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

November 21, 2017

Abstract

In consumer goods markets, theory shows that it is generally profitable for sellers to use search-deterrence strategies to alter buyer search. These results rely on agents' reacting solely to the economic content of these pressure tactics, ignoring any behaviorally based responses search deterrence may evoke. To test the validity of this assumption, this paper examines an experimental market where profit-maximizing strategy dictates that sellers should exercise one form of search deterrence, exploding offers. Sellers demonstrate a reluctance to use such offers against human buyers, but they are less reluctant to use them against computerized buyers. Human buyers are three times more likely to deviate from optimal strategy by rejecting rather than accepting these offers. Survey responses are consistent with other-regarding-preference-based reasons for sellers' actions but not buyers'. Taken together, these results suggest the benefits of tactics that rely on pressuring decision-makers may be more nebulous than previously thought.

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, Daniel Fragiadakis, Ian Krajbich, Silvana Krasteva, Vikram Manjunath, 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

Many markets are modeled as dynamic. For instance, buyers search and learn new information about their valuations over time before finalizing purchase decisions. Recent theory developments (Courty and Hao, 2000; Nocke et al., 2011; Armstrong and Zhou, 2016) show sellers often earn more than a static price discrimination theory would predict. They do so by using dynamic contracts to alter buyers' information structure. Implicit in the execution of these contracts is the pressure sell. Sellers prod buyers to make a decision immediately rather than at their leisure. While perfectly rational agents will respond to this pressure in a manner consistent with theory, other agents may respond to increased pressure—or even the thought of increased pressure on someone else—with a suboptimal, behaviorally based response.[1]

This paper examines this question in a laboratory experiment designed to implement a simplified consumer search model (Armstrong and Zhou, 2016): two sellers simultaneously choose one of three prices and make either an exploding or non-exploding offer.[2] 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 doing so, they receive the seller's offer. The buyers 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 seller behavior is conditional on perceived buyer response, two treatments are used to isolate seller behavior. In one, sellers knowingly interact with

[1]Beginning with Güth et al. (1982), there is a substantial body of literature showing that actions that violate some form of fairness are met with a negative response, even in a single-shot setting (Cooper and Kagel, 2016). We interpret a plausible extension of this research that buyers would negatively reciprocate against a pressure sell, and sellers would be reluctant to make such sells. There is empirical precedent for such thinking (see Kahneman et al., 1986). Similarly, a large body of literature is generally consistent with the idea that pressure exacerbates optimal decision-making (for example, see Shah et al., 2012).

[2]The exploding offer is an offer that is 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. In the experiment, we assume the offer expires instantly after a buyer's visit.

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.[3] Buyers will reject or accept exploding offers, unbiasedly, based on expected value and profit maximization.

In contrast to the equilibrium prediction, human buyers are three 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's 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 two-thirds as often against human buyers as they do against computer buyers. Overall, sellers play the equilibrium strategy five times more often against computer buyers than human buyers.

Though our design features only exploding offers, we suspect the basic flavor of these results—that buyers and sellers respond more negatively to pressure sales tactics than standard theory predicts—applies to all forms of search deterrence, a tactic that is continually evolving and expanding in the field.[4] Originally, cases of consumer search deterrence were confined to specific markets and situations. Literature suggests use of exploding offers was common among professional salespeople (Cialdini, 2003) and health clubs offering memberships (Xiong and Chen, 2014). Buy-now discounts were noted in home improvement service contracts (Bone, 2006) and housing rentals (Robinson, 1995). Internet commerce expanded

[3]Exploding offers will cause high-valuation buyers to buy immediately from the seller and may drive away medium-valuation buyers who would have returned to buy the product if their search was not deterred. If the distribution of buyers is increasing in valuation, these tactics will be more profitable, compared to regular free-recall offers.

[4]Exploding offers are special limiting cases of other forms of search deterrence, "buy-now" discounts and deposits. If a deposit or discount is sufficiently high, a buyer will be forced to either accept the offer immediately or reject it entirely. Further, with sufficient return costs, there exist equilibria such that any pricing decision is effectively an exploding offer, because buyers will not return to sellers whose offer they have rejected (Armstrong and Zhou, Proposition 4).

its scope, in part, by greatly reducing the costs of maintaining, tracking, and notifying customers for search deterrence. Tracking “cookies” may lead to increases in prices should a consumer leave a website and return. ”Daily deals” or ”flash sale” websites now market limited-time offers to registered customers by email, text or social networks. The initial market valuations of flash sale websites and single-day revenues of the most popular daily-sale events are on the order of magnitude of billions of dollars.[5] Recent patents suggest search-deterrence technology will only improve.[6] Closely related to the expansion of search-deterrence strategies is the movement in consumer protection policies. In 2005, the European Union began prohibiting sellers from making false claims about product availability “in order to elicit an immediate decision and deprive consumers of sufficient opportunity or time to make an informed choice.”(Armstrong and Zhou, 2016)

However, there need not be one single underlying mechanism for sellers’ and buyers’ departures from equilibrium play. Our survey results find separate patterns behind buyer and seller responses. Buyers who answer the fewest questions correctly on the Cognitive Reflection Test (CRT; 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. 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. These sellers are far less likely to use exploding offers. Interestingly, not one buyer indicates a corresponding concern about the effects of exploding offers. Our interpretation is that sellers’ reluctance to use exploding offers is driven by other regarding prefer-

[5]See Rao (2016); Soper (2017); Lavin (2016). Recent trends suggest flash sale websites may be fading from popularity after initial fanfare, while scheduled daily deals are thriving. This is surprising, as there is little distinction between the two methods. We suspect the slight difference in the framing of the two types of sales may be responsible. In our conclusion, we discuss this issue more broadly as a topic for future research.

[6]See patents numbers US 8543470 B2, US 7418405 B1 and US 20100023407 A1 (Utter et al., 2008; Grady and Orttung, 2010, 2013).

ences, but buyers are intuitively (rather than deliberately) biased toward rejecting exploding offers. Having two separate triggers for suboptimal play under pressure tactics broadens the implications of our findings. It is well known that fairness and norm considerations can alter strategies used in markets (Kahneman et al., 1986); less is known about the implications of buyers' rejecting pressure sales innately, even when the sellers' intentions are benevolent. In our conclusion, we explore such implications within and outside the context of search deterrence.

2 Related Literature

Investigations into search deterrent strategies initially attracted the attention of researchers and policy makers in the labor market. Most work concerns buyers' making exploding offers to sellers, where the numbers of offers and purchases may both be extremely constrained. These unique features, not found in consumer goods markets, make 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 that 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—and with binding acceptances—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 than do traditional hires, leading to welfare losses for both sides. Tang et al. (2009) examine how employers select the duration of 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.

