deriving the pricing power of product features by mining consumer reviews. *Management Science*, 57(8), 1485-1509

Aviv, Y., A. Pazgal. 2008. Optimal pricing of seasonal products in the presence of forward-looking consumers. *Manufacturing and Service Operations Management*, 10(3), 339-359

Bearden, J. N., R. O. Murphy, A. Rapoport. 2008. Decision biases in revenue management: Some behavioral evidence. *Manufacturing and Service Operations Management*, 10(4), 625-636

Bendoly, E. 2012. Real-time Feedback and Booking Behavior in the Hospitality Industry: Moderating the Balance between Imperfect Judgment and Imperfect Prescription. *Journal of Operations Management*, 31, 62-71

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

Brown, R. G., R. F. Meyer. The fundamental theorem of exponential smoothing. *Operations Research*, 9(5), 673-685.

Buhrmester, M., T. Kwang, S. D. Gosling. 2011. Amazon's Mechanical Turk A New Source of Inexpensive, Yet High-Quality, Data? *Perspectives on Psychological Science*, 6(1), 3-5

Cachon, G. P., R. Swinney. 2009. Purchasing, pricing, and quick response in the presence of strategic consumers. *Management Science*, 55(3), 497-511

Camerer, C., T.H. Ho. 1999. Experience-weighted attraction learning in normal form games. *Econometrica*, 67(4), 827-874.

Chen, K. Y., T. Hogg. 2009. Modeling Risky Economic Decision-Making with Bounded Rationality. In AAAI Spring Symposium: Human Behavior Modeling, 1-6

Chen, Y., X. Su, X. Zhao. 2012. Modeling Bounded Rationality in Capacity Allocation Games with the Quantal Response Equilibrium. *Management Science*. Articles in Advance, 1-11

Dasu, S., C. Tong. 2010. Dynamic pricing when consumers are strategic: Analysis of posted and contingent pricing schemes. *European Journal of Operational Research*, 204(3), 662-671

Elmaghraby, W., A. Gülcü, P. Keskinocak. 2008. Designing optimal preannounced markdowns in the presence of rational customers with multiunit demands. *Manufacturing and Service Operations Management*, 10(1), 126-148

Elmaghraby, W., P. Keskinocak. 2003. Dynamic pricing in the presence of inventory considerations: Research overview, current practices, and future directions. *Management Science*, 49(10), 1287-1309

Gallego, G., G. Van Ryzin. 1994. Optimal dynamic pricing of inventories with stochastic demand over finite horizons. *Management Science*, 40(8), 999-1020.

Güth, W., S. Kröger, H. T. Normann. 2004. Durable-Goods Monopoly with Privately Known Impatience: A Theoretical and Experimental Study. *Economic Inquiry*, 42(3), 413-424

Güth, W., P. Ockenfels, K. Ritzberger. 1995. On durable goods monopolies an experimental study of intrapersonal price competition and price discrimination over time. *Journal of Economic Psychology*, 16(2), 247-274

Ho, T. H., N. Lim, T. H. Cui. 2010. Reference dependence in multilocation newsvendor models: A structural analysis. *Management Science*, 56(11), 1891-1910

Ho, T. H., J. Zhang. 2008. Designing pricing contracts for boundedly rational customers: Does the framing of the fixed fee matter?. *Management Science*, 54(4), 686-700

Holt, C.A., S. K. Laury. 2002. Risk aversion and incentive effects. *The American economic review*, 92(5), 1644-1655

Kahn, B. E., J. Baron. 1995. An exploratory study of choice rules favored for high-stakes decisions. *Journal of Consumer Psychology*, 4(4), 305-328.

Katok, E. 2011. Laboratory experiments in operations management. TutORials in Operations Research. INFORMS, Hanover, MD, 15-35

Kaufmann, C., M. Weber, E. Haisley. 2013. The role of experience sampling and graphical displays on one's investment risk appetite. *Management Science*, 59(2), 323-340

Lim, N., T. H. Ho. 2007. Designing price contracts for boundedly rational customers: Does the number of blocks matter?. *Marketing Science*, 26(3), 312-326

Liu, Q., Van Ryzin. G. 2008. Strategic capacity rationing to induce early purchases. *Management Science*, 54(6), 1115-1131

McFadden, D. L. 1976. Quantal Choice Analaysis: A Survey. *Annals of Economic and Social Measurement*, 5(4), 363-390

McKelvey, R. D., T. R. Palfrey. 1995. Quantal response equilibria for normal form games. *Games and economic behavior*, 10(1), 6-38

Osadchiy, N., E. Bendoly. 2011. Are consumers really strategic? Implications from an experimental study. Working paper, Emory University, Atlanta, GA.

