

Testing a Purportedly More Learnable Auction Mechanism

Katherine L. Milkman
Harvard Business School, Harvard University, Boston, MA 02163
kmilkman@hbs.edu

James Burns
Department of Economics, Harvard University, Cambridge, MA 02138
jburns@fas.harvard.edu

David C. Parkes
School of Engineering and Applied Sciences, Harvard University, Cambridge, MA 02138
parkes@eecs.harvard.edu

Greg Barron
Harvard Business School, Harvard University, Boston, MA 02163
gbarron@hbs.edu

Kagan Tumer
Mechanical Engineering Department, Oregon State University, Corvallis, OR 97331
kagan.tumer@oregonstate.edu

We describe an auction mechanism in the class of Groves mechanisms that has received attention in the computer science literature because of its theoretical property of being more “learnable” than the standard second price auction mechanism. We bring this mechanism, which we refer to as the “*clamped second price auction mechanism*,” into the laboratory to determine whether it helps human subjects learn to play their optimal strategy faster than the standard second price auction mechanism. Contrary to earlier results within computer science using simulated reinforcement learning agents, we find that both in settings where subjects are given complete information about auction payoff rules and in settings where they are given no information about auction payoff rules, subjects converge on playing their optimal strategy significantly faster in sequential auctions conducted with a standard second price auction mechanism than with a clamped second price auction mechanism. We conclude that while it is important for mechanism designers to think more about creating learnable mechanisms, the clamped second price auction mechanism in fact produces slower learning in human subjects than the standard second price auction mechanism. Our results also serve to highlight differences in behavior between simulated agents and human bidders that mechanism designers should take into account before placing too much faith in simulations to test the performance of mechanisms intended for human use.

Keywords: auctions; mechanism design; learning; experimental economics

Author’s Note: We thank Al Roth, Max Bazerman and Itay Fainmesser for their insightful feedback on this work. Thanks to Harvard Business School and the Alfred P. Sloan Foundation (for Parkes) for funding support.

I. Introduction

Each year, auctions are used to determine how billions of dollars of goods and services will be allocated across the globe. On eBay alone, \$52.5 billion in merchandise was exchanged in 2.4 billion auctions conducted during FY2006.¹ Considerable attention has been paid in the academic literature to the question of how to design auctions with efficient allocation (for example, see Vickrey, 1961; Dasgupta and Maskin, 2000; Ausubel, 2004) and revenue maximizing (for example, see Myerson, 1981; Milgrom and Weber, 1982; Armstrong, 2000) properties. However, in part because auction rules are typically published and standard theory assumes economic agents are capable of computing optimal strategies from published rules, little attention has been paid to the question of how to design auctions whose optimal strategies are easy to learn. In fact, experimental evidence suggests that even when auction rules are published and dominant strategies exist, people nonetheless struggle and sometimes fail to learn to play their optimal strategy (Kagel et al., 1987; Kagel and Levin, 1993; Ariely et al., 2005).² As a result, we argue that the question of how to design a learnable, strategy-proof auction mechanism is an important one.

In this paper, we describe an auction mechanism in the class of Groves mechanisms whose design was inspired by recent work in the computer science literature aimed at producing more “learnable” utilities for agents in multiagent systems (Tumer et al., 2002; Agogino and Tumer, 2005; Wolpert and Tumer, 2001). This auction mechanism, which we call the “*clamped second price auction mechanism*,” has received attention in the computer science literature because of its theoretical property of being more “learnable” than the standard second price auction mechanism (Parkes, 2004). It was designed to provide agents with payoffs that are more responsive to their bids than the payoffs in a standard second price auction. Rather than determining agents’ payoffs by charging them the total loss in value they

¹ Company Annual Reports, January 18, 2006 earnings announcement.