[7] There are also clear-cut examples of labor markets in which 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).

Armstrong and Zhou (2016) do not include these preceding features (i.e., matching quality, reciprocity after market transaction, variable exploding offer length), and neither does our paper. Their paper, the theoretical basis for ours, involves multiple firms' applying strategies of search deterrence to increase consumers' search cost and affect their search behavior in the market. The authors concede their model does not incorporate "behavioural factors," but they speculate these factors could either make search-deterrent strategies either 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. On one hand, they find, the search outcome is usually very efficient; on the other hand, subjects tend to stop earlier than the risk-neutral optimal strategy, which is consistent with risk aversion. More recently, the sequential searching behavior has been evaluated under more general experimental markets—where sellers make price offers and buyers make purchase decisions (Grether et al., 1988; Cason and Friedman, 2003). Our paper builds on this strand of literature by allowing for the possibility of search deterrence. Contrary to previous findings, our main result implies buyers generally search *longer* than the risk-neutral optimum, that is, in the opposite direction of the risk-averse prediction.

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 who 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 that has a private value for the buyer drawn from the same ex-ante value distribution: $V^i_k \in \{V_1, V_2, ..., V_K\}$ (where $i = 1, 2$ represents sellers and $k = 1, 2, ..., K$ represents K possible values) with probability $v_1 \equiv prob(V_1), v_2 \equiv prob(V_2), ..., v_K \equiv prob(V_K)$. 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.
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).

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

[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.

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 ends; 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 ends. If he rejects and the first offer was an exploding offer, no transaction occurs and the game ends. 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 ends. If there is no transaction, all players receive zero payoff. If there is a transaction, the buyer receives a payoff equal to the difference between his value and the price of the good he bought; that seller receives a payoff equal 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 buyer will have full information when visiting the second seller, whether the second seller uses an exploding offer has no effect on the buyer's strategy; he will maximize surplus regardless. Thus we need to consider only two cases: (1) the first offer is a free-recall offer, and (2) the first offer is an exploding offer.

If the first offer is a free-recall offer, visiting S^2 does not prevent the buyer from revisiting S^1; the buyer always searches.[11] After visiting both sellers, the buyer

[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

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 will decide by comparing the payoff from accepting the first offer to 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 he rejects the first offer, the buyer will accept the second offer only if $V^2 > P^2$.

3.3 Seller Strategies

Like 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 will visit first. There are three possible cases to consider: (1) both sellers use exploding offers; (2) both sellers use free-recall offers; and (3) one seller uses an exploding offer and the other seller uses a free-recall offer.

First, consider the case in which both sellers use exploding offers. Consider seller i with a price P^i, who plays with seller j with a price P^j. Two possible situations occur with equal probability:[14]

---

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 will accept the offer and gain $V_k^2 - P^2$; however, if $V_k^2 < P^2$, he will reject the offer and earn zero payoff. So, for each value k of the second good, the buyer will 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.

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

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

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

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

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

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

Since seller i receives a payoff P^i only if the buyer purchases from her, seller i's expected payoff is $P^{i*}[\frac{1}{2}\text{Prob(accept } i_1) + \frac{1}{2}\text{Prob(accept } i_2)]$.

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

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

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

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 $P^{i*}\text{Prob}(\text{accept } i_3)$.

Last, consider the case in which one seller uses an exploding offer and the other seller uses a free-recall offer. Because the second seller's offer type 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 $P^{i*}[\frac{1}{2}\text{Prob}(\text{accept } i_1)+\frac{1}{2}\text{Prob}(\text{accept } i_3)]$.[15] If seller i uses a free-recall offer while seller j uses an exploding offer, seller i's expected payoff is $P^{i*}[\frac{1}{2}\text{Prob}(\text{accept } i_3)+\frac{1}{2}\text{Prob}(\text{accept } i_2)]$.[16]

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

[16]The case in which the buyer visits seller i first is equivalent to the case in which both sellers use free-recall offers, and the case in which the buyer visits seller j first is equivalent to the case in which 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 $V^i_k \in \{V^i_1, V^i_2, ..., V^i_K\}$, probabilities $\upsilon_1, ..., \upsilon_K$, and prices $P^i \in \{P^i_1, P^i_2, ..., P^i_L\}$, 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 for the high- and low- buyer-heterogeneity distributions. The high-buyer-heterogeneity distribution offers a greater variance of buyers' values for items than does the the low-buyer-heterogeneity distribution.[17] One interpretation of this heterogeneity is that it represents the underlying competitiveness of the market depending on the similarity of the sellers' products.

The valuation structures provide different incentives for buyers across distributions. Suppose, for example, a buyer faces exploding offers from both sellers at a price of 35. Under the high-buyer-heterogeneity distribution, a buyer will accept the first offer only if his value for the first item is either 65 or 70. In contrast, under the low-buyer-heterogeneity distribution, a buyer will accept the first offer if his value for the first item is either 55, 65 or 70. Under both distributions, if the first offer is rejected, the second offer will 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 in normal form.[18] The table introduces our convention of referring to seller strategies. The price a seller offers will be followed by a letter ("E" or "F") to indicate whether the offer is an exploding or free-recall offer. Under both distributions, the unique

[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 greatly modifying the other moments. However, we were constrained by the equilibrium prediction and the need to use 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 of accepting each seller's offer; (2) a buyer rejects an offer with zero payoff.

Table 2: Two-Seller Game Table with Expected Payoffs for One Seller's Strategy Choice (Row) Given Other Seller's Strategy Choice (Column). Payoffs provided for row player only. Because the game is symmetric, the payoffs for the column player are found in the transpose cell. 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. The equilibrium payoffs are in bold.

	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

equilibrium for sellers involves an exploding offer with the highest price of 35 (points), that is, (35E, 35E).

Finally, these results rely on the assumption of risk neutrality. Risk-averse buyers should act differently in two main ways: they should more often accept exploding offers, and they should show preference for a sure surplus over a random distribution of surpluses from the second seller. This will increase the profitability of exploding offers for sellers, even if both sellers choose to use exploding offers. If sellers rather than buyers are risk averse, there is a possibility that sellers may offer lower prices in equilibrium but still use exploding offers. However, this result depends on the assumption of one buyer in the model, an assumption chosen without loss of generality in a risk-neutral framework. With 24 buyer-valuation draws instead (as in our experiment), risk-averse sellers have sufficient diversification over buyer valuation levels that (35E, 35E) is the unique equilibrium.

