Exploding Offers with Experimental Consumer Goods*

Alexander L. Brown
Department of Economics
Texas A&M University

Ajalavat Viriyavipart
Department of Economics
Texas A&M University

Xiaoyuan Wang
Department of Economics
Texas A&M University

August 29, 2014

Abstract

Recent theoretical research indicates that search deterrence strategies are generally optimal for sellers in consumer goods markets. Yet search deterrence is not always employed in such markets. To understand this incongruity, we develop an experimental market where profit-maximizing strategy dictates sellers should exercise one form of search deterrence, exploding offers. We find that buyers over-reject exploding offers relative to optimal. Sellers underutilize exploding offers relative to optimal play, even conditional on buyer over-rejection. This tendency dissipates when sellers make offers to computerized buyers, suggesting their persistent behavior with human buyers may be due to a preference rather than a miscalculation.

Keywords: exploding offer, search deterrence, experimental economics, quantal response equilibrium
JEL: C91 D21 L10 M31

*We thank the Texas A&M Humanities and Social Science Enhancement of Research Capacity Program for providing generous financial support for our research. We have benefited from comments by Gary Charness, Catherine Eckel, Silvana Krasteva and seminar participants of the 2013 North American Economic Science Association and Southern Economic Association Meetings. Daniel Stephenson and J. Forrest Williams provided invaluable help in conducting the experimental sessions.

1 Introduction

One of the first applications of microeconomic theory is to the consumer goods market: basic supply and demand models can often explain transactions in centralized markets quite well. When markets are decentralized, often additional modeling is required to explain the interactions of consumers and producers. One popular model, a search model, suggests consumers sample prices from a variety of producers, buying once the price of goods falls below a certain threshold (Stigler, 1961; Rothschild, 1974). If decisions of sellers can affect buyer search, the model becomes more complicated. Armstrong and Zhou (2013) show under relatively mild conditions, it is unilaterally profitable for sellers to deter search. Specifically the strategies of exploding offers (i.e., "take-it-or-leave-it" offers), "buy-now" discounts, and requiring deposits for the option to buy later are profitable for sellers. Given the profitability of such strategies, a natural question to ask is why they are not seen more often market transactions. One possibility is that producers are justifiably concerned that consumers may respond more negatively to these tactics than theory predicts. The focus of this paper will be an experimental investigation of this question: whether producers are hesitant to use exploding offers and whether consumers respond more negatively to such offers than theory predicts. We choose exploding offers over the other two search-deterrence tactics (i.e.,"buy-now" discounts, deposits) simply because we believe this tactic is most likely of the three to generate a negative reaction from buyers.

Prior to Armstrong and Zhou (2013), the focus of most economic research on exploding offers concerned labor markets, a type of market where there are specific cases of exploding offers being the norm.[1] There may be distinct features of labor markets—not found in consumer goods markets—that are responsible for the prevalence of exploding offers. Exploding offers may be an essential tool in the unraveling of matching markets, as employers compete to lock down new employees earlier and earlier (Niederle and Roth, 2009). Additionally, if employers are required to have only one outstanding offer to a candidate at a time, exploding offers can be seen as a technique to "forestall the event that no one is hired" (Lippman and Mamer, 2012). Evidence of negative responses to exploding offers also relies on features unique to labor markets. Lau et al. (forthcoming) find experimental evidence that employees will respond negatively to employers' use of exploding offers by reducing effort after being hired; there is no equivalent reaction that could occur in consumer goods markets.

[1]For instance, law students applying for appellate court clerkships frequently receive exploding offers (Roth and Xing, 1994; Avery et al., 2001, 2007; Niederle and Roth, 2009).

Thus it is an open question how individuals will react to exploding offers in consumer goods markets, a setting where the use of such offers is notably different from labor markets. Observing behavior of buyers and sellers in situations with exploding offers would be difficult. Firms do not have incentives to record or publicly release their use of exploding offers. Moreover, markets may offer more durable products or long term services, and receive certain amounts of new demand at each period, which makes information less transparent and traceable at the individual consumer level.

For these reasons we turn to experimental analysis, the first such analysis of exploding offers in the consumer goods setting. We implement a simplified version of Armstrong and Zhou (2013): two sellers simultaneously choose from one of three prices and either make exploding or non-exploding offer. Buyers, previously unaware of their personal value for either seller's good, randomly visit one seller and learn their value for that seller's good. In doing so, they receive the seller's offer. The buyer must then decide whether to visit the other seller. If the first seller makes an exploding offer, a visit to the second seller will terminate the opportunity to buy from the first seller. If our intuition and previous experimental results are any guide, buyers will over-reject exploding offers. They will choose to visit the second seller more often than theory would predict. Sellers may also make decisions inconsistent with theory.

We can observe the behavior of buyers directly. Sellers' behavior, however, is conditional on perceived buyer response. To isolate this effect, we use two treatments. In one, sellers knowingly interact with computer buyers programmed to follow optimal strategy; in the other they interact with human buyers.

Our results are striking. Consistent with our intuition, buyers reject exploding offers roughly 20% more often than optimal theory dictates. The tendency persists through all twenty periods of the experiment. Sellers make exploding offers less often to human buyers than computer buyers. Adjusting payoffs for sellers to account for the increased propensity of buyers to reject exploding offers, we find that sellers still are hesitant to give human buyers exploding offers, a tendency we term "exploding-offer aversion." The net result of these differences from optimal strategy is a transfer of surplus from sellers to buyers. That is, seller payoffs in the computer-buyer treatment are higher than in the human-buyer treatment, and human buyers earn more than computer buyers.

These results can help explain why exploding offers in consumer markets may not be as prevalent as theoretically predicted. Buyers will likely reject them more than the optimal level.

Sellers make fewer exploding offers than optimal even after accounting for this buyer behavior. While one should be cautious in making field predictions directly from the results of laboratory data, it is important to note that many of the features that might make sellers hesitant to use exploding offers and buyers quick to reject them are not present in this laboratory environment. Sellers make offers anonymously and are somewhat insulated from the negative feelings of having their offer rejected. Buyers may find exploding offers less objectionable on a computer interface than through an actual human seller. So it may be reasonable to conclude that if anything the tendencies found in this experiment may be amplified in actual field market situations.

The remainder of our paper is organized as follows: Section 2 provides the theoretical model used in our experiment. Section 3 discusses our experimental design. Section 4 presents the results. Section 5 closes the paper with a brief discussion of related work and concluding remarks.

2 The Model

The experiment in this paper implements a simplified model based on Armstrong and Zhou (2013). The following section describes the basic setup of the model and the assumptions and simplifications used in the experiment. Section 2.1 explains how the sequential search game takes place. Sections 2.2 and 2.3 explain the optimization problem for the buyer and seller, respectively. Section 2.4 describes simplifying assumptions and parameter choices that will be used in the experiment. The main changes from the literature are to discretize buyer valuations and seller pricing. This change reduces the number of decisions for subjects, simplifying the problem. Assuming optimal play by buyers, the end result is a 6 x 6 symmetric normal-form game between two sellers. Table 1 provides payoffs for a seller given a fixed offer and pricing choice strategy, conditional on the other seller's pricing and offer choice strategy. The table will be used as a theoretical benchmark for analysis of sellers choices in the experimental game.

2.1 The Search

This model represents an experimental search market of two sellers with one buyer who visits each seller sequentially in a random order.[2] Each seller offers a good which has a private value for the buyer drawn from the same ex-ante value distribution: $V_k^i \in \left\{V_1^i, V_2^i, ..., V_K^i\right\}$

[2]Several identical buyers were used in our experiments for a larger sample size. Each seller can only choose one strategy for all buyers in each period.

Table 1: Expected Payoffs for One Seller's Strategy Choice Given Other Seller's Strategy Choice

	E25	E30	E35	F25	F30	F35
E25	11.91	13.28	13.28	**15.04**	15.82	16.21
E30	11.72	13.59	15	13.36	**17.58**	18.52
E35	**13.67**	**13.67**	**15.86**	14.49	15.59	**20.51**
F25	9.18	12.7	13.48	12.3	15.23	16.41
F30	9.14	10.08	13.83	10.78	14.06	17.34
F35	10.12	10.66	11.76	10.94	12.58	16.41

Note: For the strategy labels letters "F" and "E" denote free-recall and exploding offers, respectively. The number indicates price. For example, "E25" indicates the strategy of offering price 25 with an exploding offer. This convention is used throughout the paper.