² Sponsored search provides another motivation for the study of learnable mechanisms (Edelman et al., 2007; Varian, 2007). In sponsored search, auctions are used to allocate advertisers to web page slots adjacent to the results of a search query. Not only are the rules of the auction protocols used by companies such as Google and Yahoo! partially hidden from participants in that they depend in part on the predicted “click through rates” made by proprietary machine learning algorithms, but there is considerable empirical evidence of adaptive bidding behavior in these auction environments (Feng and Zhang, 2007; Edelman and Ostovsky, 2007; Borgs et al., 2007).

impose on other participants by bidding in an auction, the clamped second price mechanism charges agents the difference between the hypothetical total value to other participants if the agent in question had not bid in an auction and the hypothetical total value to all participants if the bid of the agent in question were replaced with an alternative bid. Simulations conducted with reinforcement learners have suggested that this auction mechanism has the potential to help agents converge on their optimal, Nash strategy faster than the standard second price auction mechanism (Parkes, 2004). We bring the clamped second price auction mechanism into the laboratory and conduct two studies to determine whether it helps human subjects in three-player auctions learn to play their optimal strategy faster than the standard second price auction mechanism.

Our findings suggest that both in settings where subjects are given complete information about auction payoff rules and in settings where they are given no information about auction payoff rules, subjects converge on playing their optimal strategy significantly faster in sequential auctions conducted with a standard second price auction mechanism than in auctions conducted with a clamped second price auction mechanism. We conclude that while it is important for mechanism designers to think more about creating learnable mechanisms, the clamped auction mechanism studied by Parkes (2004) in the context of simulated reinforcement learning agents in fact produces *slower* learning in human subjects than the standard second price auction mechanism. Our results also allow us to explore some of the ways in which the learning behaviors exhibited by simulated reinforcement agents differ from those exhibited by human subjects.

The rest of this paper is organized as follows. Section II reviews the relevant literature on auctions and the computer science literature on learnable mechanism design. In Section III we present the results of study one, in which subjects participate in a series of three-player auctions whose payoff rules are explained, and subjects are either assigned to a standard second price auction mechanism or a clamped second price auction mechanism. In Section IV we present the results of study two, in which subjects participate in a series of three-player auctions whose payoff rules are not explained, and subjects are again either assigned to a standard second price auction mechanism or a clamped second price

auction mechanism. Section V discusses the results of our two experimental studies, potential explanations for and implications of these results, and concludes.

II. Literature Review

A. The Class of Groves Mechanisms

The payment and allocation rules of the auctions discussed in this paper are drawn from the set of Groves mechanisms (1973). In an auction environment, a Groves mechanism is any mechanism that satisfies the following two properties:

- (1) The auction selects an allocation that maximizes the total reported values of all bidders.
- (2) The difference between the payment a bidder would owe given any two, different bids is equal to the effect of the resulting change in allocation (if any) on all other bidders' reported values.

These properties make truthful revelation of preferences a dominant strategy in such a mechanism, and they also ensure that allocations are efficient (Groves, 1973). A generic Groves mechanism payment function is given by:

$$p_{groves,i}(\hat{\theta}_i, \hat{\theta}_{-i}) = \sum_{j \neq i} v_j(g_{eff}(\hat{\theta}_i, \hat{\theta}_{-i}), \hat{\theta}_j) - h_i(\hat{\theta}_{-i})$$

where $p_{groves,i}$ is agent i 's payment to the mechanism, $\hat{\theta}$ is an agent's reported values, $g_{eff}(\hat{\theta}_i, \hat{\theta}_{-i})$ is the allocation rule that would maximize global utility given all agents' reported values, $v_j(k, \theta)$ corresponds to agent j 's value for a given allocation, k , given his reported values, θ , and $h_i(\hat{\theta}_{-i})$ is an arbitrary function on the reported values of all agents except agent i . Bidder i 's payoff is determined by the value $v_i(g_{eff}(\hat{\theta}_i, \hat{\theta}_{-i}), \theta_i)$ to the bidder for the allocation, summed with the first component of the payment term, $\sum_{j \neq i} v_j(g_{eff}(\hat{\theta}_i, \hat{\theta}_{-i}), \hat{\theta}_j)$, and net of the second component, $h_i(\hat{\theta}_{-i})$. In considering the incentive properties of the mechanism, this latter component can be ignored, and the bidder's payoff will be aligned with the sum of her own true value for the allocation and the total reported value of the allocation to all

other bidders. This is what makes truthful preference revelation optimal in a Groves mechanism: if an agent reports her true value to the mechanism, it will select an allocation that maximizes her true value, a characteristic that follows from the first property of a Groves mechanism stated above. Note that the above representation of a Groves mechanism allows for a degree of freedom in the choice of h_i while preserving the desirable properties of the mechanism. The auction mechanisms studied in this paper differ only in the choice of this function, h_i .