4 Experimental Design and Procedures

Experimental methodology is essential for research questions such as these. Because search-deterrence strategies are often offered casually and not publicly announced, "It is hard to obtain empirical evidence about this form of price discrimination."(Armstrong and Zhou, 2016, p. 26) The laboratory setting also addresses the question of credibility with search-deterrence offers. An exploding offer in the laboratory will expire with certainty. In the field, search-deterrence offers may be cheap talk. It is in the interest of firms to make statements that prod customers to buy quickly and without looking at competing offers. If a customer refuses their initial offer and return, firms will face incentives to sell to the customer rather than keeping their word on their initial search-deterring offer.

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 which type of buyers they were matched against. Each group consisted of eight sellers (for all treatments) and 16 buyers (only for the HB treatment). For the CB treatment, a session contained two or three 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 24 buying decisions from either computer buyers programmed to play optimal strategies (CB treatment) or four human buyers (HB treatment). Each human buyer played six different possible buying decisions with randomly determined item-value pairs.[19] In each

[19]Experimental economics often relies on techniques (e.g., multiple price lists, strategy method)

market, half of the buyers visited one seller first and the other half visited the other seller first.

There were 20 total periods. In 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). A seller's price and offer type were the same for all buyers that encountered the seller. Buyers would observe the prices and offer types of both sellers in the market, but they would see only 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 allowed 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.

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, because of both time constraints and buyer fatigue. Instead we use six randomly selected item-value combinations. We view six as the highest number of decisions we could ask buyers to make each round without encountering fatigue or time limit issues. Our intent on the seller side is to use the 24 buying decisions to reduce the risk associated with a single buyer-valuation draw that can alter the equilibrium of the game for risk-averse sellers (see Section 3.4).

[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 all four buyers had completed their six buying decisions, screens showed sellers the outcome of all 24 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 24 buying decisions in the market. Sellers were provided this large amount of feedback to give them the best opportunity to respond optimally 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 departed a little from this design for pragmatic reasons. First, sellers were 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 six different possibilities of this same game with identical seller strategies to provide us with

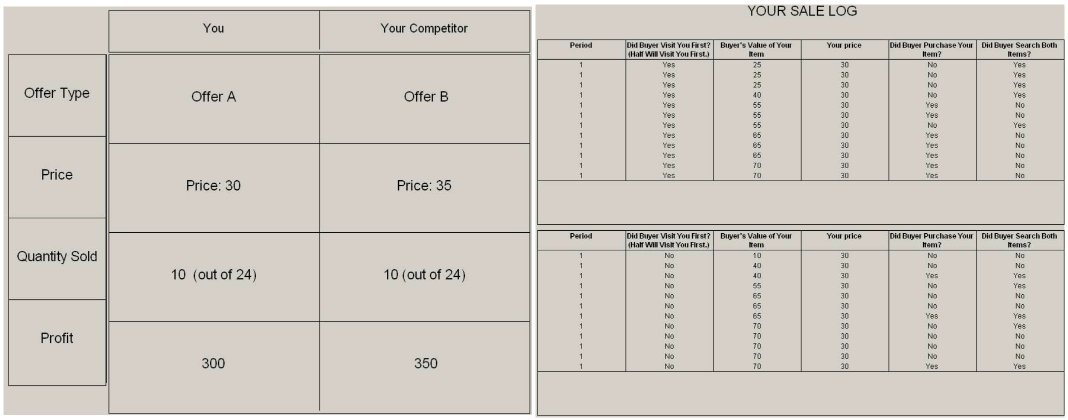

Figure 2: A Seller's Feedback Screen at the End of the Period. Sellers could toggle between each of the screens. (a, left) Both sellers are informed on the performance of each other in the market in aggregate. (b, right) Each seller observes all 24 buying decisions.

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

Before each session began, the instructions were shown on screen and read aloud to ensure the game was common knowledge among the subjects. After seeing and hearing the instructions, the subjects answered a multiple-choice quiz about the game to ensure that they understood how to play it. 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 20 periods 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).

The risk-preference question asked subjects to choose their preferred 50/50 lottery from the following five outcome pairs, increasing in order of risk: ($10, $10),

($18, $6), ($26, $2), ($34, $-2), ($42, $-6).[22] Since economic theory predicts clear departures from the risk-neutral equilibrium in our game for risk-averse agents, these survey responses are relevant to categorize subject deviations from that equilibrium.

The CRT consists of three questions with apparently easy, but incorrect, answers.[23] Answering the questions correctly requires more deliberative thinking. Frederick (2005) proposed the test as a measure of an individual's propensity to use System 1 ("quick") or System 2 ("deliberative") thinking (for more on dual-system theory see Stanovich and West, 2000; Kahneman, 2011). Results are shown to correlate with susceptibility to cognitive biases (e.g., Oechssler et al., 2009; Hoppe and Kusterer, 2011; Besedeš et al., 2012), strategic thinking (e.g., Brañas-Garza et al., 2012; Carpenter et al., 2013; Kiss et al., 2016), and other-regarding preferences (e.g., Corgnet et al., 2015; Cueva et al., 2016; Peysakhovich and Rand, 2015).[24] Because all of these behavioral tendencies might cause systematic deviations from equilibrium play, we hypothesized the test would correlate with such deviations.

Subjects were then privately paid their earnings (plus a five dollar show-up bonus) in the session in cash. Each seller in both treatments was paid based on one randomly selected period.[25] 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

[22]By the same order, the distribution of subjects' most preferred pairs is 29%, 26%, 23%, 5%, 17%.

[23]The questions are: "A bat and a ball cost $1.10 in total. The bat costs $1.00 more than the ball. How much does the ball cost (in cents)?" (easy: 10, correct 5); "If it takes 5 machines 5 minutes to make 5 widgets, how long would it take 100 machines to make 100 widgets (in minutes)?" (easy: 100, correct: 5); "In a lake, there is a patch of lily pads. Every day, the patch doubles in size. If it takes 48 days for the patch to cover the entire lake, how long would it take for the patch to cover half the lake (in days)?" (easy: 24, correct: 47). Respectively, 46%, 25%, 16%, and 13% of subjects answered zero, one, two, and three questions correctly.

[24]See Brañas-Garza et al. (2015) for a full survey.