Paolacci, G., J. Chandler, P. Ipeirotis. 2010. Running experiments on amazon mechanical turk. *Judgment and Decision Making*, 5(5), 411-419

Pötzelberger, K., L. Sögner. 2003. Stochastic equilibrium: learning by exponential smoothing. *Journal of Economic Dynamics and Control*, 27(10), 1743-1770.

Schwarz, G. 1978. Estimating the dimension of a model. *The annals of statistics*, 6(2), 461-464.

Shen, Z. J. M., X. Su. 2007. Customer behavior modeling in revenue management and auctions: A review and new research opportunities. *Production and operations management*, 16(6), 713-728

Su, X. 2007. Intertemporal pricing with strategic customer behavior. *Management Science*, 53(5), 726-741

Talluri, K. T., G. Van Ryzin. 2005. The theory and practice of revenue management. Springer, New York.

Toubia, O., E. Johnson, T. Evgeniou, P. Delquie. 2013. Dynamic experiments for estimating preferences: An adaptive method of eliciting time and risk parameters. *Management Science*, 59(3), 613-640.

Tversky, A., D. Kahneman. 1973. Availability: A heuristic for judging frequency and probability. *Cognitive psychology*, 5(2), 207-232

Wu, D., K. Y. Chen. 2012. Supply chain contract design: Impact of bounded rationality and individual heterogeneity. *Production and Operations Management*, 2013

Yin, R., Y. Aviv, A. Pazgal, C. S. Tang. 2009. Optimal markdown pricing: Implications of inventory display formats in the presence of strategic customers. *Management Science*, 55(8), 1391-1408

Appendix A. SIM Instructions

Welcome to an experiment in decision making. You will be participating in a simulation. Your earnings will consist of participation compensation and an additional compensation that is based on your performance in this simulation. The performance based compensation can double your earnings.

Please read the instruction carefully. Your level of understanding can significantly affect your performance. After you read the instruction, you will be asked to take a test. If your score in this test is below the qualification level, you will be disqualified and you will only earn a basic compensation of $0.50.

THE SIMULATION GAME

You will play the role of a retailer who buys one unit of a product and then resells it to a customer.

You can purchase the product in one of two-periods. Price of the product in the first period is $120. In the second period, there is a discount, and the price is $80.In the first period, the product is always available. In the second period, however, you may not get the product because it may be sold out.

If you can purchase the product, you will resell it to a customer at a given resale price. And your profit = resale price - your purchase price

If you chose to wait and the product is sold out, then your profit is zero.

It is guaranteed that you can resell it, so you are only supposed to focus on when to buy the product.

MULTIPLE ROUNDS

The game consists of 50 to 70 rounds. In each round you will be given information that includes your resale price (From $125 to $200), the first period price ($120), and the second period price ($80). In each round, you have to decide if you want to (a) buy in the first period or (b) wait and purchase in the second period. After you make your decision, you will learn whether or not the product is available in the second period. You will also be informed of your profit. Then the simulation will advance to the next round.

Your objective is to maximize your profits. The additional performance based compensation you earn depends on your profit.

OTHER RETAILERS

In addition to you there are 35 other retailers. In the first period, only 10 retailers including you are present. In the second period, 26 additional retailers join the market. In each round, the retailers who are present in the first period have to decide whether to buy in the first period or wait until the second period.

The resale prices, however, are not the same for all of these retailers. For each retailer, a computer program randomly generates the resale price. The resale prices lie between $125 and $200 in $5 increments. As a result, different retailers will have different resale prices. The 26 additional retailers who join the market in the second period have resale prices between $81 and $120; therefore, they only want to buy the product in the second period.

The purchase prices (first period price and second period price), however, will be the same for all the retailers. Total number of retailers is 36 but the total number of units available for sale is 20. This is why you may not get the product in the second period. Because there are only 10 retailers in the first period, if you decide to buy the product in the first period, you are guaranteed to get the product.