B. The Standard Second Price Auction Mechanism

The standard second price auction mechanism discussed in this paper, commonly referred to as the Vickrey-Clarke-Groves (VCG) mechanism, is a Groves mechanism with the attractive property that the payment made by each auction participant is equal to the cost that she imposes on other auction participants by placing her bid in the auction. This is accomplished by setting $h_i(\cdot)$ equal to the net value to all bidders other than bidder i of the efficient allocation that would be achieved if bidder i were removed from the auction, given all participants' reported values. Specifically, the second price auction mechanism assigns agents' payments as follows:

$$P_{\text{standard_second_price},i}(\hat{\theta}) = \sum_{j \neq i} v_j(g_{\text{eff}}(\hat{\theta}), \hat{\theta}_j) - \sum_{j \neq i} v_j(g_{\text{eff}}(\hat{\theta}_{-i}), \hat{\theta}_j)$$

where $g_{\text{eff}}(\hat{\theta}_{-i})$ is the efficient allocation, according to agents' reported values, that would be computed if agent i were not participating in the mechanism. In an environment in which a single good is being auctioned off, this implies that the item is allocated to the highest bidder at a price equal to the second highest bid and that all other bidders will make no payments to the auction mechanism. As is the case with all Groves mechanisms, it is a dominant strategy for agents faced with this mechanism to truthfully reveal their preferences, and allocations determined by the mechanism are efficient.

C. Learning in Auctions

The failure of bidders to learn to use Nash bidding strategies when interacting with a standard second price auction mechanism (or with other auction mechanisms boasting efficient allocation or

revenue maximizing properties) is well documented in the experimental literature (for a survey, see Kagel, 1995). Early contributors to the experimental study of the standard second price auction include Kagel, et. al (1987), who find that in an affiliated private values setting, anonymous, randomly matched bidders in a second price auction bid, on average, significantly above the dominant strategy bid of their true value (experiments previous to this, such as Cox et al. (1982), documented very different bidder behavior, but these authors restricted bids to be below the bidder's value). Not only do Kagel, et al. (1987) find that players persistently overbid relative to their dominant strategy in a second price auction, they also find that final auction prices are significantly higher than the predicted dominant strategy price in 80% of the auctions in their experiment. In addition, Kagel et al. (1987) find that significant overbidding relative to the dominant strategy continues even after subjects participate in 30 successive auctions. This pattern of results is replicated in Kagel and Levin (1993) with randomly matched bidders in an independent private values auction setting. The deviations from the Nash strategy observed in the second price auction experiments conducted by Kagel et al. (1987) and Kagel and Levin (1993) represent strategic errors by bidders, as the Nash strategy is dominant regardless of players' risk preferences given random, anonymous matching. In another experiment conducted in a second price auction setting, Ariely, Ockenfels, and Roth (2005) examine bidding behavior in auctions with different ending rules and find that in a second-price, sealed bid auction, the median bidder's bid does converge to the dominant strategy bid, but only after 15 to 20 rounds.

These experiments demonstrate that players fail to learn their dominant strategies instantaneously in second price auctions, contradicting the prediction of standard economic theory. However, to our knowledge, there has been no previous experimental work that has attempted to modify the standard second price auction mechanism to increase the rate at which subjects learn to play their optimal strategy. Both the theoretical and experimental literatures on auction design have largely focused on exploring auction designs that maximize revenue or have efficient allocation properties, ignoring a mechanism's

learnability as a design criterion.³ We propose that learnability may be a relevant, additional criterion to consider when designing an auction mechanism.