[25]We chose to pay for one random decision to eliminate any subject complementarities that might occur across decisions or periods, most notably income effects. See Azrieli et al. (forthcoming) and Brown and Healy (forthcoming) for a greater discussion.

between the value and the price of that particular item purchased. The earnings were zero if no purchase was made. The conversion rate for a buyer was $1 for two points.[26] 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 2013 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).

5 Results

Result 1 *Sellers played equilibrium strategy five times more frequently in the computer-buyer treatment than in the human-buyer treatment. That is, they used exploding offers more often and offered higher prices against computer buyers. Over time, sellers increased prices, used exploding offers more often, and played equilibrium strategy more often in the computer-buyer treatment. In contrast, in the human-buyer treatment, sellers reduced prices and played equilibrium strategy less often over time.*

[26]We chose a larger conversion rate for buyers to make it similar to many field settings. Often in field settings, 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 were compensated based on six buying decisions from each of the four buyers, while buyers were compensated based on one buying decision. The differential rates of payment were common knowledge to all participants.

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

Table 3, Panel A provides a breakdown of seller exploding offer use by treatment, collapsed to the subject level. Sellers used exploding offers roughly two-thirds as often against human buyers as they did against computer buyers (66% CB vs. 47% HB; t-test and Mann-Whitney-Wilcoxon, $p < 0.01$).[27] This difference is more pronounced under the low-buyer-heterogeneity than the high-buyer-heterogeneity distribution, though both are significant at the 5 percent level when evaluated separately. Table 3, Panel B shows seller pricing decisions have similar differences across treatments. The prices that sellers offered to human buyers were roughly three points lower than the ones they offered to 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 percent 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.[28] Appendix Table A.1 provides a listing of all 18 cohorts and their respective averages of our three dependent variables. In 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[29] for

[27]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. For statistical tests (i.e., t-tests and Mann-Whitney-Wilcoxon) we accomplish this by comparing subject averages (collapsing 20 seller or 120 buyer decisions). For regressions, we cluster at the subject level.

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

[29]Since there are seven human- and 11 computer- buyer cohorts, there are 31,824 possible mappings

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 percent level (see Appendix Tables A.3 and A.4).

Figure 3 displays the dynamics of seller decisions across treatments. Over 20 periods, sellers in the CB treatment appear to increase their use of exploding offers. The relation is inconclusive for sellers in HB treatment. Linear trend analysis confirms this finding: sellers' 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 details). By the final five periods, about 73 percent of sellers in the CB treatment used an exploding offer, whereas only about 51 percent 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, as 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 that both trends are significant (see Appendix Table A.7).

Result 2 *Buyers deviated from optimal, profit-maximizing strategies more often when they encountered an exploding offer compared to when they encountered 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 six buying decisions in each period over 20 periods. Pooling the results from the seven sessions of 16 buyers each, there were a total of 13,440

of our dependent variable to these cohorts.

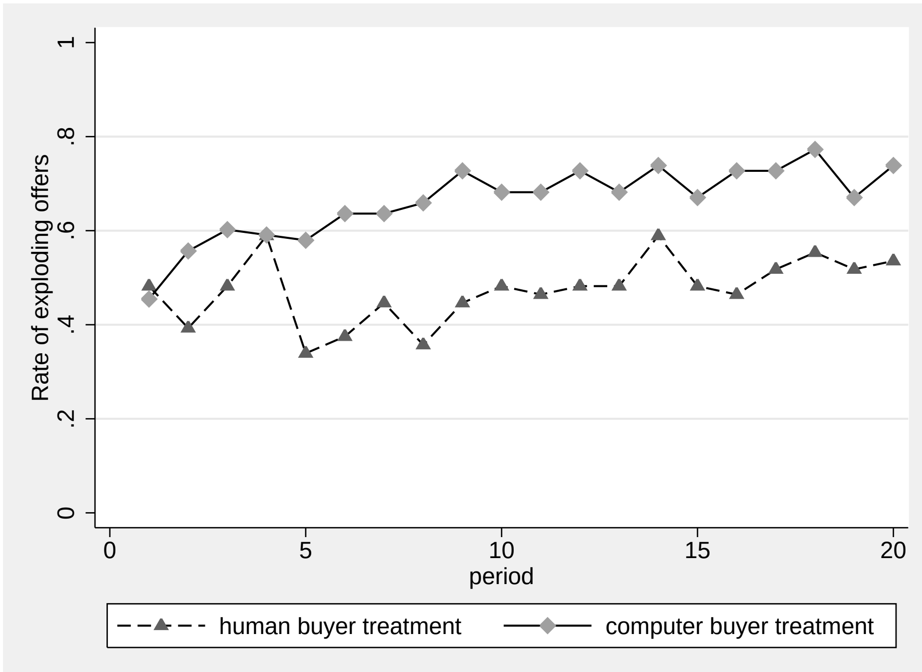

Figure 3: Rate of Exploding Offers Used by Sellers by Period, HB and CB Treatments