(where $i = 1, 2$ represents sellers and $k = 1, 2, ..., K$ represents K possible values) with probability $\lambda_1 \equiv prob(V_1), \lambda_2 \equiv prob(V_2), ..., \lambda_K \equiv p(V_K)$. The game is as follows:

1. Each seller sets a price from a possible price range: $P^i \in \left\{P^i_1, P^i_2, ..., P^i_L\right\}$ and chooses an offer type as either an exploding or a free-recall offer.

2. Nature randomly selects which seller the buyer will visit first (S^1).[3]

3. The buyer observes the prices of both sellers (P^1 and P^2) and his value of the first good he[4] visits (V^1).

4. The buyer chooses whether to accept the first offer or to visit S^2. If he chooses to accept, the transaction occurs and the game is ended; otherwise, the game continues to the next step.

5. The buyer visits S^2 and observes the value of the good (V^2).

6. The buyer chooses whether to accept or reject the offer from S^2. If he accepts, the transaction occurs and the game is ended. If he rejects and the first offer was an exploding offer, no transaction occurs and the game is ended. If he rejects and the first offer was a free-recall offer, the game continues to the next step.

7. The buyer chooses whether to accept or reject the offer from S^1 (if it is a free-recall offer).

Each player's payoff is determined after the game is ended. If there is no transaction, all players receive zero payoff. If there is a transaction, the buyer receives a payoff equal to the difference between his value and the price of the good he bought; that seller receives a payoff equals to that price; the (other) seller with no transaction receives zero payoff.

[3]We denote the first seller S^1 and the other seller S^2.

[4]As a convention, we assume female sellers and a male buyer.

2.2 Buyer Best Response

We assume that the buyer is rational and has an objective to maximize his expected payoff. Since the offer type of the second seller has no effect on a strategy of the buyer, we only need to consider two cases; (1) the first offer is a free-recall offer and (2) the first offer is an exploding offer.

If the first offer is a free-recall offer, visiting S^2 does not prevent the buyer from revisiting S^1, the buyer always searches.[5] After visiting both sellers, the buyer chooses an option that provides him the highest payoff from three possible options. The options are (1) accepting the first offer ($V^1 - P^1$), (2) accepting the second offer ($V^2 - P^2$), and (3) rejecting both offers (zero payoff).

If the first offer is an exploding offer, the buyer would make a decision by comparing the payoff from accepting the first offer and the expected payoff from rejecting the offer. The payoff from accepting the first offer is the difference between the value and the price of the first offer or $\Pi^1 = V^1 - P^1$ while the expected payoff from visiting S_2 is

$$E(\Pi^2) = \sum_{k=1}^{K} \lambda_k{}^* \max(0, V_k^2 - P^2).^6 \tag{1}$$

The buyer accepts the first offer if $\Pi^1 < E(\Pi^2)$ and rejects otherwise.[7] If the first offer was rejected, the buyer accepts the second offer as long as $V^2 > P^2$.

2.3 Seller Strategies

Similar to the buyer, we assume that each seller is rational and has an objective to maximize her expected payoff. In this market, each seller is required to choose a price and an offer type before knowing which seller each buyer visits first. There are three possible cases to be considered: (1) both sellers use exploding offers, (2) both sellers use free-recall offers, and (3) one seller uses an exploding offer and another seller uses a free-recall offer.

First, consider a case where both sellers use exploding offers. Consider seller i with a price P^i,

[5] In some cases, it is not necessary for the buyer to search. For example, if V^1 is the highest possible value from the distribution and $P^1 \leq P^2$. In which case, there is no gain or loss from searching, so we assume for simplicity that the buyer always visits the second seller if the first offer was a free-recall offer. Different assumptions do not change the equilibrium of the game in our experiments.

[6] If a value of the good from the second seller is higher than the price, the buyer would accept the offer and gain $V_k^2 - P^2$; however, if $V_k^2 < P^2$, he would reject the offer and earn zero payoff. So, for each value k of the second good, the buyer would earn the greater of 0 and $V_k^2 - P^2$. The expected payoff is calculated from the sum of the multiplication of $\max(0, V_k^2 - P^2)$ and its probability as shown above.

[7] If $\Pi^1 = E(\Pi^2)$, we assume that the buyer would search with probability $\frac{1}{2}$. Different tie-breaking rules do not change the equilibrium of the game.

who plays with seller j with a price P^j. There are two possible situations with an equal probability:

1. A buyer visits seller i first. The buyer will accept the offer if the difference between his valuation of the first good and its price is greater than the expected payoff from the second offer; i.e., $V_k^i - P^i > E(\Pi^j) = \sum_{l=1}^{K} \lambda_l{}^*\max(0, V_l^j - P^j)$ and rejects otherwise. The probability that he will accept the offer is

$$\text{Prob(accept } i_1) = \sum_{k=1}^{K} \lambda_k{}^* D_k^i \tag{2}$$

   where $D_k^i = 1$ if $V_k^i - P^i > E(\Pi^j)$ and $= 0$ otherwise.

2. A buyer visits seller j first. Similar to the first case, the buyer will accept the offer from j with probability $\sum_{l=1}^{K} \lambda_l{}^* D_l^j$ where $D_l^j = 1$ if $V_l^j - P^j > E(\Pi^i) = \sum_{k=1}^{K} \lambda_k{}^*\max(0, V_k^i - P^i)$ and $= 0$ otherwise. If the buyer rejects the offer from seller j, he will visit seller i. Upon visiting i, he will accept the offer as long as his value is above P^i or with probability $\sum_{k=1}^{K} \lambda_k{}^* B_k^i$ where $B_k^i = 1$ if $V_k^i > P^i$ and $= 0$ otherwise. So, the probability that the buyer will purchase from seller i is

$$\text{Prob(accept } i_2) = (1 - \sum_{l=1}^{K} \lambda_l{}^* D_l^j)^* \sum_{k=1}^{K} \lambda_k{}^* B_k^i. \tag{3}$$

   Therefore, seller i's expected payoff is $P^{i*}[\frac{1}{2}\text{Prob(accept } i_1) + \frac{1}{2}\text{Prob(accept } i_2)]$.

Second, consider the case where both sellers use free-recall offers. Again, consider seller i with price P^i who plays with seller j with price P^j. The order of seller visits has no effect here because a buyer always searches in this scenario. Therefore, the buyer will purchase from seller i if (1) $V_k^i - P^i > V_l^j - P^j$ and (2) $V_k^i - P^i > 0$. The probability that the buyer will purchase from seller i is

$$\text{Prob(accept } i_3) = \sum_{k=1}^{K} \sum_{l=1}^{K} \lambda_k \lambda_l{}^* A_{kl}^{ij}, \tag{4}$$

where $A_{kl}^{ij} = 1$ if (1) $V_k^i - P^i > V_l^j - P^j$ and (2) $V_k^i - P^i > 0$ and $A_{kl}^{ij} = 0$ otherwise. Therefore, his expected payoff is $P^{i*}\text{Prob(accept } i_3)$.

Last, consider a case where one seller uses an exploding offer and another seller uses a free-recall offer. Since an offer type of the second seller has no effect on the buyer' strategy, we can use the expected payoffs from the previous two cases. If seller i uses an exploding offer while seller j uses a free-recall offer, seller i's expected payoff is $P^{i*}[\frac{1}{2}\text{Prob(accept } i_1)+\frac{1}{2}\text{Prob(accept } i_3)]$.[8] If seller i uses a free-recall offer while seller j uses an exploding offer, seller i's expected payoff is $P^{i*}[\frac{1}{2}\text{Prob(accept } i_3)+\frac{1}{2}\text{Prob(accept } i_2)]$.[9]

2.4 Parameter Choice for an Experimental Search Market

The previous analysis shows how payoffs are calculated in this game. For any sets of values $V_k^i \in \{V_1^i, V_2^i, ..., V_K^i\}$, probability $\lambda_1, ..., \lambda_K$, and prices $P^i \in \{P_1^i, P_2^i, ..., P_L^i\}$, we can calculate payoffs for any combinations of strategies for each seller. Because we are interested in a case where using an exploding offer is an optimal strategy, we choose parameters for our experimental market as follows:

$V \in \{10, 25, 40, 55, 65, 70\}$

$\lambda(10) = \lambda(25) = \lambda(40) = \lambda(55) = 0.125$ while $\lambda(65) = \lambda(70) = 0.25$

$P \in \{25, 30, 35\}$

In this case, there exists a unique equilibrium in which both sellers in the market choose an exploding offer with the highest price of 35 (points). In this equilibrium, a buyer would accept the first offer only if his value for the first item is either 65 or 70 and reject all other values. If the first offer was rejected, the second offer would be accepted as long as his value for the second item is above 35 (40, 55, 65, 70). All other combinations of choices cannot be established as a Nash Equilibrium. We provide expected payoffs for all decisions in Table 1.