D. The Theory of Collectives

A central problem in the design of multiagent systems is the coordination of actions of independent agents so that their collective behavior optimizes a system-wide objective. This is a particularly complex problem when communication is restricted and agents are faced with an unknown or uncertain environment. The theory of collectives (Wolpert and Tumer, 2001) approaches this problem by focusing on environments in which each agent is assumed to behave in such a way that he maximizes his own private utility using a boundedly-rational decision making algorithm. Under this assumption, the system designer's problem becomes one of imposing appropriate incentives or utility functions on individual agents such that the agents not only converge to an optimal strategy but do so rapidly.

The theory of collectives proposes two properties that are key to deriving agent utilities that will lead to coordinated system behavior (Wolpert and Tumer, 2001; Tumer and Wolpert, 2004). The first property, dubbed "*factoredness*," measures the degree of alignment between an agent's utility and the system utility. Intuitively, the higher the degree of factoredness between two utilities, the more likely it is that a change of strategy by an agent will have the same impact on the agent's utility and the system's utility. This alignment is key in ensuring that actions taken by an agent that are beneficial to that agent are also beneficial to the system as a whole. The second property, dubbed "*learnability*" measures the sensitivity of an agent's utility to its own strategies as opposed to the strategies of others. Intuitively, learnability measures the signal-to-noise ratio for an agent's utility where its own strategies represent the signal, and the strategies of other agents represent the noise. Agents have a hard time learning strategies

³ Mathevet (2007) considers the explicit goal of the design of learnable mechanisms and introduces the concept of *supermodular Bayesian implementation* in which mechanisms are constructed to induce supermodular games that are learnable in the sense of Milgrom and Roberts (1990). Mathevet (2007) suggests an approach for *approximate* supermodular implementation in the context of auctions but the mechanism remains untested in an experimental setting. Mathevet (2007) also provides a survey of the related literature on learning and stability in the context of implementation and mechanism design.

when their own utilities are affected by the actions of others, or in other words, when their signal becomes corrupted by the “noise” in the system.

Based on the goal of maximizing these two properties of factoredness and learnability, the theory of collectives provides a class of agent utility functions called “difference utilities” that are fully factored and have generally high learnability (Wolpert and Tumer, 2001; Tumer and Agogino, 2005). For a given system-wide utility function, $G(s)$, where s represents the strategies of the agents in the system (e.g., reported types in an auction context), the difference utility for agent i is given by:

$$DU_i(s) = G(s) - G(s_{-i}, \bar{s}_i)$$

where s_{-i} represents the strategies of all the agents except agent i , and \bar{s}_i is a constant strategy to which agent i 's strategy has been clamped (or fixed). This approach yields an agent utility that is the difference between the actual system utility and the system utility if agent i were replaced with a “neutral” agent, hence the name. The difference utility is fully factored because the second term of this equation does not depend on the actual strategies of agent i . Therefore, any change of strategy by agent i will impact DU and G in a similar manner (for differentiable DU and G , the two utilities have the same derivative). Furthermore, the second term removes some noise from agent i 's signal, giving DU generally higher learnability than G .

DU has been successfully applied in various multiagent domains, including air traffic flow management, data routing, and robot coordination (Tumer and Agogino, 2007; Wolpert and Tumer, 2002; Agogino and Tumer, 2005). As we discuss in the next section, the theory of collectives has also been used to provide the motivation for the purportedly more learnable auction protocol introduced in Parkes (2004) that we examine experimentally in this paper. It has, however, primarily been used for multiagent coordination in domains where *artificial* learning agents explore their strategy space to improve their own utility.

