Search Deterrence in Experimental Consumer Goods Markets*

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

Ajalavat Viriyavipart
Department of Economics
American University of Sharjah

Xiaoyuan Wang†
School of Management and Economics
University of Electronic Science and Technology of China

January 26, 2017

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. Sellers demonstrate a reluctance to use such offers against human buyers that is lessened when facing computerized buyers. Human buyers are three times more likely to deviate from optimal strategy by rejecting rather than accepting these offers. The differential rate of buyer suboptimal play shifts the equilibrium of the game to a point where seller gains from the use of exploding offers are greatly reduced. In sum, the results suggest the benefits of search deterrence may be substantially less than what theory predicts.

Keywords: exploding offer, search deterrence, experimental economics, game theory
JEL: C91 D21 L10 M31

*We thank the Texas A&M Humanities and Social Science Enhancement of Research Capacity Program and the National Science Foundation of China (Funding number 71502023 and 71531003) for providing generous financial support. This paper received the Most Outstanding Paper Award at the 5th Annual Xiamen University International Workshop on Experimental Economics. We benefited from helpful comments by Gary Charness, James Cox, Catherine C. Eckel, Ernst Fehr, Ian Krajbich, Silvana Krasteva, Nikos Nikiforakis, Charles R. Plott, Michael Price, seminar participants at the Consumer Financial Protection Bureau, Georgia State University, Thammasat University, Chulalongkorn University, Southwest University of Finance and Economics and conference participants of the 2013 North American Economic Science Association, 2013 Southern Economic Association Meetings, and the 5th Annual Xiamen University International Workshop on Experimental Economics. Daniel Stephenson, J. Forrest Williams and Wei Zhan provided invaluable help in conducting the experimental sessions.

†Corresponding author. Please contact at wangxy@uestc.edu.cn.

1 Introduction

Basic supply and demand models can often explain transactions in centralized markets quite well. When markets are decentralized and more detail is required, search models are commonly used. Consumers sample prices from a variety of producers, buying once the price of goods falls below a certain threshold (Stigler, 1961; Rothschild, 1974). The basic model assumes sellers cannot affect buyer search. However, Armstrong and Zhou (2016) relax this assumption and show under relatively mild conditions, it is unilaterally profitable for sellers to deter search. Specifically the strategies of exploding offers, "buy-now" discounts, and requiring deposits for the option to buy later are profitable for sellers. Assuming sellers can implement such strategies, a natural question to ask is why such offers are not used more often in market transactions. One possibility is that producers are justifiably concerned that consumers may respond more negatively to these tactics than theory predicts.[1]

To test this line of reasoning, this paper utilizes an experiment[2] to implement a simplified version of Armstrong and Zhou (2016): two sellers simultaneously choose from one of three prices and either make an exploding or non-exploding offer.[3] Buyers, previously unaware of their personal value for either seller's good, randomly visit one seller and learn their (personal) value for that seller's good. In

[1]Other possibilities include the difficulty for sellers to identify customers and the credibility issues associated with such an offer. While these two are independent from our research question and not issues in our experimental design, we will not ignore these issues when discussing search deterrence in the field.

[2]For research questions such as these, experimental methodology is essential. The laboratory setting solves the issue of credibility with search-deterrent offers. An exploding offer in the laboratory will expire with certainty. In the field, firms face strong incentives to appear to use search deterrence, prodding customers to buy quickly without looking at competing offers. However, should a customer refuse their initial offer and return, firms would face incentives to sell to the customer at a new price (based on updated information), ensuring a profitable sale, rather than keeping their word on their initial search-deterring offer. It would be impossible to interpret a consumer's decision in response to a search-deterring offer without first knowing his beliefs on the credibility of the offer.

[3]Exploding offers are offers which are valid for only a short time period; the period is short enough to ensure that the receiver of the offer cannot find any new competing offers before deciding on the current offer. Exploding offers are limiting cases of "buy-now" discounts and deposits, the other two types of search deterrence tactics. If a deposit or discount is sufficiently high, a buyer will be forced to either accept the offer immediately or reject it entirely.

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.

Because sellers' behavior is conditional on perceived buyer response, two treatments are used to isolate seller behavior. In one, sellers knowingly interact with computer buyers programmed to follow optimal strategy; in the other, they interact with human buyers. Common across both treatments is the equilibrium prediction. The opportunity to use exploding offers removes the incentives for sellers to compete on price; they will play a pure strategy of charging the highest price with an exploding offer. Buyers will reject or accept exploding offers, unbiasedly, based upon expected value and profit maximization.

In contrast to the equilibrium prediction, human buyers are 3-times more likely (i.e., 15 probability points more) to violate profit-maximizing strategy by rejecting an exploding offer than accepting it. This differential rate of suboptimal play alters sellers' incentives; seller best-response is to charge the lowest price, albeit with an exploding offer. In a deviation from this best response, sellers use exploding offers about 2/3 as often against human than computer buyers. In sum, sellers play the equilibrium strategy five times as often against computer than human buyers.

The reasons behind these deviations are likely different for buyers and sellers. Buyers who score lowest on the Cognitive Reflection Test (Frederick, 2005) are most likely to reject an exploding offer they should accept. Such buyers are no more or less likely to accept an exploding offer they should reject, creating a differential negative response to exploding offers. The same test does not predict seller exploding offer use, however. On exit surveys, about a third of sellers indicate a reliance on free-recall offers, often justified by how an exploding offer would affect buyers. Sellers that give this response are more likely to be female and are far less likely to use exploding offers.

The results of this experiment provide an explanation as to why search deterrence in consumer markets may not be as widespread as theory might imply. Buyers' deviation from profit maximizing behavior—primarily due to the intuitive (also "fast" or "lazy", see Kahneman (2011)) response to reject an exploding offer without corresponding forces to accept exploding offers—disproportionately harms this sales tactic. Sellers may also demonstrate additional reluctance to use search deterrence out of concern for buyers. Buyer deviations can alter the equilibrium for seller pricing, removing any profit sellers gain from having the ability to use search deterrence. Provided the use of search deterrent strategies requires a cost to firms (perhaps in the form of selecting/training a sales staff to use this tactic without reservation or in developing technology to make offers credible and/or track customers), it would not necessarily make sense for a firm to use search deterrence.

Thinking of all the markets where search deterrence is conceivable requires some imagination by the reader (for more examples, see Armstrong and Zhou, 2016). Markets with professional salespeople (e.g., door-to-door sales, automobiles, real-estate, club memberships, fine jewelry, mortgages, insurance policies) are the most obvious examples. Electronic transactions may provide the best opportunity for these offers. The internet greatly reduces the costs of maintaining, tracking, and notifying customers for search deterrence. "Daily deals" or "flash sale" websites usually market limited-time offers to registered customers by email, text or social networks (each day or at customers' desired frequencies). Since the customer does not see inventory, it can be manipulated or misrepresented to make this limited-time frame credible. One could also envision extensions to markets where demand exhausts supply (i.e., tickets to events, travel, and restaurants during busy hours). Even grocery stores, long thought of as the one area where search deterrence would be impossible, can notify shoppers of personal limited time coupons on their shopper's card, essentially a personal, buy-it-now discount.[4]

[4] The theoretical model of Armstrong and Zhou (2016) and experimental design used here both assume prices are known at both sellers, but value can only be observed when one encounters a seller. A possible alternative is that *value* is known, but *prices* can only observed when one encounters a seller. Admittedly, some of the previous examples (often the ones with professional salespeople)

Granted our experimental design focuses on a limiting case of search deterrence, exploding offers. We use this case because we think it is the most transparent way to model search deterrence to our subjects in an experimental game.[5] However, we are examining the validity of the assumptions of a general theory of search deterrence and its implications on an entire market. Because the theory applies to all forms of search deterrence, we believe our results—through the lens of theory—also have relevance to all forms. In the above cases, our results suggest buyers require a great deal of deliberation (and possibly risk-aversion) to avoid the bias towards rejecting offers. Since the source of the bias appears to be cognitive in nature—it is associated with quick, intuitive thinking—we would expect the problem to be present in a variety of field situations. Further, the distinction between exploding offers and other forms of search deterrence may be quite ambiguous in the field. A buyer simply may not know whether he can return to the seller again after the offer expires to acquire the good. It is not clear this distinction is particularly relevant. If the initial offer from the seller appears to be below the market price and the item can be easily obtained at the market price, a buyer simply may not care if she can return to the same seller to buy later at the market price.

2 Related Literature

Exploding offers, as well as other search deterrence strategies, may be commonly observed in bilateral negotiation markets. Cialdini (2003) has documented the adoption of deterrence strategies in search markets.

> "A prospective health-club member or automobile buyer might learn that the deal offered by the salesperson is good for that one time only;

are more likely to resemble the former case, while some (often the internet cases with non-financial goods) are more likely to resemble the latter. This distinction makes little difference theoretically, but conceivably might alter the behavioral responses of buyers. We leave this option open for future research.

[5]There are possible psychological differences in the framing of an exploding offer compared to other forms of search deterrence. We return to this issue in our final section.

> should the customer leave the premises, the deal is off...A home vacuum cleaner operation I infiltrated instructed its sales trainees to claim that, 'I have so many other people to see that I have the time to visit a family only once. It's company policy that even if you decide later that you want this machine, I can't come back and sell it to you. (page 208)"

Xiong and Chen (2014) provide indirect evidence of exploding offers; they demonstrate such offers would be feasible and profitable (as a form of price-discrimination to some customers) in industries (i.e., health clubs, spa, and education) that give customers free trials.[6] Empirical work on this issue is lacking. Armstrong and Zhou (2016) note "Because inducements for quick decisions are often offered casually during the course of a one-to-one sales encounter, and because opportunist price hikes are not publicly announced, it is hard to obtain empirical evidence about this form of price discrimination." For these reasons we suspect experimental methods may be the only way to use data to investigate issues of search deterrence.

Most theoretical and experimental work on search deterrence examines exploding offers in labor markets, where buyers make exploding offers to sellers, and may be extremely constrained in both their number of offers and purchases. These unique features, not found in consumer goods markets, makes it difficult, if not impossible, to extrapolate the preceding works to consumer goods markets.[7] Lippman and Mamer (2012) theoretically model this setup and conclude the optimal offer choice for buyers varies greatly depending on underlying assumptions. The labor market focus also produces experimental results that are not particularly relevant to consumer goods markets. For instance, Niederle and Roth (2009) show that matching markets with exploding offers—together with binding acceptances—

[6] They provide 24 Hour Fitness, Rancho La Puerta, and New Orleans School as examples of firms in these three types of industries that use the free-trial practice, though not direct evidence of any search deterrence by these firms.

[7] There are also clear cut examples of labor markets where exploding offers are the norm. Law students applying for appellate court clerkships frequently receive exploding offers (Roth and Xing, 1994; Avery et al., 2001, 2007; Niederle and Roth, 2009).

create early and dispersed transactions and lower match quality. Lau et al. (2014) find experimental employees hired through exploding offers exhibit less effort for their employers, leading to welfare losses for both sides. Tang et al. (2009) examine the decision of employers to select the length of time on an exploding offer to a prospective employee. They find experimental proposers tend to set deadlines that are too short, and their offers are frequently rejected.

Neither Armstrong and Zhou (2016) nor our paper includes these preceding features (i.e., matching quality, reciprocity after market transaction, variable exploding offer length). Their paper, the theoretical basis for ours, involves sequential consumer search where multiple firms choose whether or not to use search deterrence and set prices accordingly. The authors concede their model does not incorporate "behavioural factors," but speculate these factors could either make search-deterrent strategies more or less profitable. The primary purpose of this paper concerns these factors and their implications on market outcomes.