3 Experimental Design

The experiment consists of two treatments. In the computer-buyer treatment (CB), human sellers were matched against computer buyers programmed to play an optimal strategy. In the human-buyer treatment (HB), human sellers were matched against human buyers. Each session consisted of eight sellers (for either treatment) and sixteen buyers. In each period, four markets were randomly formed from each pair of sellers. Each market consisted of two sellers and four buyers who made six independent decisions each. Twenty periods were played in all sessions and the role of each participant was fixed for the entire session. In the every market, half of buyers visited

[8]The case where the buyer visits i first is equivalent to the case where both sellers use exploding offers and the case where the buyer visits j first is equivalent to the case where both sellers use free-recall offers.

[9]The case where the buyer visits i first is equivalent to the case where both sellers use free-recall offers and the case where the buyer visits j first is equivalent to the case where both sellers use exploding offers.

one seller first and the other half visited the other seller first.

A sequence of values of items were randomly generated for every period and were identical across sessions and treatments. In addition, the same random matching was used in every session and treatment.[10] Using human rather than computer buyers was the only difference across treatments.

The instructions were both shown on screen and read aloud to ensure the game was common knowledge among the participants. After the instructions, the participants answered a quiz, in multiple choice form, to establish that they understood how to play the game. Each participant needed to answer all questions correctly before the game started.

Each seller in both treatments was paid based on one randomly selected period.[11] Seller earnings were determined by the price chosen in that period multiplied by the quantity sold and the conversion rate was four cents for one point. Each buyer in the HB treatment was paid based on one random decision in one random period. The earnings were calculated from the difference between the value and the price of that particular item purchased or zero if no purchase was made. The conversion rate for a buyer was a dollar for two points.[12]

At the end of each session, participants were privately paid their earnings in the session (plus a five dollar show-up bonus) in cash. For an 80 minute session, participants earned $18, on average.

The experiment was conducted in the Economic Research Laboratory at Texas A&M University, in April and October 2013. Four sessions (32 sellers) of the CB treatment and three sessions (24 sellers, 48 buyers) of the HB treatment were conducted. All 104 participants were Texas A&M University undergraduate students recruited campus wide using ORSEE (Greiner, 2004). The experiment was programmed and conducted with the software z-tree (Fischbacher, 2007).

4 Results

Result 1 *Sellers play different strategies against computer and human buyers. Sellers offer lower prices and chose to use exploding offers less often against human buyers. Both tendencies persist, if not intensify, over the course of the experiment.*

[10] If in one session, participant i was matched with participant j in period n; in all other sessions, participant i would be matched with participant j in period n as well.

[11] We choose to pay for one random decision to eliminate any subject complimentarities that might occur across decisions or periods, most notably income effects. See Azrieli et al. (2014) for a greater discussion.

[12] After the decision making portion of the session was completed, participants filled out a questionnaire consisting of demographics information, a Risk Preference test (Eckel and Grossman, 2008), and a Cognitive Reflection test (Frederick, 2005). We did not find any correlation between questionnaire responses and subject decision making.

Table 2: Summary Table of Sellers' Decisions

Buyer type	Observations	Exploding Offers	Free Recall Offers	30E	35E	25F	30F	35F	Average Price[1]
	640	435	71	157	207	110	75	20	30.36
computer	100.00%	67.96%	11.09%	24.53%	32.34%	17.19%	11.72%	3.13%	(0.16)
	480	262	165	75	22	164	40	14	26.95
human	100.00%	54.58%	34.38%	15.62%	4.58%	34.17%	8.33%	2.92%	(0.14)

[1] Standard error is reported in this column (in parentheses) rather than percent of observations.
Note: For the data labels letters "F" and "E" denote free-recall and exploding offers, respectively. The number indicates price. For example, "E25" indicates the strategy of offering price 25 with an exploding offer. This convention is used throughout the paper.

We first compare sellers' decisions in the computer-buyer treatment (CB) with those in the human-buyer treatment (HB). Table 2 provides a summary of all seller decisions across both treatments. Over all periods, sellers used exploding offers more often (67.96% in CB vs. 54.58% in HB) and offered lower prices (30.36 on average in CB vs. 26.95 on average in HB). Pooling these values at the subject level and comparing across treatment, a rank sum test suggests that these values are significantly different ($p < 0.035$ and $p < 0.001$, respectively).

Table 2 also shows the frequency that each combination of strategy and price was used over 20 periods. The modal response (used in 32% of all observations) in the CB treatment was the equilibrium strategy, an exploding offer with a price of 35 (points). This strategy was used in less than 5% of all observations in the HB treatment. Thus, sellers play the theoretically predicted, equilibrium strategy when it is optimal against computers programmed to play as theory would suggest. They do not play such strategy against human buyers, who we will see later are not playing the theoretically optimal strategy. The modal response in the HB treatment was an exploding offer with a price of 25 (used in 34% of all observations), a strategy only used in about 11% of all observations in the CB treatment. We will show later (in Result 3) the importance of this specific strategy in the HB treatment.

We can also observe the dynamics of subject decisions. Over the twenty periods in the experiment, both the sellers in the CB and HB treatments increased their use of exploding offers (Figure 1). The percentage of sellers who used an exploding offer in the CB treatment is higher than in the HB treatment in most periods. In the last 5 periods, about 76% of sellers in CB treatment used an exploding offer, while only about 65% of sellers in HB treatment used an exploding offer. A joint test of the period dummy variables suggests that the average exploding offer usage was significantly different between two treatments ($p \approx 0.010$). Yet a linear trend test reveals that the increase rates for both treatments were similar ($p \approx 0.476$).

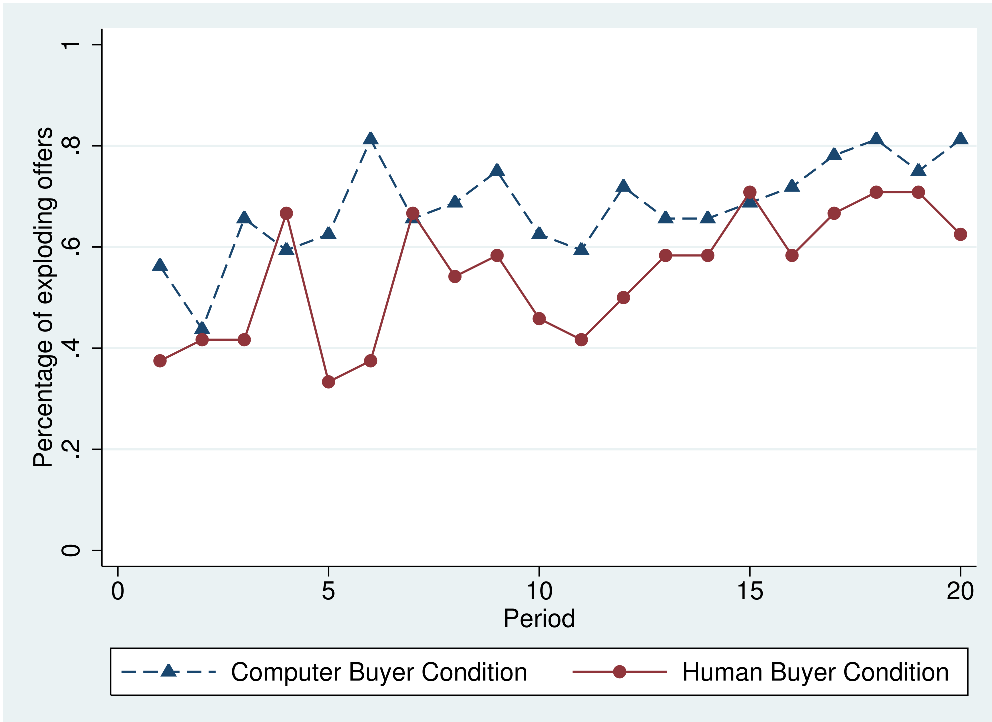

Figure 1: Proportion of Exploding Offers Used by Sellers by Period, HB and CB Treatments

Selling price dynamics are quite different as seen in Figure 2. In the first 2 periods, average prices across treatment are nearly identical. After that, they diverge. While seller prices in the CB treatment remain the same (if not increase), in the HB treatment, they quickly drop (p-values for linear trend coefficients are 0.176 and 0.000, respectively). In the last 5 periods, seller prices were on average 30.49 in the CB treatment and 26.65 in the HB treatment; the price difference is significant ($p < 0.001$).