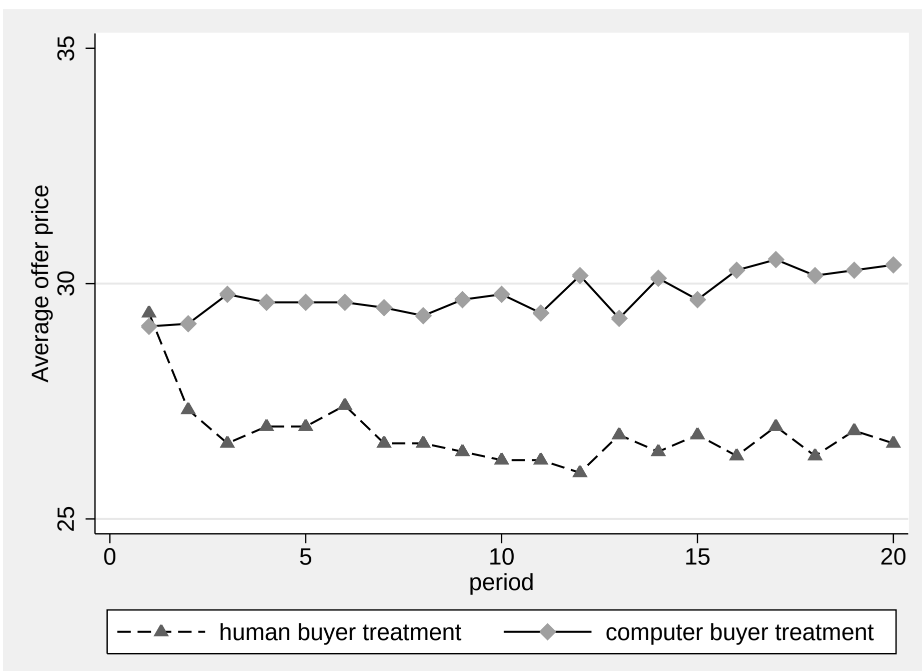

Figure 4: Average Price Offered by Sellers by Period, HB and CB Treatments

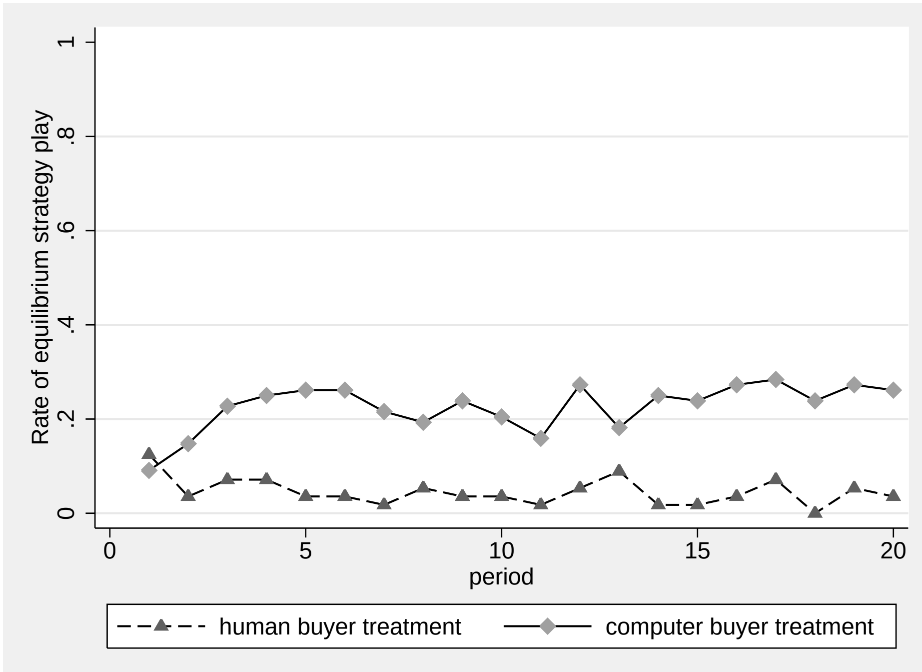

Figure 5: Rate of Equilibrium Play by Sellers by Period, HB and CB Treatments

($6 \times 20 \times 16 \times 7$) buying decisions. Sellers used exploding offers in 47 percent of these decisions. This resulted in 6,372 exploding offer buying decisions and 7,059 decisions with free-recall offers.[30]

The ways buyers could have deviated 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 violate the profit-maximizing strategy only 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 made 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 was negative (positive). Table 4, Panel A shows subject-level frequency of these two types of suboptimal play. Perhaps

[30]Due to a computer glitch, nine buying attempts were not 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 affects 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 (mean, standard deviation, number of observations)			
Valuation Type	Exploding Offer	Free-recall Offer	Difference (Paired)
Overall	0.157	0.058	0.100***
(112 subjects)	(0.125)	(0.108)	(0.121)
High Buyer-Heterogeneity	0.103	0.047	0.056***
(48 subjects)	(0.080)	(0.082)	(0.090)
Low Buyer-Heterogeneity	0.198	0.066	0.132***
(64 subjects)	(0.138)	(0.123)	(0.132)
Panel B: Buyer Deviation Rate, Exploding Offer Rejection vs. Acceptance (mean, standard deviation, number of observations)			
Valuation Type	Exploding Offer Should Be Accepted	Exploding Offer Should Be Rejected	Difference (Paired)
Overall	0.214	0.077	0.137***
(112 subjects)	(0.199)	(0.110)	(0.223)
High Buyer-Heterogeneity	0.143	0.036	0.107***
(48 subjects)	(0.126)	(0.065)	(0.154)
Low Buyer-Heterogeneity	0.266	0.108	0.159***
(64 subjects)	(0.226)	(0.126)	(0.262)

unsurprisingly given the greater difficulty of making a decision with incomplete information, buyers are roughly 10 probability points (or three times) more likely to deviate from the optimal play with an exploding offer than with a free-recall offer. The average rate of suboptimal play with an exploding offer is 16 percent compared to 6 percent 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 instead they reject it. In other cases, buyers should reject an exploding offer, but they accept it. Overall, buyers encountered slightly more exploding offers that they should have accepted than ones they should have rejected (3,754 vs. 2,618, 59% vs. 41%).[31] 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 made suboptimal decisions 14 percentage points (or three times) more often by rejecting an exploding offer than by 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 one in four exploding offers *that should have been accepted* were rejected.

As a robustness check we also can examine these results at the cohort level. As Appendix Table A.2 shows, there were seven human buyer cohorts. In each of the seven, the rate of suboptimal play was higher when buyers faced an exploding offer compared to when they faced a free-recall offer. Similarly, in each of the seven, the rate of suboptimal play was higher when buyers faced an exploding offer they should accept compared to when they faced 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 that 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 illustrates that suboptimal play is more likely in rejection of exploding offers than

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

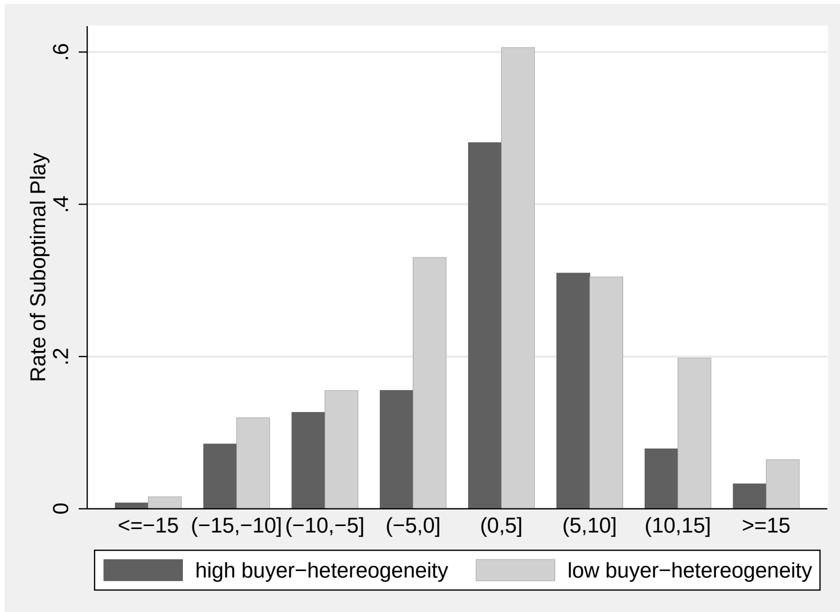

Figure 6: Rate of Buyer Suboptimal Play by Cost of Rejecting Exploding Offer

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 of the model is

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

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 20 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 encountered an exploding offer that should be accepted rather than rejected. To better understand this phenomenon, we focus on what buyer characteristics are most correlated with this incongruity. Our specific focus is gender, elicited risk preferences, and CRT scores.[32]