Also related are experimental studies in sequential search markets. Early studies 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 evaluate 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. More general experimental markets—where sellers make price offers and buyers make purchase decisions—involve testing equilibrium price and evaluating market performance (Grether et al., 1988; Cason and Friedman, 2003). For example, Cason and Friedman test "noisy search equilibrium" using both computer buyers and real buyers. Our paper builds on this strand of literature by augmenting these traditional search designs with the possibility of search deterrence. Contrary to previous findings, our main result implies buyers generally search *longer* than what theory predicts.

3 The Model

The experiment in this paper implements a simplified model based on Armstrong and Zhou (2016). There are a few major changes from the literature. First, the only search deterrence sellers may use is an exploding offer. Second, buyers are aware of all sellers' pricing decisions immediately; this results in optimal buyer strategies that do not require assumptions on the distribution of seller strategies. To avoid buyers strategically avoiding sellers that use exploding offers (and potentially eliminating the use of such strategies in equilibrium), buyers randomly encounter sellers, so that search order is exogenous.

We also discretize buyer valuations and seller prices. 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 2 (at the end of this section) provides payoffs for a seller given a fixed offer and a pricing strategy, conditional on the other seller's offer and pricing strategy. The table will be used as a theoretical benchmark for analysis of sellers' choices in the experimental game.

3.1 The Search

This model represents an experimental search market of two sellers with one buyer who visits each seller sequentially in a random order.[8] Each seller offers a good which has a private value for the buyer drawn from the same ex-ante value distribution: $V_k^i \in \{V_1, V_2, ..., V_K\}$ (where $i = 1, 2$ represents sellers and $k = 1, 2, ..., K$ represents K possible values) with probability $\upsilon_1 \equiv prob(V_1), \upsilon_2 \equiv prob(V_2), ..., \upsilon_K \equiv prob(V_K)$. The game is as follows:

1. Each seller sets a price from a possible price range: $P^i \in \{P_1, P_2, ..., P_L\}$ and chooses an offer type as either an exploding or a free-recall offer.

[8] Several identical buyers were used in our experiments to reduce noise in seller realized payoffs. Each seller can only choose one strategy for all buyers in each period.

2. Nature randomly selects which seller the buyer will visit first (S^1).[9]

3. The buyer observes the prices of both sellers (P^1 and P^2) and his value of the first good he[10] 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 equals 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.

3.2 Buyer Best Response

We assume that the buyer is rational and has an objective to maximize his expected payoff. Because 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.

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

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

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.[11] 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 $\Pi^1 = V^1 - P^1$ whereas the expected payoff from visiting S^2 is

$$E(\Pi^2) = \sum_{k=1}^{K} \upsilon_k{}^{*}\max(0, V_k^2 - P^2).^{12} \tag{1}$$

The buyer accepts the first offer if $\Pi^1 > E(\Pi^2)$ and rejects otherwise.[13] If the first offer was rejected, the buyer would accept the second offer only if $V^2 > P^2$.

3.3 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 the buyer would visit 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.

[11] 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.

[12] 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.

[13] 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.

First, consider a case where both sellers use exploding offers. Consider seller i with a price P^i, who plays with seller j with a price P^j. There are two possible situations that occur with equal probability:[14]

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., $V^i_k - P^i > E(\Pi^j) = \sum_{l=1}^{K} v_l{}^*\max(0, V^j_l - P^j)$ and reject otherwise. The probability that he will accept the offer is

$$\text{Prob}(\text{accept } i_1) = \sum_{k=1}^{K} v_k{}^* D^i_k \tag{2}$$

   where $D^i_k = 1$ if $V^i_k - P^i > E(\Pi^j)$ 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 $\sum_{l=1}^{K} v_l{}^* D^j_l$ where $D^j_l = 1$ if $V^j_l - P^j > E(\Pi^i) = \sum_{k=1}^{K} v_k{}^*\max(0, V^i_k - P^i)$ and $= 0$ otherwise. If the buyer rejects the offer from seller j, he will visit seller i. Upon visiting seller i, he will accept the offer if his value of product i (V^i) is above P^i or with probability $\sum_{k=1}^{K} v_k{}^* B^i_k$ where $B^i_k = 1$ if $V^i_k > P^i$ and $= 0$ otherwise. So, the probability that the buyer will purchase from seller i is

$$\text{Prob}(\text{accept } i_2) = (1 - \sum_{l=1}^{K} v_l{}^* D^j_l)^* \sum_{k=1}^{K} v_k{}^* B^i_k. \tag{3}$$

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

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) $V^i_k - P^i > V^j_l - P^j$ and (2) $V^i_k - P^i > 0$. The

[14] For simplicity, we assume the same probability of visiting each seller first. It is possible to assume different probabilities.

Table 1: Choices of Valuation Distributions

High buyer-heterogeneity
$V \in \{10, 25, 40, 55, 65, 70\}$, $v_1 = v_2 = v_3 = v_4 = 0.125$, and $v_5 = v_6 = 0.25$, $P \in \{25, 30, 35\}$.
Low buyer-heterogeneity
$V \in \{10, 25, 40, 55, 65, 70\}$, $v_1 = v_2 = v_5 = v_6 = 0.1$, and $v_3 = 0.2$, and $v_4 = 0.4$, $P \in \{25, 30, 35\}$.

probability that the buyer will purchase from seller i is

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

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*}Prob(accept i_3).

Last, consider a case where one seller uses an exploding offer and another seller uses a free-recall offer. Because an offer type of the second seller has no effect on the buyer's 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 $P^{i*}[\frac{1}{2}$Prob(accept i_1)+$\frac{1}{2}$Prob(accept i_3)].[15] If seller i uses a free-recall offer while seller j uses an exploding offer, seller i's expected payoff is $P^{i*}[\frac{1}{2}$Prob(accept i_3)+$\frac{1}{2}$Prob(accept i_2)].[16]

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

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

3.4 Parameter Choices for Experimental Search Markets

The previous analysis shows how payoffs are calculated in this game. For any sets of values $V_k^i \in \{V_1^i, V_2^i, ..., V_K^i\}$, probabilities $\upsilon_1, ..., \upsilon_K$, 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. Table 1 provides parameter choices used in our experiment. We use the same sets of values and prices but different probabilities between the high and low buyer-heterogeneity distributions. The high-buyer-heterogeneity distribution offers a greater variance of buyers' values for items than the low-buyer-heterogeneity distribution.[17] One interpretation of this heterogeneity is that it represents the underlying competitiveness of the market depending on whether sellers have greater similarity between products.

The valuation structures provide different incentives for buyers across distributions. Suppose, for example, a buyer faces the strategy of 35E from both sellers. Under the high-buyer-heterogeneity distribution, 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. In contrast, under the low-buyer-heterogeneity distribution, a buyer would accept the first offer if his value for the first item is either 55, 65 or 70. Under both distributions, if the first offer was rejected, the second offer would be accepted as long as his value for the second item is above 35 (i.e., the values of 40, 55, 65, 70).

Buyer optimization provides different expected payoffs for each seller-strategy pair across each treatment. Table 2 provides a matrix of these values.[18] Under both distributions, the unique equilibrium for sellers involves an exploding offer with the highest price of 35 (points).

[17] Specifically, the discrete high-buyer-heterogeneity distribution has a mean of 50, a variance of 450 and a skewness of -0.76; the discrete low-buyer-heterogeneity distribution has a mean of 47, a variance of 306 and a skewness of -0.78. Our intent was to change the variance without modifying the other moments greatly. However, we were constrained by the equilibrium prediction and the necessity of using relatively simple distributions with our subjects.

[18] The payoff matrices are calculated based on two assumptions: (1) when facing the same net values from two sellers, a buyer has the same probability in accepting each seller (2) a buyer rejects an offer with zero payoff.

Table 2: Expected Payoffs for One Seller's Strategy Choice Given Other Seller's Strategy Choice

	High buyer-heterogeneity Distribution					
	25E	30E	35E	25F	30F	35F
25E	11.33	12.50	12.50	13.67	15.04	15.63
30E	11.72	13.59	15.00	12.89	16.41	18.05
35E	13.67	13.67	**15.86**	14.22	15.04	19.14
25F	9.38	11.91	12.50	11.72	14.45	15.63
30F	9.61	11.25	14.30	10.78	14.06	17.34
35F	10.39	11.21	13.13	10.94	12.58	16.41
	Low buyer-heterogeneity Distribution					
	25E	30E	35E	25F	30F	35F
25E	11.50	11.50	15.50	13.50	14.94	15.25
30E	13.80	13.80	13.80	14.48	16.20	17.93
35E	9.10	16.10	**16.10**	9.45	16.89	18.90
25F	10.00	11.44	15.75	12.00	14.88	15.50
30F	10.28	12.00	13.73	10.95	14.40	17.85
35F	11.55	11.99	14.00	11.90	12.78	16.80

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

4 Experimental Design

The experiment consisted of two treatments. In the computer-buyer treatment (CB), human sellers were matched against computer buyers programmed to play optimal strategies. In the human-buyer treatment (HB), human sellers were matched against human buyers. Sellers were fully informed about the type of their buyers. Each group consisted of eight sellers (for all treatments) and sixteen buyers (only for the HB treatment). For the CB treatment, a session contained 2-3 groups (16 or 24 subjects) while for HB treatment, a session contained one group (24 subjects). In each period, four markets were randomly formed. Each market consisted of two sellers and twenty-four buying decisions from either computer buyers programmed to play optimal strategies (CB treatment) or four human subjects (HB treatment), each played six different possible buying decisions with randomly determined item-value pairs.[19] In each market, half of the buyers in the market visited one seller first and the other half visited the other seller first.

There were twenty total periods. Each period, buyers and sellers were randomly rematched into new markets, but the role of each subject (i.e., buyer or seller) was fixed for the entire session. In addition, the same random matching was used in every session and treatment.[20]

Each period began with sellers choosing a price and an offer type (i.e., exploding or free-recall offer). The seller's price and offer type were the same for all buyers that encountered the seller. Buyers would observe prices and offer types of both sellers in the market, but would only see the value of the item from the first seller they encountered. Each buyer's valuation for each of the six possible buying decisions was drawn independently from the known valuation distribution.

In the HB treatment, buyers played their six possible buying decisions sequentially (see Figure 1 for a screenshot of the first of six buying decisions).[21] In each decision, they chose whether to buy the item from the first seller immediately or visit the second seller. Visiting the second seller would allow the buyer to observe his personal value of the item from the second seller. If the first seller made a free-recall offer, the buyer could choose to visit the second seller and still have the opportunity to buy the item from the first seller. If the first seller used an exploding offer, the buyer could not buy the item from the first seller after observing his valuation from the second seller.

[19] Experimental economics often relies on techniques (e.g., multiple price lists, strategy method) where subjects are given several choices that might occur and paid for the one that does occur. Our intent on the buyer side is to learn what buyers would do for all 36 item-value combinations. Unfortunately, having buyers play through all possible buying decisions each period is not feasible, both because of time constraints and buyer fatigue. Instead we use 6 randomly selected item-value combinations. We view 6 as the highest number of decisions we could ask buyers to make each round without encountering fatigue or time limit issues.

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

[21] There was no time limit placed on buyers to make these decisions. In most cases buyers made their decision in less than a minute.

Figure 1: Decision on One of Six Items, Human-Buyer Treatment. The first seller has used an "Offer B" (a free-recall offer) so the buyer can choose to search and observe his value of the second seller's item without losing the option to buy from the first seller.