E. The Clamped Second Price Auction Mechanism

To apply the theory of collectives in the context of auction design, we define efficiency with respect to allocation as the system-wide utility-maximizing criteria and adjust payment functions rather than utility functions with the goal of increasing agents' speed of convergence to their optimal bidding strategies. The class of Groves mechanisms lends itself naturally to this application, as individual payments are aligned with efficient allocations and some flexibility exists in the determination of agents' payment functions. Parkes (2004) modifies the standard VCG payments according to the collectives model and runs a series of simulations with computer-generated reinforcement learning agents. Parkes compares the speed of learning in simulated auctions with adjusted VCG payments, which we refer to as clamped second-price auction payments, to the speed of learning in simulated auctions with standard second price payment rules and finds that reinforcement learning agents converge to their optimal bidding strategy more rapidly in the clamped second price auction environment than in the standard second price auction environment. This paper extends Parkes (2004) by examining whether *human bidders* converge to playing their optimal strategies more rapidly in the clamped second price auction environment than in the standard second price auction environment.

The payment rule instantiated by the clamped second-price auction mechanism is based on the *difference utility* (DU) described above, which is designed to provide payoffs that are less affected by variations in the bids of other agents, and thus more responsive to a bidder's own bid, than payoffs in a standard second price auction. According to the theory collectives, the increased sensitivity of agents' payoffs to their own bidding behavior should increase agents' speed of convergence to their optimal bidding strategies. In clamped second price auctions, the payment made by bidder i is equal to the difference between the net value to other bidders of the efficient allocation based on bidder i 's reported value and the net value to other bidders of the efficient allocation if bidder i 's bid were replaced with another hypothetical bidder's value – the auction's clamped value. The mathematical payment rule for the clamped second price auction, which is in the class of Groves mechanisms, is defined as:

$$P_{clamped_second_price,i}(\hat{\theta}) = \sum_{j \neq i} v_j(g_{eff}(\hat{\theta}), \hat{\theta}_j) - \left(v_i(g_{eff}(\bar{\theta}, \hat{\theta}_{-i}), \bar{\theta}) + \sum_{j \neq i} v_j(g_{eff}(\bar{\theta}, \hat{\theta}_{-i}), \hat{\theta}_j) \right)$$

where $\bar{\theta}$ is the “clamped value” of the auction, and $g_{eff}(\bar{\theta}, \hat{\theta}_{-i})$ is the efficient allocation that would be computed if bidder i were replaced in the mechanism with an agent who bid $\bar{\theta}$ for the item being auctioned off. The other variables and functions are defined as in Section II.A above. The value $v_i(g_{eff}(\hat{\theta}), \theta_i)$ provides a payoff to bidder i that is equal to the difference between the total value to all bidders when bidder i is present and the total value to all bidders when bidder i ’s bid is replaced with a clamped value, which maps onto the definition of a DU payoff from the theory of collectives. Wolpert and Tumer (2001) argue that setting the clamped value to the mean of a bidder’s value distribution will best approximate payments that optimize agents’ speed of convergence to their optimal strategy, and we adopt this procedure to determine the clamped value in our experiments with human subjects.

To illustrate the mechanics of the clamped second price auction, we briefly discuss how such an auction would work in an environment resembling the one we rely upon in the experiments presented in this paper. Consider a single good auction environment with three bidders whose values for the item being auctioned off are drawn independently from a uniform distribution over $[0,1]$. Setting the auction’s clamped value equal to $\frac{1}{2}$ as Tumer and Wolpert recommend, we consider the outcomes of the clamped second price auction in the three cases that result in distinct payoff rules. First, in the case where at least two of the bidders’ reported values for the object are greater than the clamped value of $\frac{1}{2}$, the payment function reduces to that of the standard second price auction where the winning bidder, who is always the bidder with the highest reported value, will pay a price equal to the second highest bidder’s reported value, and the losing bidders will pay nothing. In the second case, where only one of the bidders’ reported values is greater than the clamped value of $\frac{1}{2}$, the highest bidder will win the auction but will pay the clamped value ($\frac{1}{2}$) instead of the second highest bidder’s reported value. The losing bidders will pay nothing. Finally, in the third case where all bidders’ reported values are below the clamped value of

$\frac{1}{2}$, the winning bidder will pay the clamped value ($\frac{1}{2}$), and each losing bidder will pay the difference between $\frac{1}{2}$ and the highest bidder's reported value. An example of each case is given in Table 1 below.