[32]Sections 3.4 and 4 provided background on how risk and systems of thinking might explain

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,

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

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 and correct responses on the CRT (0, 1, 2, or 3). Variables are taken from non-incentivized survey questions (see Section 4 for more detail). Gender is elicited directly from subjects in a demographic survey. The grouping of preferred gambles, 1 or 2 vs. 3, 4, or 5—roughly half the subjects fall into each group (see footnote 22)—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) is used as a proxy for subjects' propensity to inhibit "quick" (System 1) responses and engage in deliberative (System 2) decision-making (see Stanovich and West, 2000; Kahneman, 2011).

Table 6 provides the results of this regression. To correctly interpret interaction terms, average marginal effects are calculated when the 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 rather than rejecting it. The coefficients of risk aversion are in the direction one would suspect (see Section 3.4 for more detail). Acceptance of an exploding offer leads to a certain payoff, while visiting the second seller involves

deviations from the risk-neutral theoretical predictions, respectively. We include gender because it is known to correlate with risk and social preferences (Croson and Gneezy, 2009), as well as the CRT (Frederick, 2005), and we have imperfect proxies of these measures.

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	

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

As one would expect, subjects who answered more questions correctly 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.

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. Seller play indicates a significant aversion to the use of exploding offers against human buyers that cannot be explained by payoff differences alone.*

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 on the empirically observed play of subjects; they are based on the predicted rejection probabilities from equation (5), omitting period effects. In each of the 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.[33]

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 Equation (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 (or 30E) was most profitable for sellers. In the HB treatment, 25E was most profitable for sellers.

[33] 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 (Row Player Payoff Shown), 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

Payoffs for the row player shown. Because the game is symmetric, the payoffs for the column player are found in the transpose cell. The "human buyer" payoff matrix is calculated the same way as the theoretically optimal matrix, except that the observed rejection rate of exploding offers is used rather than the theoretical optimum. The equilibrium payoffs are in bold.

Result 3 provides insight 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 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.

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

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 1 if the current strategy action involves using an exploding offer.

The results show that, compared to their computer buyer counterparts,[34] 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 distribution than it is under the low-buyer-heterogeneity distribution.[35] The results also indicate that sellers are more sensitive to variations of expected payoffs because of the other competitors or buyers in the human-buyer treatment.

The types of sellers that refrain from using exploding offers are explained in the

[34]Note that we use the same valuation draws across different sessions with high- and low- buyer-heterogeneity distributions.

[35]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, yield 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

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, a few sellers mentioned the use of free-recall offers as part of their general strategies. The reason they give for not using exploding offers is 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 toward 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, risk measures and CRT scores both 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,[36] gender effects are quite substantial. Male

[36]Female subjects generally exhibit higher degrees of risk aversion on experimental tasks than do

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	

sellers are 17 probability points more likely to use exploding offers than female sellers to use exploding offers with human buyers. There is no difference with computer buyers. Confirming previous results, sellers were much less likely to use exploding offers on human buyers than on computer buyers.

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 32 sellers noted the use of free-recall offers as part of their general strategies. Most justified this choice with a concern for 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 percent of their offers; the others used exploding offers in 50 percent of their offers ($p < 0.01$, t-test and Mann-Whitney-Wilcoxon). Interestingly, there is no significant difference between the average price offered (26.8 vs. 26.6) by these two groups, suggesting

males (see Croson and Gneezy, 2009, for a survey).

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

this tendency is unrelated to the desire to transfer income to buyers.

The survey response may also explain the difference between genders in the use of exploding offers. Female subjects were 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 percent of male subjects (4 of 21) indicated a concern for buyers ($p \approx 0.056$ on Fisher exact test). Differences on exploding offer use among the subjects who indicated this concern (23 percent exploding offer use for female subjects vs. 30 percent for men) and those who do not (51 percent female vs. 50 percent 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' s-trategies constant, buyers' deviations from best response cause slight losses to 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 increased by \$0.40 as sellers use some amount of free-recall offers against computer buyers.

Human buyers did not respond to sellers' offers the same way computer buyers did. As payoff Table 7 shows, accounting for these new strategies makes (25E, 25E) the new equilibrium for sellers instead of (35E, 35E). Sellers would have done very poorly if they played this equilibrium solely against human buyers. They would have made \$10.77 and \$11.24 in the low- and high-buyer-heterogeneity distributions of the HB treatment, respectively. However, sellers did not play this strategy exclusively. Here they increased 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 did not behave as optimal play would dictate. If we hold seller strategies as constant, buyers could have improved their earnings by \$0.40–\$0.70 by following optimal strategies. Sellers were not particularly hurt directly by these deviations; they cost sellers only about \$0.15 of their earnings. Of course, seller strategies were not constant; the way buyers responded to exploding offers greatly changes the strategic structure of the game. This had a second-order effect of moving the seller equilibrium from (35E, 35E) to (25E, 25E), which costs sellers more than \$4.00 in earnings in theory (or \$1.00–\$1.50 in practice).

Buyer behavior alone eliminated 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.

6 Discussion

Our experiment identifies two pronounced departures from equilibrium. First, buyers deviate from optimal strategies more often by rejecting than by accepting them. Second, sellers demonstrate a significant reluctance to use exploding offers on human buyers. This reluctance is 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 computer buyers than human buyers.

The behavior of buyers who answer the fewest questions correctly on the CRT is crucial to this result. This category of subjects is considered the most prone to System 1 thinking. These subjects are most likely to reject exploding offers suboptimally, but they are no more likely than other subjects to suboptimally accept them. A possible explanation is that the fast, intuitive response to an exploding offer is to run from the offer and search further. Only after deliberation does one understand the value in accepting such offers.

Another explanation is fairness. Exploding offers could be viewed as unjustified capture of surplus compared to a normal level of profit in this market. Kahneman et al. (1986) suggest in such instances sellers will be loath to use such tactics and if they do, buyers will react negatively to them. We see no exit-survey evidence that suggests buyers feel this way. Certain sellers who infrequently use exploding offers appear to make claims consistent with this principle. For sellers, there is no correlation between correct CRT answers and reluctance to use exploding offers.