After each of the four buyers had completed their six buying decisions, screens showed sellers the outcome of all twenty-four buying decisions in their market. One screen (Figure 2a) showed the price and strategy used by themselves and the other seller in the market, the number of items sold by each seller and the total profit for each seller. Another screen (Figure 2b) provided information about each of the twenty-four buying decisions in the market. Sellers were provided this, admittedly, large amount of feedback to provide the best opportunity for them to best respond to buyers over the course of the experiment.

It is important to note that our theoretical model featured a game with two sellers and one buyer. Our experiment has departed a little from this design for pragmatic reasons. First, sellers are compensated based on 24 buying decisions rather than just one. This reduces the noise in the realization of payoffs for sellers, though it does not alter the risk-neutral strategy of the game. Buyers play 6 different

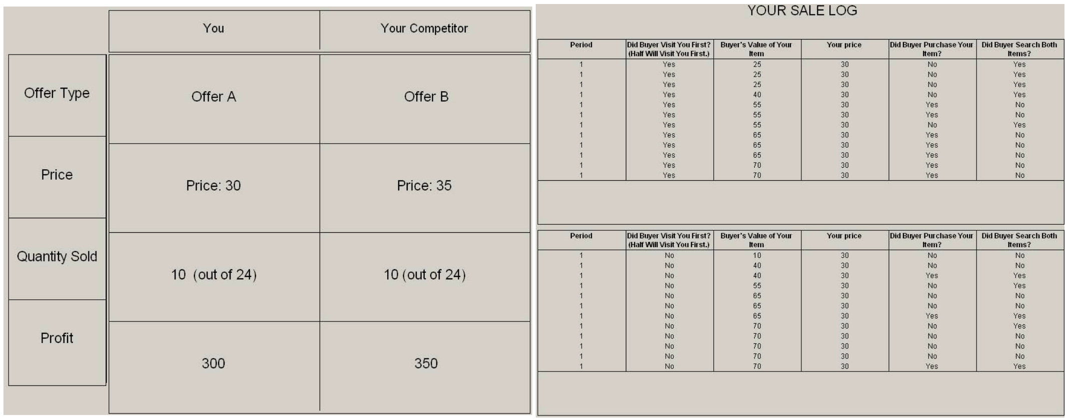

Figure 2: A Seller's Feedback Screen at the End of the Period. Sellers could toggle between each of the screens. (a, left) Both sellers are informed on the performance of each other in the market in aggregate. (b, right) Each seller observes all twenty-four buying decisions.

possibilities of this same game with identical seller strategies as a way to get more information about buyer decisions under different values for the seller items (see footnote 19). While at most one of these decisions affects buyer's payoffs, all six potentially affect sellers' payoffs. This decision was made because it would have been impractical—both in terms of cost and physical space—to pair 24 buyers with each pair of sellers.

Before each session began, the instructions were both shown on screen and read aloud to ensure the game was common knowledge among the subjects. After the instructions, the subjects answered a quiz, in multiple choice form, to establish that they understood how to play the game. Each subject needed to answer all questions correctly before the game started. Throughout the experiment, to avoid any priming effects associated with language, exploding offers were referred to as "Offer A" and free-recall offers were referred to as "Offer B."

After the twenty periods had elapsed, subjects filled out a questionnaire consisting of demographics information, a risk-preference test (similiar to Eckel and Grossman, 2008), and a Cognitive Reflection Test (Frederick, 2005). Subjects were then privately paid their earnings in the session (plus a five dollar show-up bonus) in cash. Each seller in both treatments was paid based on one randomly selected period.[22] 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.[23] For an 80 minute session, subjects earned $18, on average.

The experiment was conducted in the Economic Research Laboratory at Texas A&M University. Two (32 sellers) and three (24 sellers, 48 buyers) sessions of the CB and HB treatments, respectively, were conducted in April and October 2013 using the high-buyer-heterogeneity distributions (see Table 1). As an additional robustness check on the main results and to demonstrate that results found in both treatments are not due to specific idiosyncrasies associated with parameter choices, three (56 sellers) and four (32 sellers, 64 buyers) sessions of the CB and HB treatments, respectively, were conducted in February 2016 using the low-buyer-heterogeneity distributions. All 256 subjects were Texas A&M University undergraduate students recruited campus wide using ORSEE (Greiner, 2015). The experiment was programmed and conducted with the software Z-tree (Fischbacher, 2007).

[22]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.

[23]We chose a larger conversion rate for buyers to make it similar to the field. In the field, sellers make many individual transactions earning a small margin on each. Buyers make one transaction and earn the entire surplus. If cost and the capacity of an experimental laboratory were not binding, we might have 24 buyers for every two sellers in the lab. Pragmatically we cannot do this. So instead sellers are compensated based on six buying decisions from each of the four buyers while buyers are compensated based on one buying decision. The differential rates of payment were common knowledge to all participants.

5 Results

Result 1 *Sellers play the equilibrium strategy five times more often with computer rather than human buyers, making exploding offers less often and offering lower prices in the latter treatment. Over time, sellers increase prices, offer more exploding offers and play equilibrium strategy more often in the computer treatment. Sellers lower prices and play equilibrium strategy less often in the human buyer treatment.*

Table 3, Panel A provides a breakdown of seller exploding offer use by treatment, collapsed to the subject level. Sellers use exploding offers roughly 2/3 as often against human buyers than computer buyers (66% CB vs. 47% HB; t-test and Mann-Whitney-Wilcoxon, $p < 0.01$).[24] This difference is more pronounced under the low-buyer-heterogeneity than the high-buyer-heterogeneity distribution, though both are significant at the 5%-level when evaluated separately. Table 3, Panel B shows seller pricing decisions have similar differences across treatments. Sellers offer prices roughly 3 points lower with human buyers than with computer buyers (29.764 CB vs. 26.795 HB; t-test and Mann-Whitney-Wilcoxon, $p < 0.01$), the result and significance do not noticeably change when evaluating distributions separately.

As a result, the equilibrium strategy of 35E, charging the highest price with an exploding offer, is often utilized against computer buyers, but rarely used against human buyers. Table 3, Panel C reveals that this equilibrium strategy is played roughly one-fourth of the time in the CB treatment and less than 5% of the time in the HB treatment ($p < 0.01$, t-test and Mann-Whitney-Wilcoxon, whether distributions are pooled or evaluated separately). In sum, there is little evidence of equilibrium play by sellers against human buyers.

As a robustness check on these main results, we also categorize them at the cohort level, to alleviate the concern they could be driven by cohort-level effects.[25]

[24] Significance testing, if not explicitly mentioned otherwise, assumes independent observations at the subject level. That is, we are not counting multiple observations per subject as independent.

[25] We provide this cohort-level data and tests in the interest of full disclosure for the skeptical reader.

Table 3: Rate of Exploding Offers, Equilibrium Play and Average Offer Price Collapsed to Seller

Panel A: Rate of Exploding Offer Use by Seller (mean, standard deviation, number of observations)			
Distribution:	Both	High Buyer-Heterogeneity	Low Buyer-Heterogeneity
Human Buyer	0.474 (0.252) 56	0.546 (0.253) 24	0.420 (0.242) 32
Computer Buyer	0.663 (0.198) 88	0.680 (0.190) 32	0.654 (0.203) 56

Panel B: Average Offer Price by Seller (mean, standard deviation, number of observations)			
Distribution:	Both	High Buyer-Heterogeneity	Low Buyer-Heterogeneity
Human Buyer	26.795 (1.502) 56	26.948 (1.329) 24	26.680 (1.630) 32
Computer Buyer	29.764 (1.860) 88	30.359 (2.126) 32	29.424 (1.613) 56

Panel C: Rate of Equilibrium Play by Seller (mean, standard deviation, number of observations)			
Distribution:	Both	High Buyer-Heterogeneity	Low Buyer-Heterogeneity
Human Buyer	0.046 (0.074) 56	0.046 (0.067) 24	0.046 (0.080) 32
Computer Buyer	0.226 (0.211) 88	0.323 (0.238) 32	0.171 (0.172) 56

Appendix Table A.1 provides a listing of all eighteen cohorts and their respective averages of our three dependent variables. Using a two-tailed randomization test, the difference between means in human vs. computer buyer cohorts is the 12th ($p < 0.01$), 1st ($p < 0.01$), and 40th ($p \approx 0.01$) largest magnitude possible of 31,824 possibilities[26] for exploding offer use, price, and equilibrium play, respectively. Regression analysis using both standard and wild bootstrap (Cameron et al., 2008) clustering at the cohort level finds the differences between human and computer buyer treatments significant at the 1% level (see Appendix Tables A.3 and A.4).

Figure 3 displays the dynamics of seller decisions across treatments. Over twenty periods, sellers in the the CB treatment appear to increased their use of exploding offers, the relation is inconclusive for sellers in the HB treatment. Linear trend analysis confirms this finding; seller exploding offer use is predicted to increase by 1.5 probability points each period in the CB treatment, but remain unchanged in the HB treatment (see Appendix Table A.7 for more detail). By the last five periods, about 73% of sellers in the CB treatment used an exploding offer, whereas only about 51% of sellers in the HB treatment used an exploding offer ($p < 0.01$, for both t-test and Mann-Whitney-Wilcoxon).

Figure 4 displays seller price dynamics across treatments. In the early periods, average prices across treatments are very similar. After that, they diverge. Seller prices increase in both CB treatment but decrease in both HB treatment. A similar pattern is found in equilibrium strategy use by sellers, shown in Fiqure 5; the use of equilibrium strategies is increasing in the CB treatment and decreasing in the HB treatment. Linear trend analysis confirms both trends as significant (see Appendix Table A.7).

Result 2 *Buyers deviate from optimal, profit-maximizing strategies more often when they*

As Fréchette (2012) notes, there is little evidence of these cohort level effects existing in laboratory experiments except in cases where they are quite obvious ex ante.

[26]Since there are 7 human and 11 computer buyer cohorts, there are 31,824 possible mappings of our dependent variable to these cohorts.

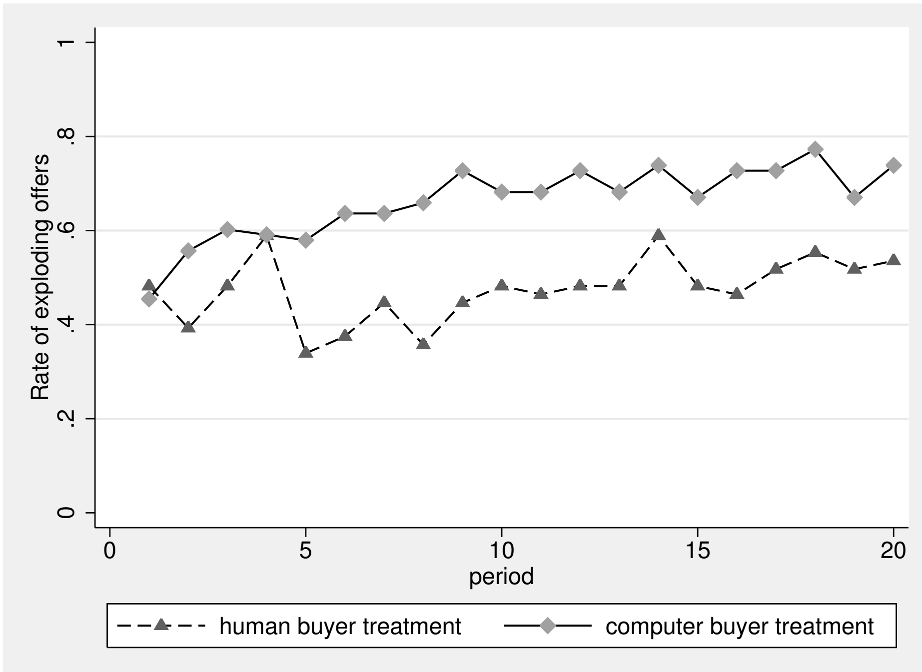

Figure 3: Rate of Exploding Offers Used by Sellers by Period, HB and CB Treatments