Result 2 *When given an exploding offer, buyers reject the offer (search for the second seller's item) more often than profit-maximizing play dictates. This tendency holds over all prices and valuations; it persists throughout the experiment.*

Buyers made 6 purchase attempts in each period over 20 periods. Pooling the results from 3 sessions of 16 buyers each we have a total of 5,760 ($6 \times 20 \times 16 \times 3$) purchase attempts. Table 3 provides summary data on all of these choices.[13] In 3,144 of these purchase attempts buyers encountered an exploding offer on the first item they searched. Optimal play (based on the price of the items and buyer valuation of the first item) dictates that buyers should have accepted this

[13]Due to a computer glitch 9 buying attempts were unable to be recorded. These affected four different buyers over two periods in one session. Given the small number of observations lost compared to the total number in the sample, we cannot envision how this loss of data would affect any results.

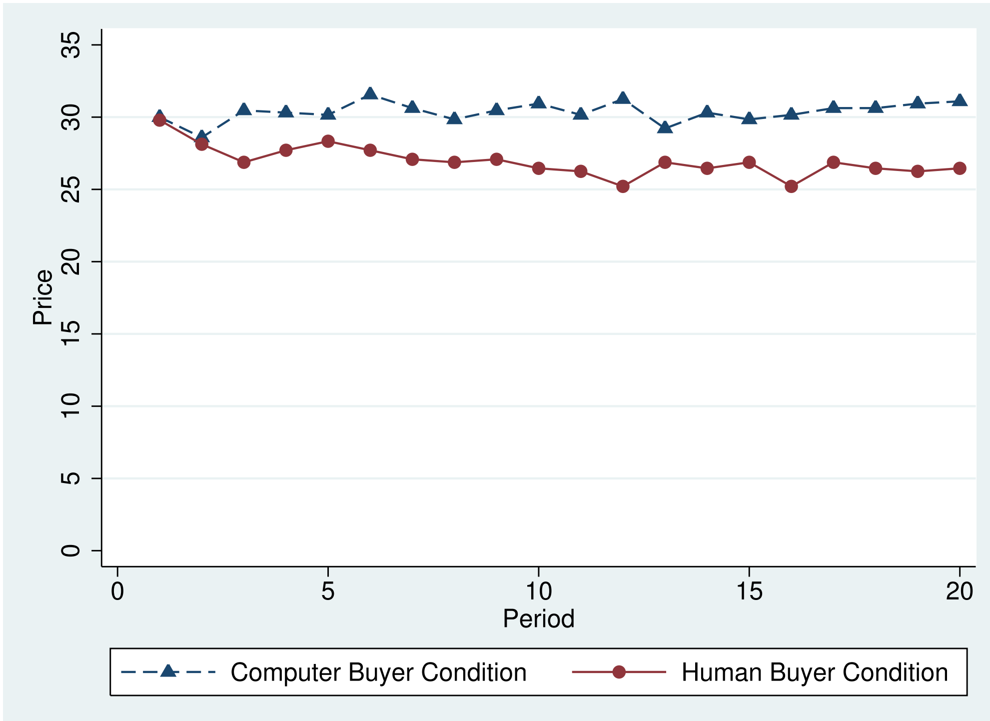

Figure 2: Average Price Offered by Sellers by Period, HB and CB Treatments

Table 3: Summary Table of Buyers' Decisions

	1st offer is exploding		1st offer is free-recall		Overall	
	Actual	Optimal play	Actual	Optimal play	Actual	Optimal play
Accepts 1st offer, immediately	1,618 51.46%	1,861 59.11%	539 20.68%	689 26.43%	2,157 37.51%	2,550 44.34%
Searches for 2nd offer	1,526 48.54%	1,283 40.81%	2,068 79.32%	1,918[2] 73.57%	3,594 62.49%	3,202 55.66%
Accepts 2nd offer	1,131 35.97%	921(+129)[1] 29.29%(+4.1%)	1,085 41.62%	1,032(+237) 39.59%(+9.09%)	2,216 38.53%	1,953(+366) 33.96%(+6.36%)
Recalls 1st offer	-	-	794 30.46%	521(+237) 19.98%(+9.09%)	794 13.81%	521(+237) 9.06%(+4.12%)
Accepts neither offer	395 12.56%	233(+129) 7.41%(+4.1%)	189 7.25%	94(+88) 3.61%(3.38%)	676 11.75%	327(+217) 5.69%(3.77%)
Total offers	3,144	3,144	2,607	2,607	5,751	5,751

[1] Numbers in parenthesis represent indifference treatments for optimal play. The subjects can receive the same net value or the subjects may receive a best offer with 0 net value. Therefore, we provide a conservative measure and its upper bound.

[2] We assume that consumers search only when the current value is strictly smaller than the difference between the highest value 70 and the other seller's price. Therefore, this measure is also a lower bound. There are 498 indifference cases.

first offer in 1,861 (59.11%) instances; instead buyers accepted in only 1,618 instances (51.46%), a difference that is statistically significant ($p < 0.001$). The net result is that buyers accept the second offer, the only offer that remains, far more often than the optimal strategy dictates. Buyers accepted the second offer 1,131 (35.97%) times after an exploding offer, higher than the 921–1,050 (29.29–33.39%) times[14] they would if they followed optimal strategy. The calculated (net) loss of earnings for such deviation (when the initial seller uses an exploding offer) is about 0.8 points per item (2.38% of earnings), a value that is significantly different from 0 ($p < 0.001$).

It should be noted that buyers also displayed a tendency to search for a second offer more often than "optimal" with free-recall offers, though these cases are very different from exploding offers. In general, buyers with a free-recall offer should continue to search for the second offer unless they will receive a surplus from the first seller than cannot be beaten by the second seller (e.g., receiving the highest possible value on an item offered at the lowest possible price). In those cases, it is unnecessary for buyers to search—the first offer is optimal—but searching produces no economic loss as buyers may recall their first offer. Buyers with free-recall offers ultimately chose the right item—the one with the highest net gain—86.74% of the time.[15]

The tendency for buyers to turn down exploding offers more often than optimal play is not isolated to a specific valuation or seller price pair. Figure 3 illustrates optimal response (dashed line) and actual response (solid line) in terms of rejection rate for buyers over different valuations for the first item when the seller uses an exploding offer. For instance when a buyer has a value higher than 55, in most instances optimal play would be to accept the offer. In the experiment, however, buyers show a significant amount of rejection under these high values. Separating the data by seller-price pairs (i.e., the price the first seller makes with an exploding offer and the price the second seller offers), the over-rejection patterns remain under all price pairs (Figure 4).

Buyers persistently over-reject exploding offers over the course of the experiment. Figure 5 plots the rejection rate from optimal play and buyer rejection rate. In every period, the actual rejection rate is greater than or equal to the rate predicted by optimal play. Both a parametric t-test and non-parametric rank sum test, collapsed to the subject level, suggest the rejection rate with human buyers is higher than optimal ($p < 0.001$).

Result 3 *Because of their propensity to reject exploding offers, human buyers present different incentives*

[14]This number varies depending on whether optimal buyers would have bought the second item if the net gain from doing so was zero (when value=price).