These two explanations highlight two major ways behavioral economics is essential to understanding markets and games. First, agents may have norms concerning what are reasonable profits and tactics in an economic environment. Numerous empirical and experimental studies suggest agents will need to restrict their actions to comply with these norms or face retaliation from other agents. Second, System 1 responses in economic transactions may alter market equilibrium as well, even when a player does not feel wronged by the other party. Here, buyers may intuitively feel that the correct response to a high-pressure offer is to reject it, thereby reducing the profitability of that technique.

The applications of this latter explanation are far less studied but may be quite meaningful. For instance, in personal finance, mortgage refinancing offers always carry a deadline because of interest rate variability. While the deadline is unlikely to be perceived as unfair, the existence of a deadline may lead those consumers who rely heavily on System 1 to not refinance. Thus, our results may provide one explanation for the anomaly that a significant portion of homeowners do not refinance their mortgages when economic conditions are unambiguously favorable (see Keys et al., 2016).

There are some suggestive policy implications from our work. The European Union prohibits sellers from making false claims about product availability. Armstrong and Zhou (2016) advocate a cooling-off period after sales as an alternative. Our findings indicate the need to make such policy changes may not be as great as theory predicts. While a firm could selectively hire, train or automate its sale staff to use search deterrence without compunction, some potential customers will be prone to fast thinking. The response of these customers will diminish potential gains from any search-deterrence strategy.

Finally, our experiment provided, as intended, an environment where sellers' pressure was clearly intentional but contained minimal emotional connotation. In different contexts, the pressure sale may be more salient and the intent of sellers may be less clear. This nuance may be tremendously important. Among search-deterrence techniques, flash sale websites are fading from popularity, while scheduled single-day sales are thriving. In terms of standard theory, there is not a substantial difference between these two methods of sales. However, it's likely that customers feel greater pressure from constant notification of temporary sales than from expecting quick sales on a single, scheduled day. There is likely a deeper, complex relationship between the nature of the pressure sell and buyers' suboptimal, negative response. 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.

Yaron Azrieli, Christopher P. Chambers, and Paul J. Healy. Incentives in experiments: A theoretical analysis. *Journal of Political Economy*, forthcoming.

Tibor Besedeš, Cary Deck, Sudipta Sarangi, and Mikhael Shor. Decision-making strategies and performance among seniors. *Journal of Economic Behavior & Organization*, 81(2):524–533, 2012.

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

Pablo Brañas-Garza, Teresa Garcia-Muñoz, and Roberto Hernán González. Cognitive effort in the beauty contest game. *Journal of Economic Behavior & Organization*, 83(2):254–260, 2012.

Pablo Brañas-Garza, Praveen Kujal, and Balint Lenkei. Cognitive reflection test: whom, how, when. 2015.

Alexander L. Brown and Paul J. Healy. Separated decisions. *European Economic Review*, forthcoming.

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.

Jeffrey Carpenter, Michael Graham, and Jesse Wolf. Cognitive ability and strategic sophistication. *Games and Economic Behavior*, 80:115–130, 2013.

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

Robert B. Cialdini. *Influence : Science and Practice*. Allyn & Bacon, Boston, 4. ed. edition, 2003.

David Cooper and John H Kagel. Other regarding preferences: a selective survey of experimental results. In John H Kagel and Alvin E Roth, editors, *The Handbook of Experimental Economics, Volume 2: The Handbook of Experimental Economics*, chapter 4, pages 217–289. Princeton university press, 2016.

Brice Corgnet, Antonio M Espín, and Roberto Hernán-González. The cognitive basis of social behavior: cognitive reflection overrides antisocial but not always prosocial motives. *Frontiers in Behavioral Neuroscience*, 9, 2015.

Pascal Courty and Li Hao. Sequential screening. *The Review of Economic Studies*, 67 (4):697–717, 2000.

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.

Carlos Cueva, Inigo Iturbe-Ormaetxe, Esther Mata-Pérez, Giovanni Ponti, Marcello Sartarelli, Haihan Yu, and Vita Zhukova. Cognitive (ir) reflection: New experimental evidence. *Journal of Behavioral and Experimental Economics*, 64:81–93, 2016.

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

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

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.

Werner Güth, Rolf Schmittberger, and Bernd Schwarze. An experimental analysis of ultimatum bargaining. *Journal of Economic Behavior & Organization*, 3(4):367 – 388, 1982.

Eva I Hoppe and David J Kusterer. Behavioral biases and cognitive reflection. *Economics Letters*, 110(2):97–100, 2011.

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

Daniel Kahneman, Jack L Knetsch, and Richard Thaler. Fairness as a constraint on profit seeking: Entitlements in the market. *The American Economic Review*, pages 728–741, 1986.

Benjamin J Keys, Devin G Pope, and Jaren C Pope. Failure to refinance. *Journal of Financial Economics*, 122(3):482–499, 2016.

Hubert Janos Kiss, Ismael Rodriguez-Lara, and Alfonso Rosa-García. Think twice before running! bank runs and cognitive abilities. *Journal of Behavioral and Experimental Economics*, 64:12–19, 2016.

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

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.

Frank Lavin. Alibaba's singles' day: What we know about the world's biggest shopping event, 2016. URL `https://www.forbes.com/sites/franklavin/2016/11/06/alibabas-singles-day-what-we-know-about-the-worlds-biggest-shopping-event/`.

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.

Volker Nocke, Martin Peitz, and Frank Rosar. Advance-purchase discounts as a price discrimination device. *Journal of Economic Theory*, 146(1):141–162, 2011.

Jörg Oechssler, Andreas Roider, and Patrick W Schmitz. Cognitive abilities and behavioral biases. *Journal of Economic Behavior & Organization*, 72(1):147–152, 2009.

Alexander Peysakhovich and David G Rand. Habits of virtue: Creating norms of cooperation and defection in the laboratory. *Management Science*, 62(3):631–647, 2015.

Leena Rao. Amazon shuts down its flash sales site, 2016. URL `http://fortune.com/2016/04/21/amazon-flash-sales/`.

Robert J Robinson. Defusing the exploding offer: The farpoint gambit. *Negotiation Journal*, 11(3):277–285, 1995.

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.

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

Anuj K Shah, Sendhil Mullainathan, and Eldar Shafir. Some consequences of having too little. *Science*, 338(6107):682–685, 2012.