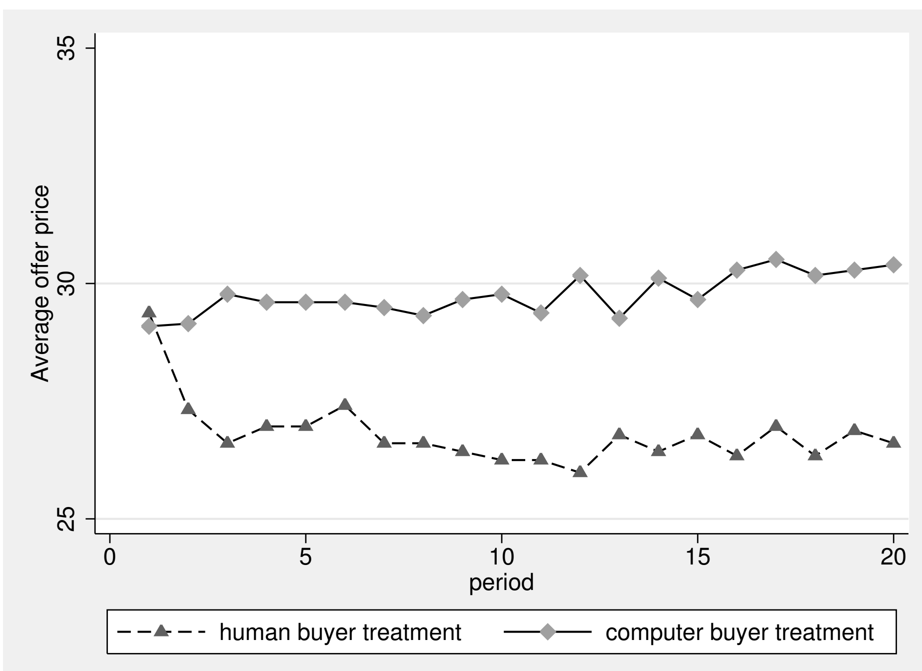

Figure 4: Average Price Offered by Sellers by Period, HB and CB Treatments

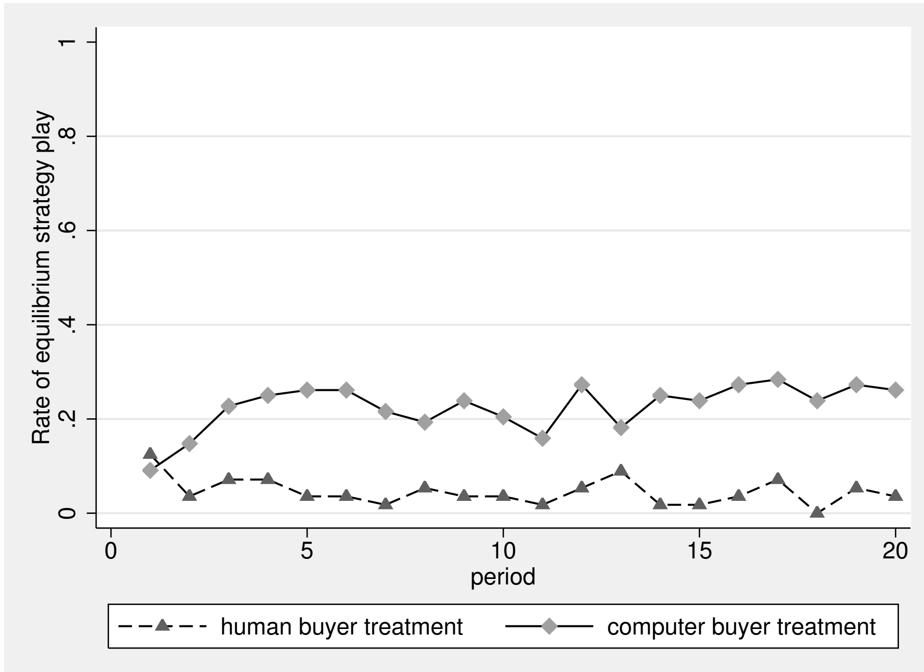

Figure 5: Rate of Equilibrium Play by Sellers by Period, HB and CB Treatments

encounter an exploding offer than a free-recall offer. With exploding offers, it is more common for buyers to reject an offer they should accept than accept an offer they should reject. This difference remains even after controlling for the costs of suboptimal play.

Buyers made 6 buying decisions in each period over 20 periods. Pooling the results from the 7 sessions of 16 buyers each, there are a total of 13,440 (6×20×16×7) buying decisions. Sellers used exploding offers in 47% of theses decisions. This results in 6372 exploding offer buying decisions and 7059 decisions with free-recall offers.[27]

The ways buyers could deviate from the optimal strategies differed based on the type of the offers they encountered. With a free-recall offer, buyers had the opportunity to learn all information about both items. They would only violate

[27]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.

Table 4: Rate of Deviations from Optimal Play by Type, Collapsed to Buyer

Panel A: Buyer Deviation Rate, Free-recall vs. Exploding Offers			
Valuation Type	Exploding Offer	Free-recall Offer	Difference (Paired)
Overall (112 subjects)	0.157 (0.125)	0.058 (0.108)	0.100*** (0.121)
High Buyer-Heterogeneity (48 subjects)	0.103 (0.080)	0.047 (0.082)	0.056*** (0.090)
Low Buyer-Heterogeneity (64 subjects)	0.198 (0.138)	0.066 (0.123)	0.132*** (0.132)

Panel B: Buyer Deviation Rate, Exploding Offer Rejection vs. Acceptance			
Valuation Type	Exploding Offer Should Be Accepted	Exploding Offer Should Be Rejected	Difference (Paired)
Overall (112 subjects)	0.214 (0.199)	0.077 (0.110)	0.137*** (0.223)
High Buyer-Heterogeneity (48 subjects)	0.143 (0.126)	0.036 (0.065)	0.107*** (0.154)
Low Buyer-Heterogeneity (64 subjects)	0.266 (0.226)	0.108 (0.126)	0.159*** (0.262)

the profit-maximizing strategy by making a choice (i.e., buy item 1, buy item 2, don't buy) that would not maximize their surplus. With an exploding offer, buyers would make decisions with imperfect information. They would violate the optimal strategy by choosing to reject (or accept) an exploding offer when the expected (net) value of continuing to the second item is negative (positive). Table 4, Panel A shows subject-level frequency of these two types of suboptimal play. Perhaps unsurprisingly given the greater difficulty of making a decision with incomplete information, buyers are roughly 10 probability points (or 3 times) more likely to deviate from the optimal play with an exploding offer than a free-recall offer. The average rate of suboptimal play with an exploding offer is 16% compared to 6% with a free-recall offer ($p < 0.01$, paired t-test and Wilcoxon signed rank, whether distributions are pooled or evaluated separately).

The suboptimal play buyers make with exploding offers can be further classified.

In some cases, buyers should accept an exploding offer, but reject it. In other cases, buyers should reject an exploding offer, but accept it. Overall, buyers encountered slightly more exploding offers they should accept than they should reject (3754 vs. 2618, 59% vs. 41%).[28] Buyers were much more likely to reject an exploding offer they should accept than accept an exploding offer they should reject. Table 4, Panel B provides subject-level frequencies of these two types of deviations from optimal play. Buyers make suboptimal decisions 14 percentage points (or 3 times) more often by rejecting an exploding offer than accepting one ($p < 0.01$, paired t-test and Wilcoxon signed rank, whether distributions are pooled or evaluated separately). The absolute levels of these rates are substantial. Under the low heterogeneity distribution, roughly 1 in 4 exploding offers *that should be accepted* were rejected.

As a robustness check we also may examine these results at the cohort level. As Appendix Table A.2 shows, there were seven human buyer cohorts. In all seven, the rate of suboptimal play was higher when buyers faced an exploding offer than a free-recall offer. Similarly, in all seven, the rate of suboptimal play was higher when buyers faced an exploding offer they should accept than one they should reject. For a two-tailed permutations test this is the (tied) highest treatment effect of 128 (2^7) possibilities, equivalent to a p-value of 0.016. Regressions with both standard clustering and wild bootstrap clustering at the buyer-cohort level find similar levels of significance ($p < 0.01$, see Appendix Tables A.5 and A.6).

The costs of suboptimal play vary by situation. Figure 6 classifies the rate of suboptimal play based on the expected costs of rejection. (Negative values indicate situations where rejection is optimal.) The figure indicates most suboptimal decisions with exploding offers occur where the cost of decision is relatively small (i.e., under 5 points). However, the skew to the right side of the graph further

[28]Note that each buyer in the experiment made 120 buying decisions. On average, a buyer encountered 54 free-recall offers, 39 exploding offers she should accept, and 27 and exploding offers she should reject under the high heterogeneity distribution. Under the low heterogeneity distribution, a buyer encountered 70 free-recall decisions, 30 exploding offers she should accept and 20 exploding offers she should reject on average. Note the ratios of exploding offers that should be accepted vs. rejected does not vary across valuation distributions.

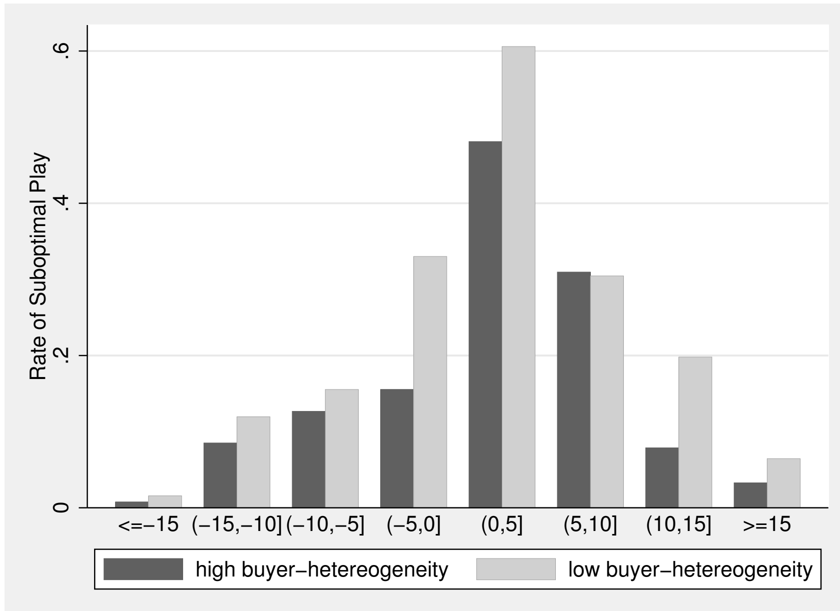

Figure 6: Rate of Buyer Suboptimal Play by Cost of Rejecting Exploding Offer

illustrates suboptimal play is more likely in rejection of exploding offers than in their acceptance. Inexplicably, suboptimal play occurs more often under the low heterogeneity distribution, even after controlling for costs.

Table 5 provides regressions of instances of buyer suboptimal play with exploding offers on the cost of suboptimality and other variables. The general form is

$$\text{logit}\left(\Pr\left(y_{ij} = 1 \left| c_{ij}, I_{ij}, d_{ij}, \alpha_j \right.\right)\right) = ac_{ij} + bI_{ij} + d'_{ij}\gamma + \alpha_i + \epsilon_{ij}. \tag{5}$$

where y_{ij} represents whether the subject i deviated from the optimal strategy when facing exploding offer j (1=yes, 0=optimal play). The variable c_{ij} is the value of the expected cost of playing the suboptimal strategy. Indicator variable I_{ij} represents whether the optimal strategy involves accepting an exploding offer (1=acceptance is optimal, 0=rejection is optimal). The term d_{ij} is a vector containing indicator dummies for each of the twenty periods and the treatment (when necessary); α_i is the random effect of subject i; ϵ_{ij} represents cluster-robust standard errors.