[15]In the remainder of these choices buyers mistakenly chose the item they valued most, ignoring price, rather than focusing on net gain.

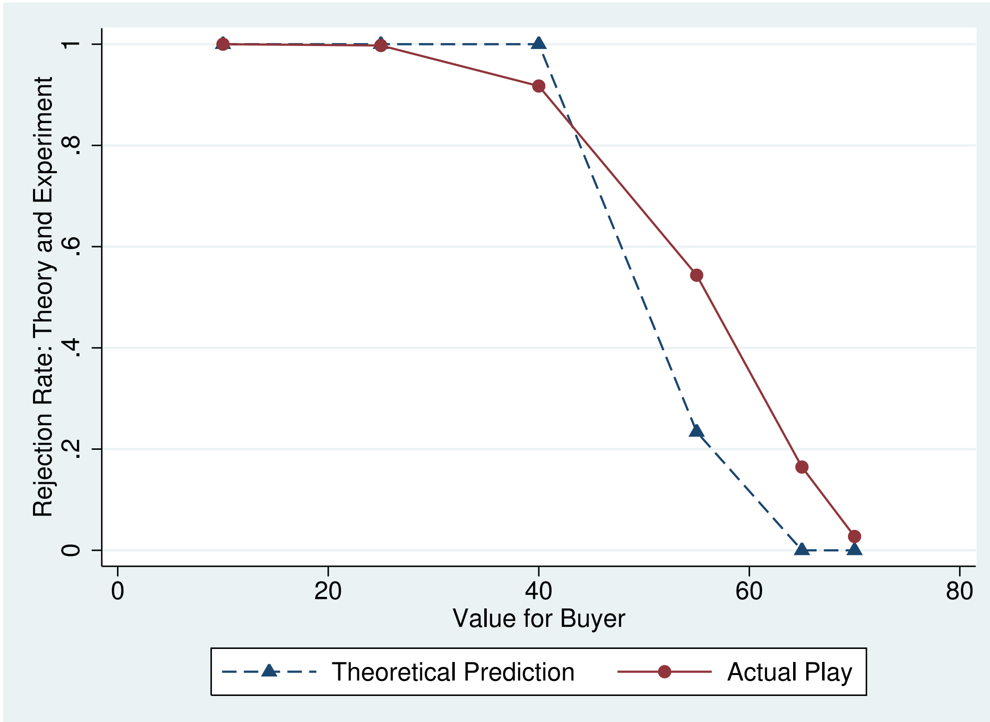

Figure 3: Rate of Buyer Rejection of Exploding Offer, Theoretical Prediction and Actual Play

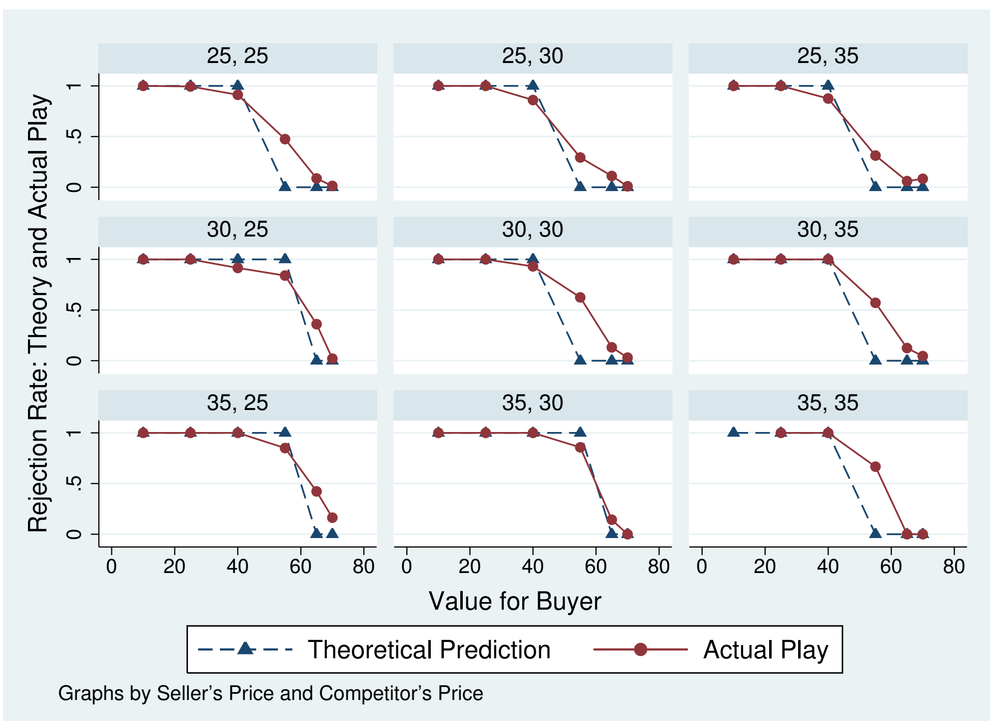

Figure 4: Aggregate Rejection Rate, Theoretical Prediction vs. Actual Play, Separated by Seller Pricing Pairs (First Seller Price, Second Seller Price)

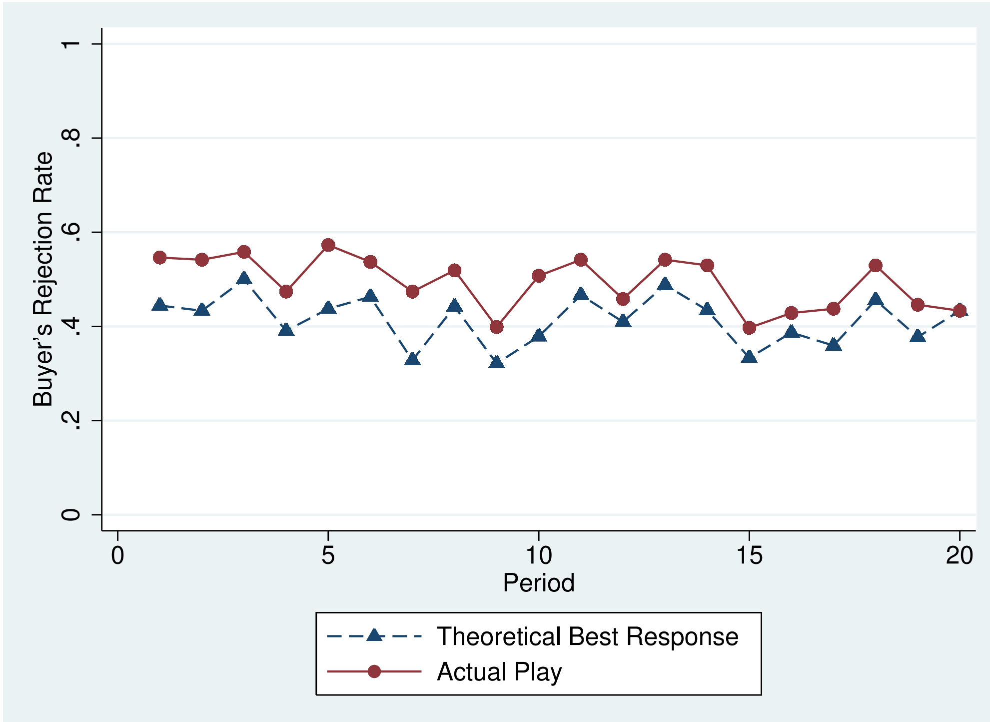

Figure 5: Rate of Buyer Rejection of Exploding Offer by Period, Theoretical Prediction and Actual Play

to sellers than computers following optimal strategy. In addition to the standard equilibrium found when buyers play the theoretically optimal strategy, the sellers' pricing game with the payoffs created by human buyers contains a second equilibrium where sellers both make exploding offers at the lowest price. The quantal response equilibrium model shows this second equilibrium is the convergent equilibrium.