In the second period all retailers who didn't buy in the first period join the 26 other retailers. All these retailers try to buy the product in the second period. Because the number of retailers exceeds the number of units for sale in the second period, the computer randomly selects the retailers who will get the product.

Total number of units for sale (20), total number of retailers present in the first period (10), additional number of retailers joining the market in the second period (26), your resale price, first period price ($120) and second period price ($80) will be displayed, and you don't have to remember them.

You will not know the resale price of the other retailers. Each retailer knows only his or her own resale price. Remember you will always get the product if you decide to buy in the first period. The computer also makes the purchasing decisions for all the other retailers. For the 9 other retailers who are present in the first period, the computer makes a decision based on a rule that depends only on the resale price observed by that retailer and the other information displayed. Using this rule, the computer decides independently for each of the 9 retailers whether they should buy in the first period or wait till the second period. The purchase decisions of these 9 other retailers will not depend on your decision. The 26 additional retailers who join in the second period always want to buy the product in the second period.

As stated earlier the game consists of 50 to 70 trials. In each trial your resale price and that of all the other retailers will change. The prices at which you can purchase, the number of retailers, and the number of units for sale will not change from one trial to the next.

Example 1:

1st Period Price	2nd Period Price	Your resale price	Number of units for sale	Number of retailers present in the first period	Number of additional retailers joining in the second period
$120	$80	$145	20	10	26

If you buy in the first period, then your profit is: $145 - $120 = $25.

If you wait and the product is available in the second period, then your profit is: $145 -$80 = $65.
If you wait and the product is not available in the second period, then your profit is: $0
If 4 of the 10 retailers decide to buy in the first period, then the number of units available in the second period: 20 – 4 = 16. And the number of retailers in the second period will be: 6 (from the first period) + 26 = 31. Only 16 out of these 31 retailers will get the product in the second period. The computer randomly picks the 16 out of the remaining 31 and assigns them the product.

PERFORMANCE COMPENSATION

Performance compensation can be $2 maximum. We will randomly pick 10 of your decisions. So out of 50 ~70 decisions, only 10 decisions will affect your performance. Because you don't know which decision will be selected, you should make each decision carefully. The profit you make in these 10 decisions will be summed up and the additional payment you receive will be proportional to this sum.

Example 2:

1st Period Price	2nd Period Price	Your resale price	Number of units for sale	Number of retailers present in the first period	Number of additional retailers joining in the second period
$120	$80	$125	20	10	26

If you buy in the first period your profit will be: Resale price = $125, Purchase price = $120, Profit = $125 - $120 = $5.
If you wait for the second period:
And if the product is out of stock, Profit = $0
And if the product is in stock: Resale price $125, Purchase price = $80 and Profit = $125 - $80 = $45.
If 9 retailers decide to buy in the first period,
Number of units left for the second period = 20 – 9 =11
Number of retailers who decide to wait for the second period = 10 - 9 = 1
Number of additional retailers joining in the second period = 26
Total Number of retailers in the second period = 1 + 26 = 27

Only 11 out of these 27 retailers will get the product in the second period.

Appendix B. LOT Instructions

In this experiment, you are given an opportunity to choose between two different options and receive additional payments based on your decisions.

One of the options (Option A) offers you a risky choice with a higher amount of points while the other option (Option B) offers a lower amount of points but there is no uncertainty. Two options will be displayed; one on the right and one on the left.

Example 1

Option A: 10% chance at winning 200 points OR Option B: 10 points.

In example 1, you can either choose option A with a 10% chance at 200 points, OR you can choose Option B that guarantees 10 points. If you select Option A, then a computer will generate a random number and determine whether or not you receive the 200 points. The chance of winning will be 10%. If you don't win you will get 0 points. If you choose the guaranteed option, you will be given the 10 points.

There are 5 decision sets. Each decision set has 20 decision questions. Please answer all of the decision questions in the 5 decision sets.

A computer simulation will randomly select one of the four decision sets and then randomly choose one of the decision questions from the chosen decision set. If you selected Option A then the computer will perform another simulation to determine whether or not you get the corresponding points. The likelihood of you winning the points will depend on the probability of success for Option A. On the other hand if you chose Option B, you will be given the corresponding points.

The points achieved will be translated to monetary value between $0 to $1.