Table 5: Logistic Regression of Buyer Suboptimal Play on Cost of Suboptimality and Type of Optimal Play. Average marginal effects shown. For an alternate version of this table, with coefficient estimates see Appendix Table A.8.

	Both valuations		High buyer-heterogeneity		Low buyer-heterogeneity	
cost of suboptimality	-0.020*** (0.001)	-0.019*** (0.001)	-0.016*** (0.002)	-0.015*** (0.002)	-0.022*** (0.002)	-0.022*** (0.002)
acceptance is optimal	-	0.069*** (0.021)	-	0.057** (0.023)	-	0.079** (0.036)
valuation dummy	Y	Y	N	N	N	N
period dummies	Y	Y	Y	Y	Y	Y
random effects	Y	Y	Y	Y	Y	Y
observations	6372	6372	3144	3144	3228	3228
subject clusters	112	112	48	48	64	64
log likelihood	-1933.892	-1899.576	-732.985	-714.777	-1189.653	-1172.800

The regressions confirm the general observations from Figure 6. For every expected point cost of deviating from the optimal play, subjects are 2 probability points less likely to deviate. However, subjects are 6-8 probability points more likely to deviate from optimal play when optimal play involves accepting an exploding offer rather than rejecting it. Another way to look at the results is that for rates of suboptimal play to be equal with rejection and acceptance of exploding offers, the cost of rejecting an exploding offer would have to be 4 points greater. This is largely consistent with what is seen in Figure 6, as the figure would be symmetric if the x-axis were shifted by that amount.

Result 3 *Measures of propensity toward "quick" (vs. "slow") thinking explain buyers' over-rejection of exploding offers; measures of risk aversion explain buyers' under-rejection.*

Result 2 showed that buyers exhibited a greater tendency for suboptimal play when they encounter an exploding offer that should be accepted rather than rejected. Here the focus will be on what features of buyers are most likely to cause this incongruity.

Gender, risk and bounded rationality all might potentially explain this behavior. To isolate the effect of each term, we expand the regression model from (5), specifically examining the interaction of these terms and the "accepting an exploding offer is optimal" dummy variable. Formally,

$$\text{logit}\left(\Pr\left(y_{ij}=1\left|c_{ij},I_{ij},X_j,d_{ij},\alpha_j\right.\right)\right)=ac_{ij}+bI_{ij}+\tilde{a}I_{ij}c_{ij}+\beta x_j+\tilde{\beta}I_{ij}x_j+d'_{ij}\gamma+\alpha_i+\epsilon_{ij}. \quad (6)$$

The variables remain the same as before. The 3x1 vector x_i is added to the model to represent subject i's gender (1=male, 0=female), risk preference (above median=1, below median=0), and correct responses on the CRT (0, 1, 2, or 3). Variables are taken from non-incentivized survey questions. Gender was directly elicited from subjects in a demographic survey. The risk preference question (similiar to Eckel and Grossman, 2008) asked subjects to pick their preferred gamble from one of five increasingly risky 50/50 gambles. The grouping of preferred gambles, 1 or 2 vs. 3, 4, or 5—roughly half the subjects fall into each group—will be used as a proxy (with unit value if a subject chose 3, 4, or 5 in the test) for subject risk-preference. The number of questions subjects correctly answered on the CRT (i.e., 0, 1, 2, or 3) will be used as a proxy for subjects' propensity to inhibit "quick" (System 1) responses and engage in deliberative (System 2) decision making (see Kahneman, 2011).

Table 6 provides the results of this regression. To correctly interpret interaction terms, average marginal effects are calculated when acceptance of an exploding offer is optimal (i.e., "acceptance is optimal" dummy variable is 1) and when the rejection of an exploding offer is optimal (i.e., "acceptance is optimal" dummy variable is 0).

The coefficients of the regression are quite telling. As before, suboptimal play is more common (7 probability points more likely) when the optimal play involves accepting an exploding offer than rejecting it. The coefficients of risk aversion are in the direction one would suspect. Acceptance of an exploding offer leads to a certain payoff while visiting the second seller involves a random distribution of possible payoffs. Fittingly, buyers whose survey responses place them below median for risk tolerance are 3 probability points more likely to accept an optimal exploding offer

Table 6: Logistic Regression of Buyer Suboptimal Play on Cost of Suboptimality and Type of Optimal Play with Gender and Survey Data Interactions. Interactions on "acceptance [of exploding offer] is optimal" dummy variable. Average marginal effects at 0 and 1 of "acceptance is optimal" dummy shown. For coefficient estimates see Appendix Table A.9.

	When acceptance is optimal (acceptance is optimal=1)	When rejection is optimal (acceptance is optimal=0)
cost of suboptimality	-0.023*** (0.002)	-0.012*** (0.002)
above median risk tolerance	0.033 (0.027)	-0.047* (0.029)
correct CRT questions	-0.040*** (0.013)	-0.008 (0.012)
male	-0.004 (0.026)	-0.007 (0.028)
acceptance is optimal	0.073** (0.021)	
valuation dummy	Y	
period dummies	Y	
random effects	Y	
observations	6372	
subject clusters	112	
log likelihood	-1871.563	

and 5 probability points more likely to accept a suboptimal one, compared with the higher risk tolerance counterparts in the sample. While the former coefficient is not significant and the latter is only marginally, the difference between the two is significant ($p < 0.05$).

As one would expect, subjects who get more questions correct on the CRT are less likely to deviate from profit maximizing strategy with exploding offers. Interestingly, this effect is isolated to cases with exploding offers that should be accepted; an additional correct CRT response is equal to a 4 probability point reduction in rejecting an exploding offer that should be accepted. There is virtually no corresponding reduction in suboptimal play with exploding offers that should be rejected.

When combined, the propensity to give quick, intuitive (albeit incorrect) responses and risk aversion appear to explain subject deviations from optimal strategies, and jointly they play a critical rule in determining whether the suboptimal play of rejection and acceptance are symmetric or not. At the extreme case of propensity for deliberation and high risk aversion, we find no difference in the rates of suboptimal play (for the acceptance or rejection of an exploding offer); yet both an increase in risk tolerance or propensity for deliberate thinking may potentially lead to a net over-rejection of exploding offers. Appendix Table A.10 shows the marginal effects in Table 6 for a specific type of subjects, those in the risk-averse group who correctly answered all 3 CRT questions. At this level, the coefficient of the "acceptance is optimal" dummy variable is essentially zero.

Another possibility for buyers' disproportionate rejection of exploding offers may be other regarding preferences (i.e., Fehr and Schmidt, 1999; Bolton and Ockenfels, 2000; Charness and Rabin, 2002). To the extent these preferences are correlated with a propensity to engage in system 1 responses (c.f., Rand et al., 2012; Krajbich et al., 2015), the correlation of low CRT correct answers and exploding offer suboptimal rejection may provide evidence of such preferences. However, free-recall mistakes are also correlated with lower CRT scores. While this correlation is, in general, not at odds with the fast vs. slow story—it is certainly a quick, intuitive answer to pick the highest valued or lower priced item negating the surplus—it does not disproportionately hurt the higher priced seller (see Appendix Table A.11). In addition, at the end of the low-buyer-heterogeneity treatments provided through surveys the reasoning behind their decisions in the experiment. None of the sixty-four buyers mentioned any type of decision making of this sort (see Appendix Table A.13), subjects appear focused on price, value, and surplus, rather than seller conduct.

Result 4 *The differential rate of suboptimal play in rejecting and accepting exploding offers by human buyers alters the pricing decision for sellers. While (35E, 35E) is the equilibrium* *with optimal buyers, (25E, 25E) is the equilibrium after accounting for human-buyer behavior. Sellers express a significant aversion to the use of exploding offers against human buyers, which cannot be explained by payoff differences.*

Table 7 provides payoff tables for sellers in both treatments under each valuation distribution. The two computer-buyer tables are identical to theory. The human buyer tables are based off the empirically observed play of subjects; they are based off the predicted rejection probabilities from equation (5) omitting period-effects. In all four cases, the equilibrium involves both sellers playing an exploding offer. The computer treatment has the equilibrium of buyers offering price 35 with an exploding offer. The low-buyer-heterogeneity HB treatment has an equilibrium of buyers offering price 25 with an exploding offer. The high-buyer heterogeneity distribution HB treatment has both equilibria.[29]

As Result 1 demonstrates, sellers do not exclusively play exploding offers even against computers. Given the observed play of human sellers and buyers, how should a single seller best respond? Appendix Figure A.1 provides the expected payoffs for seller strategies in each treatment and distribution using the regression model in (5) to model human behavior and the period-by-period empirically observed seller play. In all four cases, a strategy utilizing an exploding offer was generally most profitable, though the difference in profit between exploding offers and free-recall offers was reduced with human buyers. In the CB treatment, 35E was most profitable for sellers. In the HB treatment, 25E was most profitable for sellers.

Result 3 provides intuition on why exploding offers are still the most profitable against human buyers. While humans over-reject exploding offers, the rate of rejection decreases as the cost of the rejection increases. By reducing the price from 35 to 25, sellers increase the cost of rejection by 10. According to the regressions in

[29] In this treatment, the quantal response equilibrium model (McKelvey and Palfrey, 1995) selects (25E, 25E) as the more crucial equilibrium because the limit of two buyers with decreasingly noisy play converges to (25E, 25E).

Table 7: Payoff Matrices for Seller Strategies, Given Theoretically Optimal Play and Empirical Play of Human Buyers

	High-buyer-heterogeneity Distribution											
	Computer Buyer						Human Buyer					
	25E	30E	35E	25F	30F	35F	25E	30E	35E	25F	30F	35F
25E	11.33	12.50	12.50	13.67	15.04	15.63	**11.80**	13.22	14.97	12.74	14.66	15.68
30E	11.72	13.59	15.00	12.89	16.41	18.05	11.61	13.24	14.97	12.45	15.02	17.39
35E	13.67	13.67	**15.86**	14.22	15.04	19.14	10.99	13.25	**15.35**	11.60	14.07	17.19
25F	9.38	11.91	12.50	11.72	14.45	15.63	10.77	13.01	14.91	11.72	14.45	15.63
30F	9.61	11.25	14.30	10.78	14.06	17.34	9.94	12.28	14.92	10.78	14.06	17.34
35F	10.39	11.21	13.13	10.94	12.58	16.41	10.32	11.77	14.57	10.94	12.58	16.41
	Low-buyer-heterogeneity Distribution											
	Computer Buyer						Human Buyer					
	25E	30E	35E	25F	30F	35F	25E	30E	35E	25F	30F	35F
25E	11.50	11.50	15.50	13.50	14.94	15.25	**11.86**	14.18	16.22	12.14	14.60	15.80
30E	13.80	13.80	13.80	14.48	16.20	17.93	10.76	13.39	16.05	11.11	14.22	17.20
35E	9.10	16.10	**16.10**	9.45	16.89	18.90	9.76	12.41	15.55	10.73	12.55	16.20
25F	10.00	11.44	15.75	12.00	14.88	15.50	11.72	14.46	15.92	12.00	14.88	15.50
30F	10.28	12.00	13.73	10.95	14.40	17.85	10.60	13.57	16.69	10.95	14.40	17.85
35F	11.55	11.99	14.00	11.90	12.78	16.80	10.93	12.63	16.15	11.90	12.78	16.80

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.

table 5, this should reduce the probability of rejection by 15-20 points.