Case	Bidder 1's Bid	Bidder 2's Bid	Bidder 3's Bid	Winner	Bidder 1's Payment	Bidder 2's Payment	Bidder 3's Payment
1	0.48	0.79	0.70	Bidder 2	0.00	0.70	0.00
2	0.48	0.39	0.70	Bidder 3	0.00	0.00	0.50
3	0.48	0.39	0.40	Bidder 1	0.50	0.02	0.02

III. Study One

To determine which of two auction mechanisms is more learnable for human subjects – the standard or clamped second price auction mechanism – we study the behavior of subjects who participated in a series of 100 to 150 sequential three-player auctions governed by one of these two mechanisms. In both conditions players were told exactly how the winners and payoffs in each auction would be determined. In each auction, every player was told her private value for an imaginary good being auctioned off. Then each player was asked to submit a bid for the good being auctioned off. After all players had submitted their bids, a winner was announced, and all players learned their payoffs for the auction. This procedure was repeated for 100 or 150 successive auctions depending on the treatment condition. Players bid against changing, anonymous partners in each auction, giving auctions the characteristics of independent rather than repeated interactions.

A. Experimental Procedure

The experiment described in this section was run in the Computer Lab for Experimental Research (CLER) at Harvard Business School. Forty-two members of the standing CLER subject pool were recruited through advertisements in multiple Boston-area campus newspapers to participate in two experimental sessions. In the first session there were 24 participants, and in the second session there were 18 participants. Each participant was randomly assigned to one of two experimental conditions. Sessions lasted for 60 minutes, and procedures were identical across sessions. Players were paid based on their earnings in all of the auction games they participated in plus a base rate of US\$10. Incentive pay ranged from US\$0 to US\$18.

When players entered the laboratory they were randomly assigned to a computer terminal where they found a sheet of paper giving them instructions on the auction game in which they would be participating. All of the players received instructions explaining that they would be participating in a series of three-player auctions against changing opponents. These instructions explained that players would learn their point value for an imaginary good being auctioned off before each auction and that their point value would be based on a random draw from a uniform distribution over the interval 60 to 100. The instructions also detailed how their earnings would depend on their bid and the bids of other players. Half of the players received descriptions of how winners and payoffs would be determined using a standard second price auction rule (see Appendix A). The other half received descriptions of how winners and payoffs would be determined using a clamped second price auction rule (see Appendix B). After viewing several examples of how payoffs and winners would be determined given different sets of bids received by the auction mechanism, players were asked to fill out a comprehension check demonstrating that they understood how winners and payoffs would be determined in each auction.

Subjects then participated in a series of auction games with one another by interacting with software on their computer terminals.⁴ During each auction, subjects learned their value for the imaginary item being auctioned off, entered their bid, learned if they had won the auction (and if so what they had paid for the item being auctioned off), learned their payoff for the round, and learned their cumulative earnings for the series of auctions they had participated in (see Appendix C for game screenshots). Payoffs were determined in each auction according to the relevant auction rule, and each point earned was converted into \$0.03 of incentive pay. All players in the standard second price auction condition participated in 150 sequential auctions as did players who participated in the clamped second price auction condition in the second experimental session. However, due to a 60 minute time limit on each experimental session, the 12 players in the clamped second price auction condition during our first experimental session only participated in 100 sequential auctions.

⁴ The experiment was programmed and conducted with the software z-Tree (Fischbacher 2007).

B. Results

An initial examination of the auction data from our study revealed that the vast majority of learning took place during the first 20 auctions, with subjects settling into relatively stable patterns of play for the remaining 80 to 130 auctions. This is consistent with learning rates reported by Ariely et al. (2005) in various auction environments. As a result, we will only present an analysis of learning during the first 20 auction games subjects participated in during our experiment.