Participation fee for this experiment is $0.50 and additional payments based on the outcome in the decision you made can result in your total compensation to go up to $1.

Appendix C. Sample Screenshots of SIM experiments

Figure C1. Decision Making

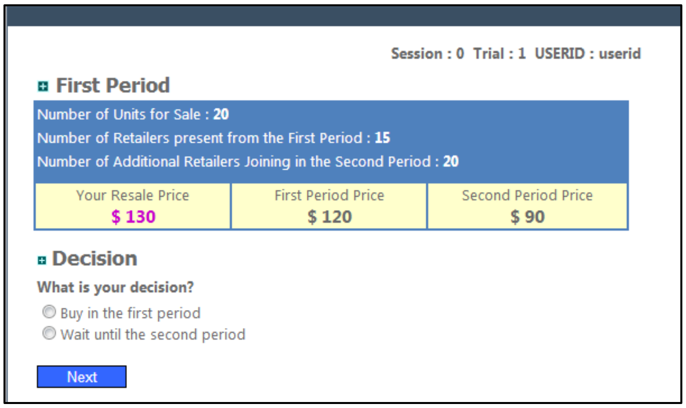

Figure C2. Result – In Stock

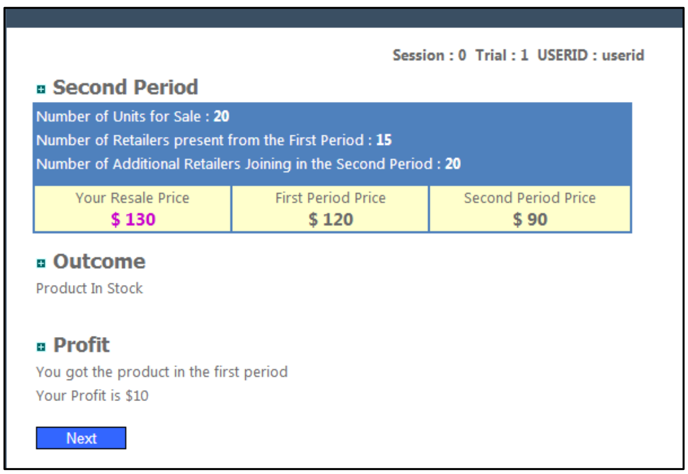

Figure C3. Result – Out of Stock

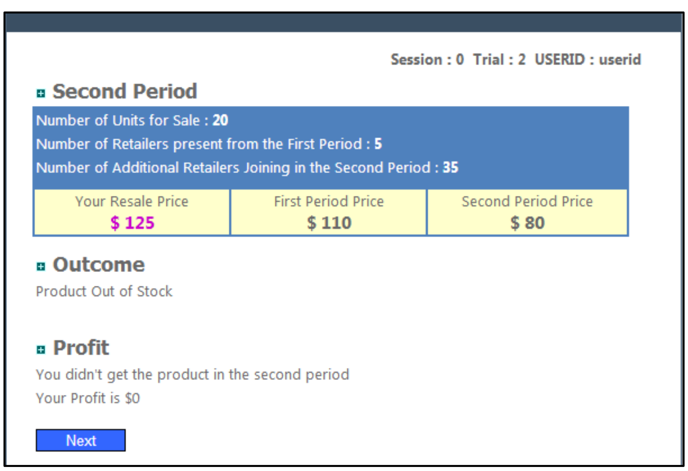

Appendix D. Sample LOT Questions

	Option A	Option B	Your Choice A or B
Decision 1	49% chance of winning $45	$5	
Decision 2	49% chance of winning $50	$10	
Decision 3	49% chance of winning $55	$15	
Decision 4	49% chance of winning $60	$20	
Decision 5	49% chance of winning $65	$25	
Decision 6	49% chance of winning $70	$30	
Decision 7	49% chance of winning $75	$35	
Decision 8	49% chance of winning $80	$40	
Decision 9	49% chance of winning $85	$45	
Decision 10	49% chance of winning $90	$50	
Decision 11	49% chance of winning $95	$55	
Decision 12	49% chance of winning $100	$60	
Decision 13	49% chance of winning $105	$65	
Decision 14	49% chance of winning $110	$70	
Decision 15	49% chance of winning $115	$75	
Decision 16	49% chance of winning $120	$80	