In all cases, sellers could maximize profits by using exploding offers, but a significant portion played free-recall offers. Given that expected payoffs for seller strategies varied across treatment, it is not clear whether sellers actually displayed an additional reluctance to play exploding offers against humans. Table 8 provides a conditional logit model of seller strategy choice in each of the four cases.

$$\text{logit}\left(\Pr\left(\text{action}_{ijt}=1\left|\text{lpayoff}_{ijt-1},\text{exploding}_{ijt}\right.\right)\right)=\beta_1\times\text{lpayoff}_{ijt-1}+\beta_2\times\text{exploding}_{ijt}. \tag{7}$$

where $\text{action}_{ijt}=1$ represents the choice of strategy $j\in\{25E,30E,35E,25F,30F,35F\}$ at period t, lpayoff_{ij} is the corresponding lagged expected payoff for strategy j (calculated based on sellers' empirical choices and buyers' optimal or empirical play in the previous period), and exploding_{ijt} is equal to one if the current strategy action involves using an exploding offer.

The results show that, compared with their computer buyer counterparts,[30] sellers had a significant, non-payoff-based reluctance to play exploding offers with human buyers. The reluctance is considerably stronger under the high-buyer-heterogeneity than low-buyer-heterogeneity distribution.[31] The results also indicate that sellers are more sensitive to variations of expected payoffs due to the other competitors or buyers in the human buyer treatment.

The types of sellers that refrain from exploding offer use are explained in the next result.

Result 5 *Neither the propensity toward "quick" (vs. "slow") thinking, risk preferences, nor gender appears to explain seller differences in exploding offer use with computer buyers. Only gender explains differences in exploding offer use with human buyers. On exit surveys,*

[30]Note that we use the same valuation draws across different sessions within high and low buyer-heterogeneity treatments.

[31]Using the quantal response equilibrium model (McKelvey and Palfrey, 1995) to represent seller decisions as noisy best responses and including a term for exploding offer aversion, yields similar results.

Table 8: Conditional Logit of Seller Action (i.e., 25E, 30E, 35E, 25F, 30F, 35F) on Lagged Payoff of Action and Type of Offer with Human-Buyer Treatment Interactions.

	Both valuations		High-buyer heterogeneity		Low-buyer heterogeneity	
lagged payoff of strategy	0.475***	0.485***	0.475***	0.619***	0.475***	0.403***
	(0.0537)	(0.0952)	(0.0790)	(0.167)	(0.0735)	(0.120)
human × lagged payoff	0.715***	0.983***	0.378**	1.587***	1.040***	1.165***
	(0.148)	(0.184)	(0.181)	(0.325)	(0.215)	(0.251)
exploding offer	-	-0.0247	-	-0.400	-	0.150
		(0.175)		(0.358)		(0.199)
human × exploding offer	-	-0.663**	-	-1.485***	-	-0.457*
		(0.261)		(0.524)		(0.276)
subject clusters	144	144	56	56	88	88
observations	16416	16416	6384	6384	10032	10032
log likelihood	-4525.996	-4481.975	-1768.664	-1705.540	-2743.742	-2736.308

a few sellers mention the use of free-recall offers as part of his/her general strategy, usually justified out of a concern for buyers. Such sellers are far less likely to use exploding offers and more likely to be women.

As explained in Result 3, subjects were classified into different levels of risk-preference and propensity towards intuitive vs. deliberative thinking based on their responses to a survey given at the end of each experimental session. Table 9 provides the results of a logit regression of exploding offer use, with risk-preference measurement, CRT scores, gender and valuation interacted on a human-buyer dummy. Surprisingly, both risk measures and CRT scores provide little explanatory power for seller exploding offer use with human or computer buyers. This finding contrasts greatly with their explanatory power for buyer behavior. Even though there is no correlation with risk,[32] gender effects are quite substantial. Male sellers are 17 probability points more likely to use exploding offers than female sellers with human buyers. There is no difference with computer buyers. Confirming previous results, sellers are much less likely to use exploding offers on human rather than computer buyers.

[32]Greater risk-aversion is normally associated with female subjects (see Croson and Gneezy, 2009).

Table 9: Logistic Regression of Exploding Offer Use on Human Buyer Treatment with Valuation, Gender, and Survey Data Interactions. Average marginal effects at 0 and 1 of "human buyer" dummy shown. For coefficient estimates see Appendix Table A.12.

	Buyer is human (human buyer=1)	Buyer is computer (human buyer=0)
above median risk tolerance	-0.024 (0.081)	0.004 (0.052)
correct CRT questions	-0.020 (0.035)	0.028 (0.022)
male	0.166** (0.083)	0.009 (0.051)
human buyer	-0.230*** (0.044)	
valuation dummy	Y	
period dummies	Y	
random effects	Y	
observations	2880	
subject clusters	144	
log likelihood	-1731.09	

In the latter sessions, under the low-buyer-heterogeneity distribution, subjects answered an open-ended question asking them to explain the reasons behind their decisions in the experiment. Ten out of thirty-two sellers expressed the use of free-recall offers as part of their general strategies. Most justified this choice with a concern on how exploding offers would affect human buyers (see Appendix Table A.14). This expression is quite predictive: these 10 subjects used exploding offers in 24.5% of their offers; the others used exploding offers in 50% of their offers ($p < 0.01$, t-test and Mann-Whitney-Wilcoxon). Interestingly, there is no significant difference in the average price offered (26.8 vs. 26.6) between these two groups, suggesting this tendency is unrelated to the desire to transfer income to buyers.

The survey response may also explain the gender differences in exploding offer use. Female subjects are more likely to indicate the use of free-recall offers in their overall strategy. Fifty-five percent of female subjects (6 of 11) and 19% of male subjects (4 of 21) indicated this concern ($p = 0.056$ on Fisher exact test). Differences

Table 10: Sellers' and Buyers' Welfare Analysis

Panel A: high-buyer-heterogeneity distribution							
buyer strategy	seller strategy	total payoff		buyer payoff		seller payoff	
		Mean	Std. Err.	Mean	Std. Err.	Mean	Std. Err.
best response	35E	27.884	0.139	12.844	0.078	15.039	0.070
best response	actual-computer	28.165	0.126	15.224	0.105	12.941	0.106
actual	25E	27.932	0.160	17.162	0.113	10.770	0.055
actual	actual-human	27.903	0.160	16.429	0.128	11.473	0.081
best response	actual-human	28.436	0.146	16.830	0.117	11.606	0.082
Panel B: low-buyer-heterogeneity distribution							
buyer strategy	seller strategy	total payoff		buyer payoff		seller payoff	
		Mean	Std. Err.	Mean	Std. Err.	Mean	Std. Err.
best response	35E	26.091	0.089	10.527	0.048	15.564	0.050
best response	actual-computer	26.510	0.078	13.464	0.076	13.046	0.074
actual	25E	26.093	0.099	14.853	0.070	11.240	0.038
actual	actual-human	26.067	0.099	14.163	0.084	11.905	0.055
best response	actual-human	26.901	0.085	14.846	0.073	12.055	0.060

on exploding offer use among the subjects who indicate this concern (23% exploding offer use for female subjects vs. 30% for men) and those who do not (51% female vs. 50% for male) do not exhibit gender effects.

Result 6 *Sellers' deviation from equilibrium strategy leads to lower earnings in the computer-buyer treatment and higher earnings in the human-buyer treatment. Holding sellers' strategies constant, buyer deviations from best response cause slight losses to both buyers and sellers; their main effect on earnings is through altering seller play.*

Table 10 provides estimates of buyer, seller, and total payoffs for a variety of buyer and seller strategies across all four treatments. Because buyers' values of items are taken from a random distribution, payoff estimates vary across simulations even when strategies are exact. The first row of the welfare calculations in each panel gives the theoretical earnings for each treatment. In both valuations, sellers should make considerably more than buyers and the disparity should be greater under the low-buyer-heterogeneity distribution. However, sellers do not play 35E (their equilibrium strategy) exclusively, even against computer buyers. This results in a loss of \$2.00-\$2.50 between seller earnings at the theoretical optimum and seller earnings in the computer-buyer treatments. Total surplus is improved by \$0.40 as sellers use some amount of free-recall offers against computer buyers.

Human buyers do not respond to sellers' offers the same way computer buyers do. As payoff table 7 shows, accounting for these new strategies makes (25E, 25E) the new equilibrium for sellers instead of (35E, 35E). Sellers would do very poorly if they played this equilibrium solely against human buyers. They would make \$10.77 and \$11.24 in the low and high buyer heterogeneity distributions of the HB treatment, respectively. However, sellers do not play this strategy exclusively. Here they increase their payoffs by roughly \$0.70 by charging higher prices or using free recall offers. This deviation from the equilibrium strategy could either be due to bounded rationality or collusion. Given the relatively poor seller earnings with computer buyers (they could collectively charge high prices as the sole equilibrium strategy, but do not), bounded rationality appears to be the more likely explanation.

Buyers do not behave as optimal play would dictate. If we hold seller strategies as constant, buyers could improve their earnings by \$0.40-\$0.70 by following optimal strategies. Sellers are not particularly hurt directly by these deviations; they only cost sellers about \$0.15 of their earnings. Of course seller strategies are not constant, the way buyers respond to exploding offers greatly changes the strategic structure of the game. This has a second-order effect of moving the seller equilibrium from (35E, 35E) to (25E, 25E) which costs sellers over \$4.00 in earnings in theory (or \$1.00-\$1.50 in practice).

6 Discussion

Our experiment identifies two pronounced departures from equilibrium. First, buyers more often deviate from optimal strategies in rejecting rather than accepting exploding offers. Second, sellers demonstrate a significant reluctance to use exploding offers on human buyers, unrelated to their best response toward suboptimal, human-buyer play. The total effect is that sellers play the equilibrium strategy five times more often against computers than human buyers.

Buyer behavior, alone, eliminates the advantage of exploding offers in this market. From table 7, the deviations of human buyers make (25E, 25E) an equilibrium. For comparison, had all exploding-offer strategies been eliminated from the game, (25F, 25F) would be the equilibrium. The payoffs for both outcomes are nearly identical.

The behavior of buyers who score lowest on the CRT is crucial for this result. These type of subjects are considered to be the most prone to system 1 (i.e., "fast," "intuitive," "lazy") thinking (Kahneman, 2011). They are most prone to reject exploding offers suboptimally, but are no more likely to suboptimally accept them. A possible explanation is that the fast, intuitive response to search deterrence is to run from the offer and search further. Only after deliberation does one understand the value in accepting such offers.

There are important implications for search-deterrence in the field. Management can select employees, but they have less control over potential customers. While a firm could selectively hire or train their sale staffs to use search deterrence without compunction,[33] some potential customers will be prone to fast thinking. In such case, our results suggest these buyers will exhibit a differentially negative response to search-deterring offers. Then, depending on the market-payoff structure, potential gains in profit from the search deterrence strategy would be reduced.

Our results show that it only takes a little bit of buyer suboptimality, directed with a slight bias, to radically alter the seller incentives in this market. We find a buyer suboptimal play rate of 0.198 and 0.103 in the low- and high-buyer-heterogeneity distribution. At a minimum, to create an equilibrium with sellers offering the lowest price (in an exploding offer) would only take a suboptimal play rate of 0.133 and 0.077, respectively, holding the rate of suboptimal rejections-to-

[33]Going further, the firm could implement a computer assisted decision management system, which could have the added benefit of removing commitment problems associated with search deterrence.

acceptances constant. Even if one is skeptical about the exact values of these rates across different field markets, the underlying game is still fragile and small differentials in suboptimal play are sufficient to alter a seller's strategy and undermine the theoretical advantages of search deterrence.