Result 2 demonstrates that human buyer behavior is significantly different from optimal behavior. As one might expect, this presents different expected payoffs for sellers in the HB than theory predicts. Table 4 presents a comparison of payoffs depending on whether the two sellers in the game face human or computer buyers. The human buyer payoffs are constructed using the buyer choice distributions in our sample. They capture the fact that an average buyer will over-reject exploding offers. Payoff values are determined by a simulation where 20,000 "human" buyers receive the offers of two sellers in random order. The payoffs for strategies that involve exploding offers are generally lower with human buyers than the theoretical prediction. This difference creates a second, pure-strategy, symmetric equilibrium where both sellers play the lowest price as an exploding offer $((25, E), (25, E))$ in addition to the pure-strategy, symmetric equilibrium of both sellers playing the highest price with an exploding offer $((35, E), (35, E))$. The latter strategy

Table 4: Payoff Matrices for Seller Strategies, Given Theoretically Optimal Play and Empirical Play of Human Buyers

	25, E	30, E	35, E	25, F	30, F	35, F
Theoretical						
25, E	11.909	13.271	13.271	13.674	15.059	15.641
30, E	11.721	13.563	14.949	12.887	16.409	18.071
35, E	13.675	13.675	**15.824**	14.200	15.035	19.143
25, F	10.016	12.740	13.309	11.781	14.528	15.679
30, F	9.650	11.291	14.312	10.816	14.137	17.434
35, F	10.486	11.258	13.173	11.012	12.618	16.493
Simulated Human Buyers[1]						
	25, E	30, E	35, E	25, F	30, F	35, F
25, E	**11.730**	13.694	14.901	12.571	14.647	15.595
30, E	10.880	13.105	14.401	11.759	14.475	16.601
35, E	10.812	13.566	**15.513**	11.510	13.624	17.839
25, F	10.940	13.575	14.985	11.781	14.528	15.679
30, F	9.937	12.767	15.234	10.816	14.137	17.434
35, F	10.313	12.561	14.167	11.012	12.618	16.493

[1] The ”human buyer” payoff matrix is calculated like the theoretically optimal matrix, except that the observed rejection rate of exploding offers is used rather than the theoretical optimum.

pair is the only pure-strategy, symmetric equilibrium that exists in theory or against computer buyers who play the theoretically optimal strategy.

There are two points about the game with simulated human buyers that require more explanation. First, it is surprising that both pure-strategy equilibria feature sellers making exploding offers. Our results this far have stressed that human buyers reject these offers more often than their computer counterparts, seemingly making such strategy less profitable. Making exploding offers, nonetheless, is still the most profitable of the two seller strategies. Note that if sellers pick equal prices with different types of offers, the seller with the exploding offer earns higher expected profits. Lowering prices against an exploding offer can lead to higher payoffs in some cases. Sellers may find it effective to offer an exploding offer with a lower price to induce a reluctant buyer to accept an exploding offer.

Second, the existence of two pure-strategy, symmetric equilibria brings up the issue of equilibrium selection. It is desirable to be able to focus on one equilibrium and there are many techniques to do so. We use one such technique, the quantal response equilibrium model (QRE) (McKelvey and Palfrey, 1995, 1998). The technique will ultimately show that the lower priced equilibrium $((25, E), (25, E))$ is the “convergent equilibrium.”

The quantal response equilibrium model (QRE) (McKelvey and Palfrey, 1995, 1998) has become a standard way to model game theoretic data where equilibrium strategies are not always played. For this reason, it is quite popular in analyzing experimental data. In a quantal response equilibrium, each player has correct beliefs about which strategies every other player will play. Expected payoffs for each strategy choice can be calculated given those beliefs.

$$\pi_i(s_i) = E\left[u_i(s_i, s_{-i})|s_{-i}\right] \tag{5}$$

Then a given player noisily best responds according to those expected payoffs. Specifically, they choose strategies with higher expected payoffs more often. As Camerer (2003) notes, typically the QRE uses a logit payoff response function

$$P(s_i) = \frac{\exp\left(\lambda\pi_i(s_i)\right)}{\sum_{s_k}\exp\left(\lambda\pi_i(s_k)\right))}. \tag{6}$$

Since each player calculates $P(s_i)$ which is dependent on all other players' values for $P(s_i)$, the system is recursive. For a given λ, the solution to the system of equations (6) is a quantal response equilibrium. The parameter λ measures each player's sensitivity in responding to these payoffs. At $\lambda = 0$ players play each strategy with equal probability. As λ increases each player best responds better, playing higher payoff yielding strategies more frequently, until at $\lambda = \infty$ each player is playing a purely best response and we are at a Nash equilibrium, "the convergent equilibrium."

To model the quantal response equilibrium in our experiment, we use an equivalent formalization of the QRE model. Rather than having players make noisy choices, we have them maximize noisy utility functions. The expected utility for a given strategy choice is subject to a random element. Given consistent beliefs and the strategy set S_i, each player i solves the problem

$$\max_{s_k \in S_i} \pi_i(s_k) + \epsilon_{ik}. \tag{7}$$

The ϵ's are independently and identically distributed type-1 extreme value, making the problem and equilibrium equivalent to the system of equations shown in (6).

Figure 6 illustrates the prevalence of different seller strategies under the quantal response equilibrium model in the HB and CB treatment. The payoffs given in table 4 are directly used

Figure 6: Seller Strategies in Quantal Response Equilibrium by λ with Payoffs Determined by Either Theoretical-Optimum or Empirically-Observed, Human-Buyer Play