The outcome variable of interest to us is the inefficiency of a player's bid, or the distance between her actual bid and the bid she would have placed had she bid her true value (following her optimal strategy). We calculate an inefficiency score for each bid as follows:

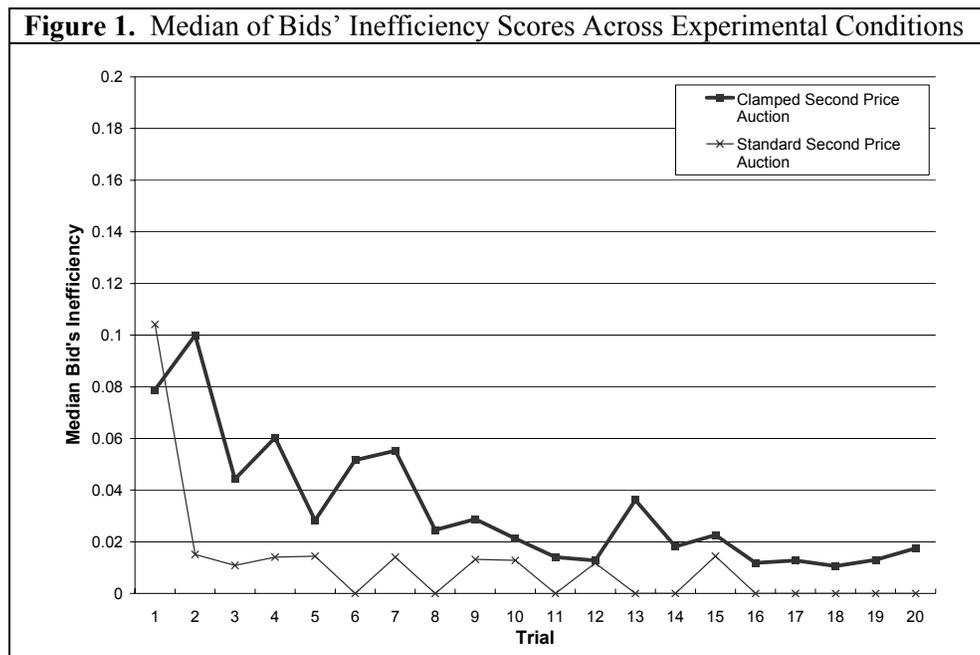
$$\text{inefficiency} = \left| 1 - \frac{\text{actual_bid}}{\text{true_value}} \right| \quad (1)$$

According to this measure of inefficiency, if a player bids optimally in an auction, her inefficiency score will be zero. The further her bid is from optimal, the larger her inefficiency score will become.

Because of our small sample of subjects (21 in each condition), extreme outliers exerted considerable pull on the average inefficiency of bids in each round. During the learning and experimentation process, some players submitted bids of zero while others submitted bids of up to twice their true value for an item. These bids were extreme outliers, yielding inefficiency scores of 1 in rounds where the median inefficiency score ranged between 0.05 and 0.02. To prevent outliers from dramatically altering the interpretation of our results, we focus our attention on the inefficiency score of the median bid in each round of auctions.

Figure 1 plots the inefficiency score of the median bid in each of the first 20 rounds of auctions in our two experimental conditions. The learning trajectories plotted on this graph suggest that players converge on their optimal bidding strategies considerably faster in auctions with a standard second price mechanism than in auctions with a clamped second price mechanism. We run a number of statistical tests whose results offer strong support for this conclusion. First, we find that the proportion of rounds in which the inefficiency of the median bid in the standard condition exceeds that in the clamped condition

(18/20) is significantly greater than the proportion of rounds in which the inefficiency of the median bid in the clamped condition exceeds that in the standard condition (2/20) (applying a binomial probability test, p -value < 0.001). Second, we find that the clamped second price auction yields significantly (p -value < 0.1) more inefficient median bids than the standard second price auction in 7 of the 20 rounds (applying a non-parametric K-sample test on the equality of medians), while the standard second price auction never yields a significantly more inefficient median bid than the clamped second price auction. Finally, we find that the clamped second price condition yields more inefficient median bids on average than the standard second price auction condition (applying a nonparametric K-sample test on the equality of medians yields significance at greater than the 0.1% level).