Using our results to answer the motivation question of this paper—why we do not see more search deterrence in the field—requires nuance. If sellers are already capable of using search deterrence in markets structured similar to ours, it may still be optimal to do so, even given buyer suboptimal play. Perhaps if sellers can learn signals of buyers' valuations (Xiong and Chen, 2014), they may use such offer conditionally—which makes such offers more "casual" and less traceable. However, if it is costly to implement or maintain such pricing strategies (either through additional technology or labor investment), firms may have little incentive to do so.

Extending this idea further, there are consumer markets where search deterrence strategies are not used, presumably because they would be too costly to implement. Rather than say that the findings of this experiment have no bearing on such markets, we come to a slightly different conclusion. As technology to track customers becomes more powerful and affordable, producers in these markets may have the opportunity to invest in technology that would allow them to implement search deterrence on each customer.[34] Our results provide a cautionary note to these producers; the increased profits from adopting this technology may not be as great as theory predicts.

Admittedly, one should be cautious in making field predictions directly from the results of laboratory data. In this case, however, the laboratory has an advantage over the field in that it can ensure all exploding offers are credible to buyers and costless for sellers to use. Among other things, field studies would need to estimate

[34]See patents numbers US 8543470 B2, US 7418405 B1 and US 20100023407 A1 (Utter et al., 2008; Grady and Orttung, 2010, 2013) for examples of new technology designed to offer customers a product for a limited time in electronic markets.

consumers' beliefs about the credibility of each search-deterring offer and ascertain the costs producers incur in using such offers before determining whether consumer or producer behavior was optimal. Given the numerous underlying assumptions required to estimate similar results to this experiment in an empirical field study, experimental data like this paper may come closest to answering our research question.

Finally, our experiment provided, as intended, a sterile environment for sellers to make exploding offers to buyers with minimal emotional connotation. In a field setting, these psychological factors may be more salient. For example, there might be psychological effects, when using such offers with less aggressive methods, credible justifications or repeated interactions outside the markets.[35] While the buyer behavior in our experiment appears to be unrelated to these psychological factors, we are not certain how these additional factors might alter our conclusions. We leave this as an intriguing topic for future research.

References

Mark Armstrong and Jidong Zhou. Search deterrence. *The Review of Economic Studies*, 83(1):26–57, 2016.

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.

[35] Armstrong and Zhou (2016) note: "It may be important for the seller to be able to rationalize its sales policy, to give the policy greater apparent credibility. For example Bone (2006) (p. 71) writes that to justify the buy-now discount customers were told "the company had so many appointments that it was difficult for our salespeople to cover them all ... if we went back to everyone twice we wouldn't see nearly as many people and would generate a lot less business."

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

Gary E. Bolton and Axel Ockenfels. ERC: A theory of equity, reciprocity, and competition. *The American Economic Review*, 90(1):166–193, 2000.

John Bone. *The hard sell : an ethnographic study of the direct selling industry*. Ashgate, Aldershot, England; Burlington, VT, 2006.

A Colin Cameron, Jonah B Gelbach, and Douglas L Miller. Bootstrap-based improvements for inference with clustered errors. *The Review of Economics and Statistics*, 90(3):414–427, 2008.

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

Gary Charness and Matthew Rabin. Understanding social preferences with simple tests. *The Quarterly Journal of Economics*, 117(3):817–869, 2002.

Robert B. Cialdini. *Influence : science and practice*. Allyn & Bacon, Boston, 4. ed. edition, 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.

Rachel Croson and Uri Gneezy. Gender differences in preferences. *Journal of Economic Literature*, 47(2):448–474, 2009.

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.

Ernst Fehr and Klaus M. Schmidt. A theory of fairness, competition, and cooperation. *The Quarterly Journal of Economics*, 114(3):817–868, 1999.

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

Guillaume R Fréchette. Session-effects in the laboratory. *Experimental Economics*, 15 (3):485–498, 2012.

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

P. Grady and M. Orttung. System and method for availability-based limited-time offerings and transactions, January 28 2010. US Patent App. 12/573,844.

P. Grady and M. Orttung. System and method for targeting limited-time offer based on likelihood of acceptance and selecting transmission media based on customer interest, September 24 2013. US Patent 8,543,470.

Ben Greiner. Subject pool recruitment procedures: organizing experiments with orsee. *Journal of the Economic Science Association*, 1(1):114–125, 2015.

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.

Daniel Kahneman. *Thinking, fast and slow*. Macmillan, 2011.

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

Ian Krajbich, Björn Bartling, Todd Hare, and Ernst Fehr. Rethinking fast and slow based on a critique of reaction-time reverse inference. *Nature communications*, 6, 2015.

Nelson Lau, Yakov Bart, J. Neil Bearden, and Ilia Tsetlin. Exploding offers can blow up in more than one way. *Decision Analysis*, 11(3):171–188, 2014.

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.

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.

David G Rand, Joshua D Greene, and Martin A Nowak. Spontaneous giving and calculated greed. *Nature*, 489(7416):427–430, 2012.

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.

B. Utter, D. Welzel, D. Irvine, D. Benefield, D. Liu, and J. Shaw. Interactive time-limited merchandising program and method for improved online cross-selling, August 26 2008. US Patent 7,418,405.

Hui Xiong and Ying-Ju Chen. Product line design with seller-induced learning. *Management Science*, 60(3):784–795, 2014.

Appendices: Not Intended for Publication

A Additional Tables and Figures

Table A.1: Exploding Offer Rate, Average Price, and Rate of Equilibrium Play by Seller Cohort

Cohort Number	Session Number	Buyer Type	Buyer Valuation Heterogeneity	Exploding Offer Rate	Average Price	Rate of Equilibrium Play
1	1	human	high	0.488	27.375	0.088
2	2	computer	high	0.638	31.406	0.456
3	2	computer	high	0.725	30.938	0.394
4	3	human	high	0.675	26.438	0.019
5	4	human	high	0.475	27.031	0.031
6	5	computer	high	0.763	29.750	0.213
7	5	computer	high	0.594	29.344	0.231
8	6	human	low	0.288	27.531	0.069
9	7	computer	low	0.625	28.844	0.125
10	7	computer	low	0.569	29.625	0.263
11	7	computer	low	0.606	29.688	0.125
12	8	computer	low	0.700	29.563	0.213
13	8	computer	low	0.731	28.250	0.063
14	9	human	low	0.481	27.250	0.094
15	10	human	low	0.519	26.000	0.006
16	11	human	low	0.394	25.938	0.013
17	12	computer	low	0.650	30.406	0.288
18	12	computer	low	0.694	29.594	0.119

Table A.2: Mistake Rate by Buyer Cohort

			Suboptimal Play Rate			
			All Offers		Exploding Offers Only	
Cohort Number	Session Number	Buyer Valuation Heterogeneity	Against Free-Recall Offers	Against Exploding Offers	When Acceptance is Optimal	When Rejection is Optimal
1	1	high	0.020	0.093	0.035	0.136
4	3	high	0.031	0.115	0.036	0.168
5	4	high	0.069	0.115	0.044	0.162
8	6	low	0.080	0.243	0.121	0.349
14	9	low	0.071	0.221	0.138	0.281
15	10	low	0.055	0.137	0.088	0.166
16	11	low	0.034	0.180	0.109	0.230

Table A.3: Regression of Dependent Variables on Main Treatment Effect with Clustering at the Seller Cohort Level.

	Exploding Offer Rate		Average Price		Equilibrium Play Rate	
human buyer treatment	-0.189*** (0.046)	-0.193*** (0.042)	-2.97*** (0.355)	-3.013*** (0.343)	-0.181*** (0.038)	-0.187*** (0.037)
low buyer heterogeneity	-	-0.066 (0.040)	-	-0.667* (0.358)	-	0.092** (0.043)
period dummies	N	Y[a]	N	Y[a]	N	Y[a]
observations	2880	2880	2880	2880	2880	2880
cohort clusters	18	18	18	18	18	18
r-squared	0.035	0.055	0.142	0.154	0.059	0.080

a. Alternate regression specifications replacing period fixed effects with a single continuous period variable or omitting the variable altogether do not change coefficients in any meaningful way.

Table A.4: Regression of Dependent Variables on Main Treatment Effect with Clustering at the Seller Cohort Level. Wild bootstrap clustering used. P-values given.

	Exploding Offer Rate		Average Price		Equilibrium Play Rate	
human buyer treatment	-0.189*** $p = 0.004$	-0.193*** $p = 0.002$	-2.97*** $p = 0.002$	-3.013*** $p = 0.002$	-0.181*** $p = 0.004$	-0.187*** $p = 0.002$
low buyer heterogeneity	-	-0.066* $p = 0.1$	-	-0.667* $p = 0.056$	-	-0.092** $p = 0.044$
period dummies	N	Y[a]	N	Y[a]	N	Y[a]
observations	2880	2880	2880	2880	2880	2880
cohort clusters	18	18	18	18	18	18
r-squared	0.035	0.055	0.142	0.154	0.059	0.080

a. Alternate regression specifications replacing period fixed effects with a single continuous period variable or omitting the variable altogether do not change coefficients in any meaningful way.

Table A.5: Regression of Mistake Rate on Offer Type with Clustering at the Buyer Cohort Level.

	Suboptimal Play Rate			
seller uses exploding offer	-0.095*** (0.016)	-0.103*** (0.146)	-0.095*** $p = 0.008$	-0.103*** $p = 0.009$
low buyer heterogeneity	-	0.049** (0.017)	-	0.049*** $p \approx 0.000$
type of clustering	standard	standard	wild bootstrap	wild bootstrap
period dummies	N	Y[a]	N	Y[a]
observations	13,431	13,431	13,431	13,431
cohort clusters	18	18	18	18
r-squared	0.025	0.044	0.026	0.044

a. Alternate regression specifications replacing period fixed effects with a single continuous period variable or omitting the variable altogether do not change coefficients in any meaningful way.

Table A.6: Regression of Mistake Rate on Offer Type with Clustering at the Buyer-cohort Level.

	Suboptimal Play Rate			
acceptance is optimal	0.123*** (0.012)	0.122*** (0.015)	0.123*** $p = 0.006$	0.122** $p = 0.011$
low buyer heterogeneity	-	0.076** (0.020)	-	0.076* $p = 0.088$
type of clustering	standard	standard	wild bootstrap	wild bootstrap
period dummies	N	Y[a]	N	Y[a]
observations	6372	6372	6372	6372
cohort clusters	18	18	18	18
r-squared	0.029	0.054	0.029	0.054

a. Alternate regression specifications replacing period fixed effects with a single continuous period variable or omitting the variable altogether do not change coefficients in any meaningful way.

Table A.7: Linear Time Trends of Dependent Variables.

	Exploding Offer Rate	Average Price	Equilibrium Play Rate
human buyer treatment × period	-0.000 (0.003)	-0.138*** (0.021)	-0.007*** (0.021)
computer buyer treatment × period	0.015*** (0.002)	0.107*** (0.021)	0.008*** (0.002)
constant	0.497*** (0.249)	28.488*** (0.208)	0.134*** (0.018)
observations	2880	2880	2880
subject clusters	144	144	144
r-squared	0.043	0.134	0.057

Table A.8: Logistic Regression of Buyer Suboptimal Play on Cost of Suboptimality and Type of Optimal Play. Coefficients shown. For an alternate version of this table, with marginal effects see Table 5.