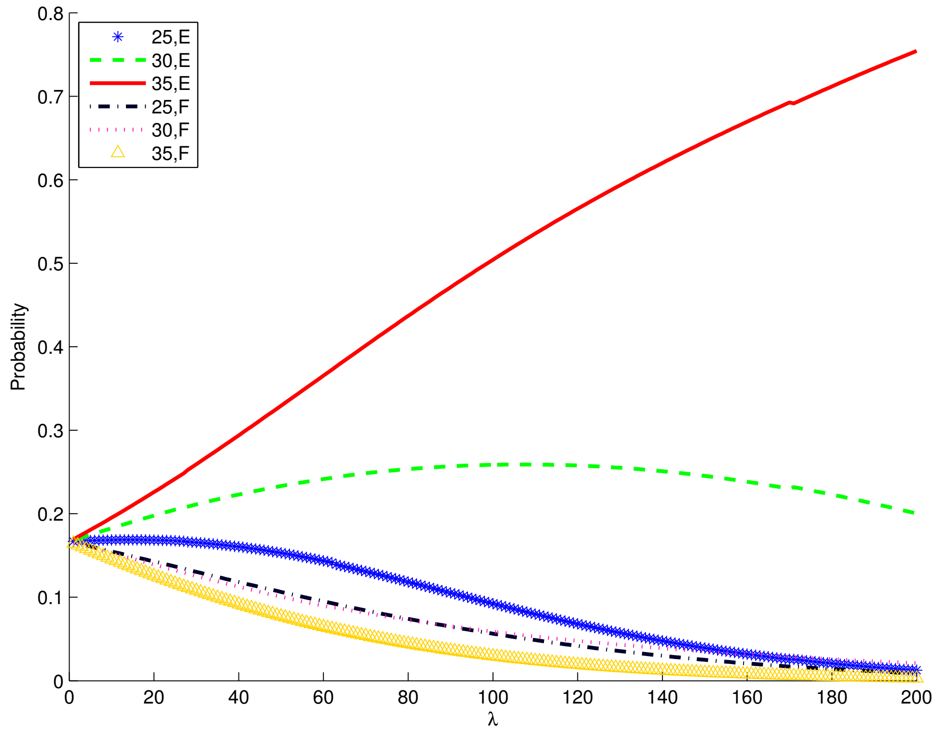

(a) Theoretical Optimum Buyer

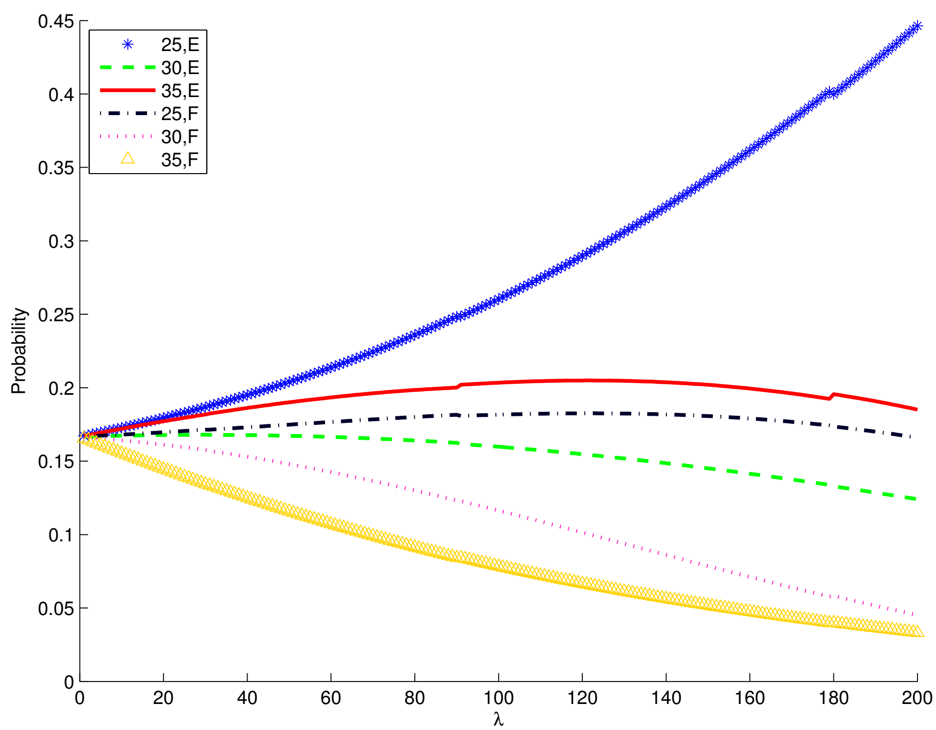

(b) Empirically Observed Buyer

as sellers' payoffs for playing different strategies.[16] In the CB treatment, exploding offers lead to consistently better payoffs than free-recall offers given the computer buyers' optimized play. In the HB treatment, however, buyers consistently over-reject exploding offers. Modifying seller payoffs to account for this over-rejection creates a new game, one where both sellers making an exploding offer with price 25 becomes an additional Nash Equilibrium. Figure 6(b) shows that the QRE model selects this equilibrium as the convergent equilibrium.

It is interesting to note that this convergent equilibrium strategy is also the modal strategy choice of sellers in the HB treatment (see Result 1).

Result 4 *In both treatments, sellers demonstrate a reluctance to play strategies that involve the use of exploding offers. The tendency persists through all twenty periods against human buyers, but dissipates against computer buyers who play optimal strategy. This analysis controls for the differential expected payoffs of both strategies in the human- and computer- buyer treatments.*

The quantal response model provides a baseline utility framework for sellers (see equation 7). In order to determine whether sellers have any preferences toward exploding offers not captured in the model, we introduce a new term δ that is included in sellers' utility only if they make an exploding offer. If δ is negative (positive), than sellers are reluctant (overeager) to use exploding offers; they derive additional negative (positive) utility from making them. If δ is zero, sellers do not have a systematic bias in their use of exploding offers. As the use of exploding offers varies between treatments and also within treatment by period (see Figure 1), we introduce four terms to capture the dynamics and session effects of exploding offers. The terms δ_{H0}, δ_{HT} represent the δ term in the first and last periods of the human buyer treatment, respectively; the terms δ_{C0}, and δ_{CT} represent the δ term in the first and last periods of the computer-buyer treatment, respectively. All other periods are convex combinations of their respective treatments' two terms. Similar terms are constructed for λ in the QRE model: λ_{H0}, λ_{HT}, λ_{C0}, and λ_{CT}. Equation (8) provides this utility model for subject i in period t.

$$u_{it}\left(s_{it}\right) = \left(\frac{20-p}{19}u_{X0}\left(s_{it}\right) + \frac{p-1}{19}u_{XT}\left(s_{it}\right)\right) \tag{8}$$

$X \in \{C, H\}$ represents two treatments.

where

[16] To be clear, this means we are analyzing a two-player, normal-form game. We set buyer behavior as fixed. That is, we are not using QRE to model buyer behavior.

$$u_{C0}(s_{it}) = \lambda_{C0}\left(\sum_{-i} \hat{u}(s_{it}, s_{-i})\, \pi_{-i}(s_{-i}) + \delta_{CO}\mathcal{I}(\text{exploding offer})\right)$$
$$u_{CT}(s_{it}) = \lambda_{CT}\left(\sum_{-i} \hat{u}(s_{it}, s_{-i})\, \pi_{-i}(s_{-i}) + \delta_{CT}\mathcal{I}(\text{exploding offer})\right)$$
$$u_{H0}(s_{it}) = \lambda_{H0}\left(\sum_{-i} \hat{u}(s_{it}, s_{-i})\, \pi_{-i}(s_{-i}) + \delta_{H0}\mathcal{I}(\text{exploding offer})\right)$$
$$u_{HT}(s_{it}) = \lambda_{HT}\left(\sum_{-i} \hat{u}(s_{it}, s_{-i})\, \pi_{-i}(s_{-i}) + \delta_{HT}\mathcal{I}(\text{exploding offer})\right)$$

Table 5 provides parameter estimates for this model. Initially, in both the HB and CB treatments sellers were reluctant to use exploding offers. Both coefficients, δ_{CO} and δ_{HO}, are significantly less than 0 ($p < 0.001$). By period 20, however, sellers' reluctance to use exploding offers on human buyers persists (δ_{HT} is significantly less than 0, $p < 0.001$), but sellers show no reluctance to use exploding offers on computer buyers (δ_{CT} is only marginally significant, $p \approx 0.054$). The terms of this "exploding-offer aversion" are economically significant. Literally interpreting the coefficients suggests that sellers experienced a disutility equivalent to \$1.10–\$1.40 in possible earnings the use of exploding offers against human buyers.[17] Full analysis of both buyer and seller earning are found in the next result.

The λ term in the HB treatment is generally greater than the corresponding term in the CB treatment. An F-test rejects the joint hypothesis of both $\lambda_{CO} = \lambda_{HO}$ and $\lambda_{CT} = \lambda_{HT}$ ($p < 0.001$). Further, the estimate of λ increases in the HB treatment over the 20 periods (λ_{HT} is significantly greater than λ_{HO}, $p < 0.001$), but if anything the estimate decreases in the CB treatment. The λ is usually interpreted as the "noisiness" parameter in a QRE model. Thus, our estimation results indicate that sellers play more "accurately" with human buyers. This may be due to two facts. First, we used empirical play information in the estimation. The empirical play is determined by seller-buyer interactions. Second, sellers may face less payoff uncertainty when playing with human buyers, given that they reject high price offers or exploding offers with higher probability.

Result 5 *Sellers in the human-buyer treatment earn less than sellers in the computer-buyer treatment. Human buyers are better off compared with computer buyers. The aggregate surplus of the human-buyer treatment is lower than that of the computer-buyer treatment.*

[17] This is a tricky point. Sellers are only paid based on one randomly-selected period of twenty so no one decision to avoid an exploding offer has an expected cost of \$1.10–\$1.40. However, the pattern of behavior of *continually* avoiding exploding offers does cost sellers losses of this magnitude.

Table 5: Parameter Estimates for Dynamic QRE Model of Exploding Offer Usage with Both Human and Computer Buyers

	Computer Buyers	Human Buyers
λ_{X0}	0.964^{***}	1.427^{***}
	(0.149)	(0.205)
λ_{XT}	0.748^{***}	3.272^{***}
	(0.139)	(0.212)
δ_{X0}	-2.017^{***}	-1.106^{***}
	(0.202)	(0.232)
δ_{XT}	-0.595^{*}	-1.442^{***}
	(0.308)	(0.269)
LL	-1036.710	-710.265

[1] $X \in \{C, H\}$ represents computer-buyer treatment or human-buyer treatment.

Result 2 shows the primary difference between human buyers and the optimal play, utilized by computer buyers, is the human buyers' increased tendency to reject exploding offers. This difference in play leads to great differences in earnings. Sellers earn more on average each period with computer buyers ($12.94) than human buyers ($11.47). Both a parametric t-test and non-parametric rank sum test, collapsed to subject level, confirm sellers earn more in the CB treatment ($p < 0.001$ for both tests). Human buyers' earnings are significantly greater than computer buyers ($16.429 vs. $15.224, $p < 0.001$ for both tests). These results cannot be due to different realizations of buyer valuations; both computer buyers and human buyers received exactly the same draws of a random distribution of valuations.