Spencer Soper. Amazon's prime day generates estimated $1 billion in sales, 2017. URL `https://www.bloomberg.com/news/articles/2017-07-12/amazon-s-prime-day-proves-to-be-biggest-shopping-day-ever`.

Keith E Stanovich and Richard F West. Individual differences in reasoning: Implications for the rationality debate? *Behavioral and Brain Sciences*, 23(5):645–665, 2000.

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 does 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 does 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 does 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 does 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***	-0.214***	-0.250***	-0.253***	-0.184***	-0.184***
	(0.016)	(0.018)	(0.022)	(0.027)	(0.021)	(0.022)
acceptance is optimal	-	0.822***	-	1.085**	-	0.699**
		(0.275)		(0.517)		(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 available only 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 available only 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

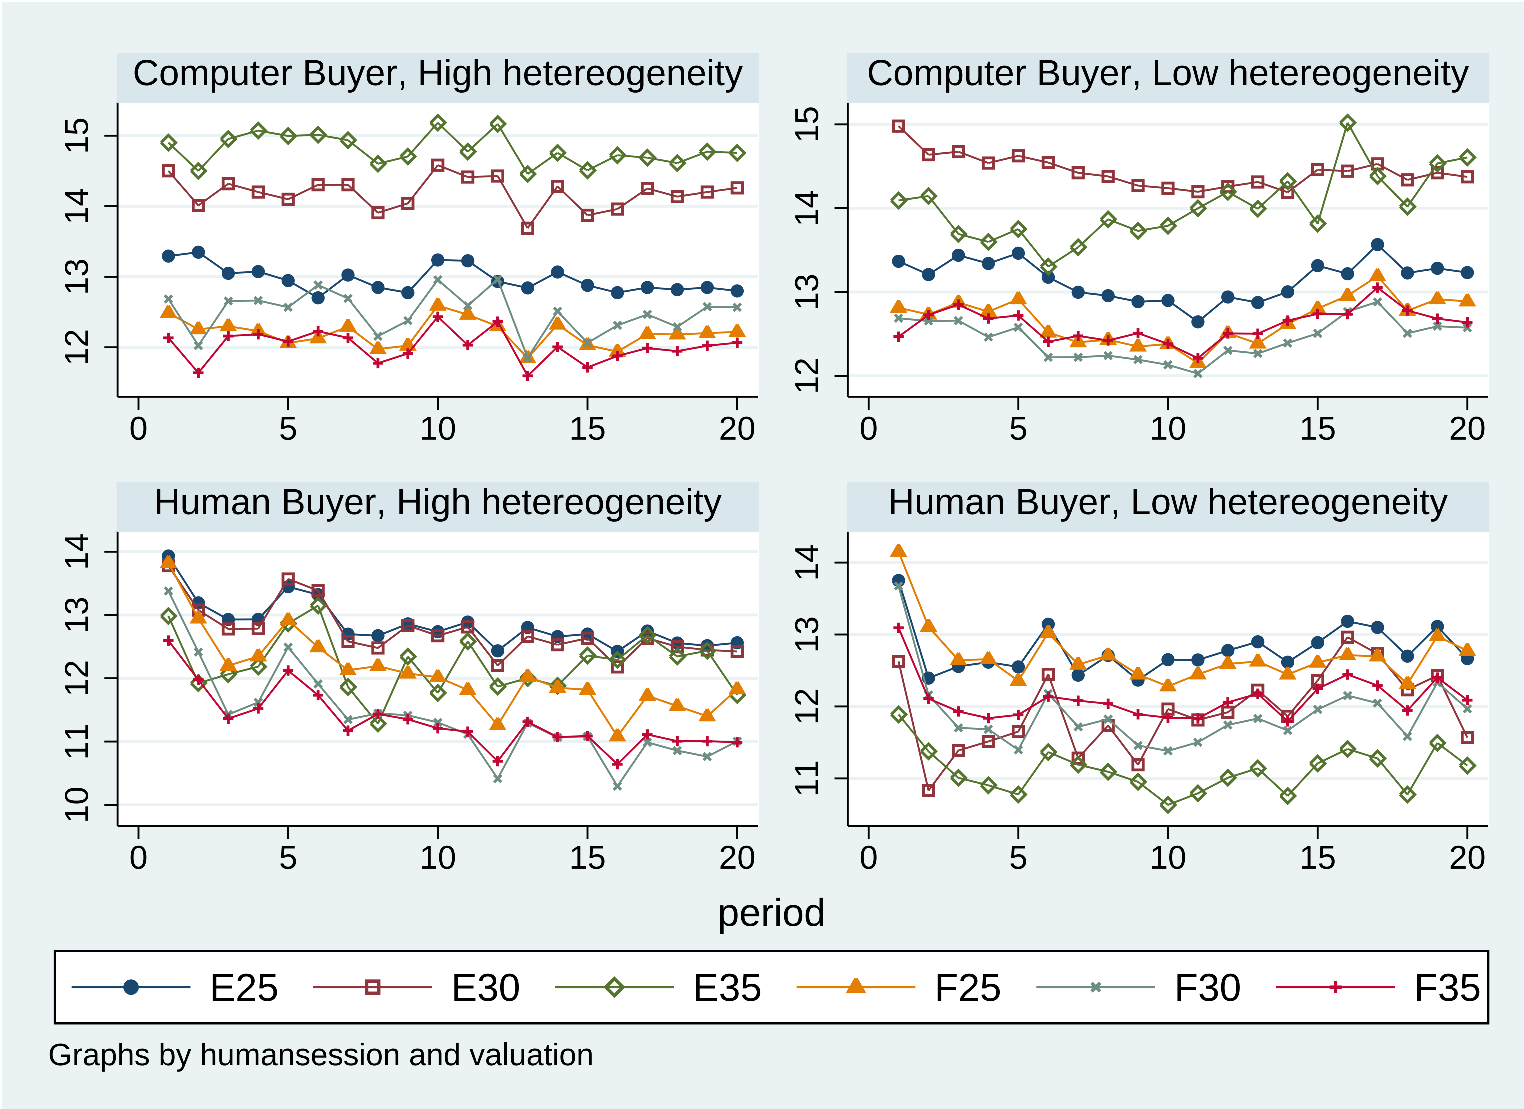

Figure A.1: Payoff Vector for: CB Treatment (high heterogeneity valuation) (left-top), CB Treatment (low heterogeneity valuation) (left-bottom), HB Treatment (high heterogeneity valuation) (right-top), HB Treatment (low heterogeneity valuation) (right-bottom)