	Both valuations		High buyer-heterogeneity		Low buyer-heterogeneity	
cost of suboptimality	-0.214*** (0.016)	-0.214*** (0.018)	-0.250*** (0.022)	-0.253*** (0.027)	-0.184*** (0.021)	-0.184*** (0.022)
acceptance is optimal	-	0.822*** (0.275)	-	1.085** (0.517)	-	0.699** (0.328)
valuation dummy	Y	Y	N	N	N	N
period dummies	Y	Y	Y	Y	Y	Y
random effects	Y	Y	Y	Y	Y	Y
observations	6372	6372	3144	3144	3228	3228
subject clusters	112	112	48	48	64	64
log likelihood	-1933.892	-1899.576	-732.985	-714.777	-1189.653	-1172.800

Table A.9: Logistic Regression of Buyer Suboptimal Play on Cost of Suboptimality and Type of Optimal Play with Gender and Survey Data Interactions. Interactions on "acceptance [of exploding offer] is optimal" dummy variable. For an alternate version of this table, with average marginal effects evaluated at 0 and 1 of "acceptance is optimal" dummy shown see Table 6. Also see A.10 for average marginal effects at (below median risk aversion=1 and correct CRT questions=3) and 0 and 1 of acceptance is optimal dummy.

	Suboptimal play
cost of suboptimality	-0.175***
	(0.602)
acceptance is optimal	0.970
	(0.601)
male	-0.103
	(0.406)
above median risk tolerance	0.701*
	(0.391)
correct CRT questions	-0.112
	(0.167)
acceptance is optimal×	
cost of suboptimality	-0.058**
	(0.0256)
male	0.637
	(0.531)
above median risk tolerance	1.041**
	(0.520)
correct CRT questions	-0.293
	(0.238)
valuation dummy	Y
period dummies	Y
random effects	Y
observations	6372
subject clusters	112
log likelihood	-1871.563

Table A.10: Logistic Regression of Buyer Suboptimal Play on Cost of Suboptimality and Type of Optimal Play with Gender and Survey Data Interactions. Interactions on "acceptance [of exploding offer] is optimal" dummy variable. Average marginal effects at (above median risk tolerance=0 and correct CRT questions=3) and 0 and 1 of acceptance is optimal dummy shown.

	When acceptance is optimal (acceptance is optimal=1)	When rejection is optimal (acceptance is optimal=0)
cost of suboptimality	-0.013*** (0.003)	-0.013*** (0.004)
above median risk tolerance	0.022 (0.019)	-0.040* (0.023)
correct CRT questions	-0.023*** (0.005)	-0.008 (0.011)
male	-0.004 (0.019)	-0.007** (0.029)
acceptance is optimal	-0.018 (0.036)	
treatment dummy	Y	
period dummies	Y	
random effects	Y	
observations	6372	
subject clusters	112	
log likelihood	-1871.563	

Table A.11: Logistic Regression of Buyer Suboptimal Play on Cost of Suboptimality and Type of Optimal Play with Gender and Survey Data Interactions, Free-Recall Offers. Interactions on "choice of higher-priced item [of free recall offer] is optimal" dummy variable.

	Suboptimal play
cost of suboptimality	-0.058***
	(0.009)
higher-priced is optimal	0.463
	(0.397)
male	-0.132
	(0.489)
above median risk tolerance	0.249
	(0.452)
correct CRT questions	-0.756***
	(0.226)
higher-priced is optimal×	
cost of suboptimality	0.008
	(0.017)
male	0.096
	(0.363)
above median risk tolerance	-0.572*
	(0.345)
correct CRT questions	0.454
	(0.176)
valuation dummy	Y
period dummies	Y
random effects	Y
observations	5920
subject clusters	112
log likelihood	-938.8467

Table A.12: Logistic Regression of Exploding Offer Use on Human Buyer Treatment with Valuation, Gender, and Survey Data Interactions. The marginal effects shown in Table 9 are derived from the regression represented in this table.

	Suboptimal play
human buyer	-1.10***
	(0.337)
male	-0.043
	(0.243)
above median risk tolerance	0.021
	(0.247)
correct CRT questions	0.133
	(0.105)
human×	
male	0.654
	(0.435)
above median risk tolerance	-0.124
	(0.428)
correct CRT questions	-0.216
	(0.186)
valuation dummy	Y
period dummies	Y
random effects	Y
observations	2880
subject clusters	144
log likelihood	-1731.09

Table A.13: Open Ended Survey Response Answers, Buyers. Data only available for Low-Buyer Heterogeneity Distribution Subjects (N = 64).

Type of Response	Example Statement(s)	Number of Subjects	Percent
Heuristic	"If I made a profit of over 30 pts then I automatically accepted the offer..."	18	0.281
Best-Response	"As the buyer, When given Offer A, I looked to see what the offer was and compared the price and chose a value from them, keeping in mind the options and probability. But for the Offer B, i looked and both options first and then calculated which..."	17	0.266
Surplus Maximization	"Y-X Which ever made the highest net profit for me."	17	0.266
Always Search	"i always searched the marked [sic]..."	17	0.266
Take Any Positive Surplus	"...and i had enough value to buy it, i would buy it because i knew forsure [sic] i would get points..."	3	0.047
Purchase Higher Value Item	"It was based on how high the value number was, if the value was below half i clicked search no matter what"	2	0.031
Miscellaneous	"I picked the item that had a low value"; "I guess"; random letters, etc.	4	0.062
Total		64	1.000

Table A.14: Open Ended Survey Response Answers, Sellers. Data only available for Low-Buyer Heterogeneity Distribution Subjects (N = 32).

Type of Response	Example Statement(s)	Number of Subjects	Percent
Free-Recall Use Without Explanation	"went with option b at lowest price"	2	0.062
Free-Recall Use With Explanation	"I TRIED TO GIVE PEOPLE THE OPTION TO ALWAYS COME BACK TO BY ITEM TO PURCHASE AND RARELY EVER SOLD FOR THE MAX PRICE"	8	0.250
Trial-And-Error	"The first several periods I tested the different offers, and found the low priced B or A offers worked best to achieve the most points. The B offer worked well regardless of the other sellers offer and the A offer only worked well when the other chose B"	16	0.500
Miscellaneous	"What felt right."; "What worked in the past."; "Least random"	6	0.188
Total		32	1.000

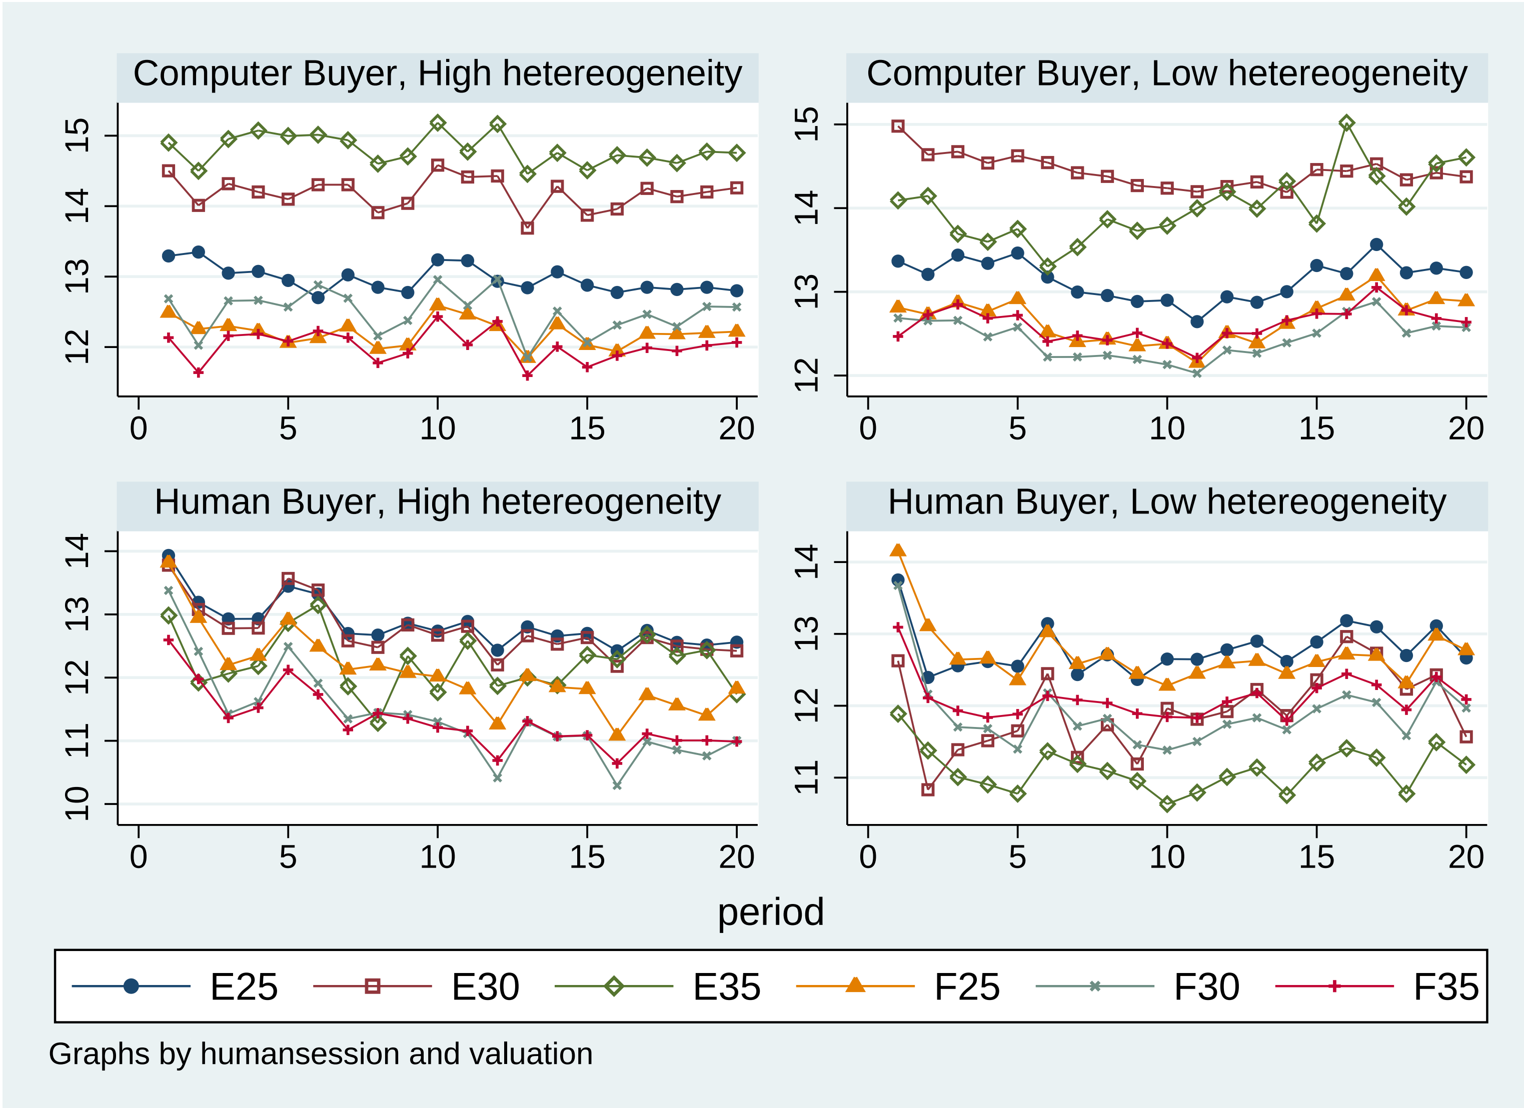

Figure A.1: Payoff Vector for: CB Treatment (high heterogeneity valuation) (left-top), CB Treatment (low heterogeneity valuation) (left-bottom), HB Treatment (high heterogeneity valuation) (right-top), HB Treatment (low heterogeneity valuation) (right-bottom)