Figures 7(a) and (b)—which show the average earnings of sellers and buyers, respectively, in each treatment, over the twenty periods of the experiment—demonstrate these differences in payoffs persist. There is evidence to suggest the difference between seller earnings in HB and CB treatments is actually *increasing* over the course of the experiment. Between the two treatments, the average difference in seller earnings is $1.172 in periods 1–10; the average difference in seller earnings is $1.771 in periods 11–20. A difference-in-difference regression reveals this result is significant ($0.60, $p < 0.001$). Similarly, buyers' profit difference on average is $1.205 and it is increasing over time ($p = 0.003$). There is also a downward trend in the average profit for computer buyers ($p = 0.029$).

The aggregate market surplus for both buyers and sellers is higher for the CB treatment. A t-test ($p = 0.075$) and a rank-sum test ($p = 0.073$) at group level show the difference is marginal, or statistical significant at 90% confidence interval. In monetary value, the average difference is

Figure 7: Profit for Buyers and Sellers in Human- and Computer- Buyer Treatments. (a, left) Profit for Sellers. (b, right) Profit for Buyers.

\$0.27, or 0.95% of the average surplus each period.

5 Conclusion

This paper provides the first experimental investigation into methods of search deterrence in the consumer goods market. While theory (Armstrong and Zhou, 2013) suggests some form of search deterrence is optimal for sellers in nearly all conditions, our suspicion was that buyers might respond negatively to such tactics, reducing the likelihood they would be used by sellers. Our findings confirm this suspicion. Buyers reject exploding offers more often than is optimal. Sellers use exploding offers less often than both the theoretical optimum and a profit-maximizing strategy based on actual buyer behavior would dictate. Sellers do not demonstrate a similar tendency with computer buyers, suggesting their aversion to exploding offers may be a preference-based phenomenon and not the result of miscalculation.

The results of this experiment provide suggestive evidence as to why search deterrence in consumer markets may not be as widespread as theory might suggest. Buyers simply do not like exploding offers and reject them more than what profit-maximizing behavior would dictate. A best-responding seller would have to take this into account and use exploding offers less often than the theoretical optimum. Further, sellers who are exploding-offer averse, like in our experiment, would use even fewer exploding offers.

The results and implications of this work fill a previously unexamined area in the literature. Most theoretical and experimental work examine exploding offers in labor markets, where buyers make exploding offers to sellers. For instance, in theoretical work, Lippman and Mamer (2012) characterize under which conditions a buyer, seeking to purchase an asset from a seller, will use exploding offers. In experimental work, Niederle and Roth (2009) show that matching markets with exploding offers—together with binding acceptances—create early and dispersed transactions and lower match quality. Lau et al. (forthcoming) find experimental employees hired through exploding offers exhibit less effort for their employers, leading to welfare losses for both sides. Tang et al. (2009) frame their experimental Ultimatum Deadline Game as a hiring problem. Proposers offer responders some amount of time to make a decision. They find experimental proposers tend to set deadlines that are too short, and their offers are frequently rejected. Only Armstrong and Zhou (2013) explicitly model a consumer goods market. Their model, the theoretical basis for this paper, involves sequential consumer search where multiple firms choose whether or not to use exploding offers and set prices accordingly.

Our paper also relates to experimental studies in sequential search markets. Early studies in sequential search markets focus on the optimal stopping rule when individuals faced price or wage offers (Schotter and Braunstein, 1981; Cox and Oaxaca, 1989; Kogut, 1990). Those experiments evaluated individuals' search behavior when uncertain price/wage offers follow a known distribution and searching involves a constant search cost. They find that consumers tend to stop earlier, compared with risk neutral consumers, who only care about marginal expected gains. The literature naturally extends to more general experimental markets where sellers make price offers and buyers make purchase decisions (Grether et al., 1988; Cason and Friedman, 2003). That research involves testing equilibrium price and evaluating market performance. For example, Cason and Friedman (2003) test "noisy search equilibrium" using both computer buyers and real buyers. Our paper builds on this strand of literature by augmenting traditional search experimental designs with the possibility of exploding offers. Unlike the previous findings, the use of exploding offers generally leads buyers to search *longer* than optimal, as buyers are more likely to reject sellers' exploding offers and continue their search.

Of course, one simplification we have made is using exploding offers as the only example of search deterrence in our experiment. We specifically chose exploding offers rather than other forms of search deterrence, because we believed it would be the strategy most likely to provoke a negative response from buyers. Other less aggressive search deterrence strategies are actually optimal under more conditions than exploding offers (Armstrong and Zhou, 2013). We leave it as a future extension of our work to experimentally investigate buyers' response to such search deterrence.

References

Mark Armstrong and Jidong Zhou. Search deterrence. *Working Paper, University of Oxford*, 2013.

Christopher Avery, Christine Jolls, Richard A. Posner, and Alvin E. Roth. The market for federal judicial law clerks. *The University of Chicago Law Review*, 68(3):793–902, 2001.

Christopher Avery, Christine Jolls, Richard A. Posner, and Alvin E. Roth. The new market for federal judicial law clerks. *The University of Chicago Law Review*, 74(2):447–486, 2007.

Yaron Azrieli, Christopher P. Chambers, and Paul J. Healy. Incentives in experiments: A theoretical analysis. 2014.

Colin F. Camerer. *Behavioral Game Theory*. Princeton University Press, 2003.

Timothy N. Cason and Daniel Friedman. Buyer search and price dispersion: a laboratory study. *Journal of Economic Theory*, 112(2):232–260, 2003.

James C. Cox and Ronald L. Oaxaca. Laboratory experiments with a finite-horizon job-search model. *Journal of Risk and Uncertainty*, 2(3):301–329, 1989.

Catherine C. Eckel and Philip J. Grossman. Forecasting risk attitudes: An experimental study using actual and forecast gamble choices. *Journal of Economic Behavior & Organization*, 68(1):1 – 17, 2008.

Urs Fischbacher. z-tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics*, 10(2):171–178, 2007.

Shane Frederick. Cognitive reflection and decision making. *Journal of Economic Perspectives*, 19(4): 25–42, 2005.

Ben Greiner. The online recruitment system ORSEE 2.0 - a guide for the organization of experiments in economics. *University of Cologne*, 2004.

David M. Grether, Alan Schwartz, and Louis L. Wilde. Uncertainty and shopping behaviour: An experimental analysis. *Review of Economic Studies*, 55(2):323–42, 1988.

Carl A. Kogut. Consumer search behavior and sunk costs. *Journal of Economic Behavior & Organization*, 14(3):381–392, 1990.

Nelson Lau, Yakov Bart, Joseph Neil Bearden, and Ilia Tsetlin. Exploding offers can blow up in more than one way. *Decision Analysis*, forthcoming.

Steven A. Lippman and John W. Mamer. Exploding offers. *Decision Analysis*, 9(1):6–21, 2012.

Richard D. McKelvey and Thomas R. Palfrey. Quantal response equilibria for normal form games. *Games and Economic Behavior*, 10(1):6 – 38, 1995.

Richard D. McKelvey and Thomas R. Palfrey. Quantal response equilibria for extensive form games. *Experimental Economics*, 1(1):9–41, 1998.

Muriel Niederle and Alvin E. Roth. Market culture: How rules governing exploding offers affect market performance. *American Economic Journal: Microeconomics*, 1(2):199–219, 2009.

Alvin E Roth and Xiaolin Xing. Jumping the gun: Imperfections and institutions related to the timing of market transactions. *The American Economic Review*, pages 992–1044, 1994.

Michael Rothschild. Searching for the lowest price when the distribution of prices is unknown. *The Journal of Political Economy*, pages 689–711, 1974.

Andrew Schotter and Yale M. Braunstein. Economic search: An experimental study. *Economic Inquiry*, 19(1):1–25, 1981.

George J Stigler. The economics of information. *The journal of political economy*, pages 213–225, 1961.

Wenjie Tang, J. Neil Bearden, and Ilia Tsetlin. Ultimatum deadlines. *Management Science*, 55(8): 1423–1437, 2009.