In addition to finding that players in the standard second price auction condition learn faster than players in the clamped second price auction condition, we also find that players in the clamped auction condition take more time to place their bids, suggesting that their slow learning may be more cognitively costly than the faster learning subjects exhibit in a standard second price auction. Players in the clamped second price auction condition spend an average of 8.1 seconds deciding what bid to place during each of the first 20 auctions they participate in, while players in the standard second price auction condition spend

significantly less time deciding what bid to place during the first 20 auctions they participate in – just 5.8 seconds on average (applying a t-test, p -value < 0.001).

C. Discussion

The evidence presented above indicates that contrary to the predictions made by the theory of collectives and earlier results on learnable mechanism design, a standard second price auction mechanism helps players learn to play their optimal strategies in auction games faster than a clamped second price auction mechanism. Not only do we find that this is the case, but we also find evidence suggesting that the slow learning we observe under the clamped second price auction mechanism requires more cognitive effort than the faster learning induced by the standard second price auction mechanism.

One concern about our findings is that our subjects may have been familiar with the standard second price auction mechanism before participating in our experiment but unfamiliar with the clamped second price auction mechanism, leading them to perform better in the standard second price auction condition. However, if this were the case, we would expect to see less inefficiency in the first round bids of players in the standard second price auction game than we see in the first round bids of players in the clamped second price auction game. The average and median inefficiencies of the first bids across conditions are, in fact, statistically indistinguishable (see Table 2 and Figure 1), suggesting that this is an unlikely explanation for our results. In order to address this concern, however, we run a second study in which subjects are not given information about how payoffs will be calculated in a series of auction games. This allows us to ensure the results from our first study were not driven by subjects' familiarity with the optimal bidding strategy in a standard second price auction.

	Clamped Second Price Auction	Standard Second Price Auction
Mean Inefficiency Score of Bids	0.12	0.10
Median Bid's Inefficiency Score	0.08	0.10

This table reports summary statistics about the inefficiency scores of subjects' first bids in our two experimental conditions. For each condition, we calculate the average inefficiency scores of subjects' bids and the median bid's inefficiency score in the first auction game. Applying a t-test to compare the mean inefficiency scores and a non-parametric K-sample test on the equality of medians to compare the median bids' inefficiency scores, we find no statistically distinguishable differences at an alpha-level of 0.1.

IV. Study Two

The design of our second study was identical to the design of our first except that subjects were told that they would participate in a series of 75 three-player auction games (as opposed to 150 games) in which the highest bidder would win the auction, and subjects were not given any information about how their payoffs would be determined in these auction.

A. *Experimental Procedure*

This experiment was again run in the CLER at Harvard Business School. Eighteen additional members of the standing CLER subject pool were recruited to participate in a single experimental session. The 18 subjects who participated in this 60 minute study were each randomly assigned to one of two experimental conditions – the standard second price auction condition or the clamped second price auction condition. Players were paid based on their earnings in all of the auction games they participated in plus a base rate of US\$10. Incentive pay ranged from US\$1 to US\$14.

As in study one, when players entered the laboratory they were randomly assigned to a computer terminal where they found a sheet of paper giving them instructions about the auction game they would be participating in. All of the players received instructions explaining that they would be participating in a series of three-player auctions against changing opponents. These instructions explained that players would learn their point value for an imaginary good being auctioned off before each auction and that their point value would be based on a random draw from a uniform distribution over the interval 60 to 100. Subjects in the two conditions received nearly identical instructions. The only difference was that subjects in the second price auction condition did not see any negative payoffs in an example demonstrating how payoffs would accumulate across rounds (see Appendix D) while subjects in the clamped second price auction condition did see one negative payoff example (see Appendix E).

As in study one, subjects then participated in a series of auction games with one another by interacting with software on their computer terminals.⁵ During each auction, subjects learned their value for the imaginary item being auctioned off, entered their bid, learned if they had won the auction (and if

⁵ The experiment was programmed and conducted with the software z-Tree (Fischbacher 2007).

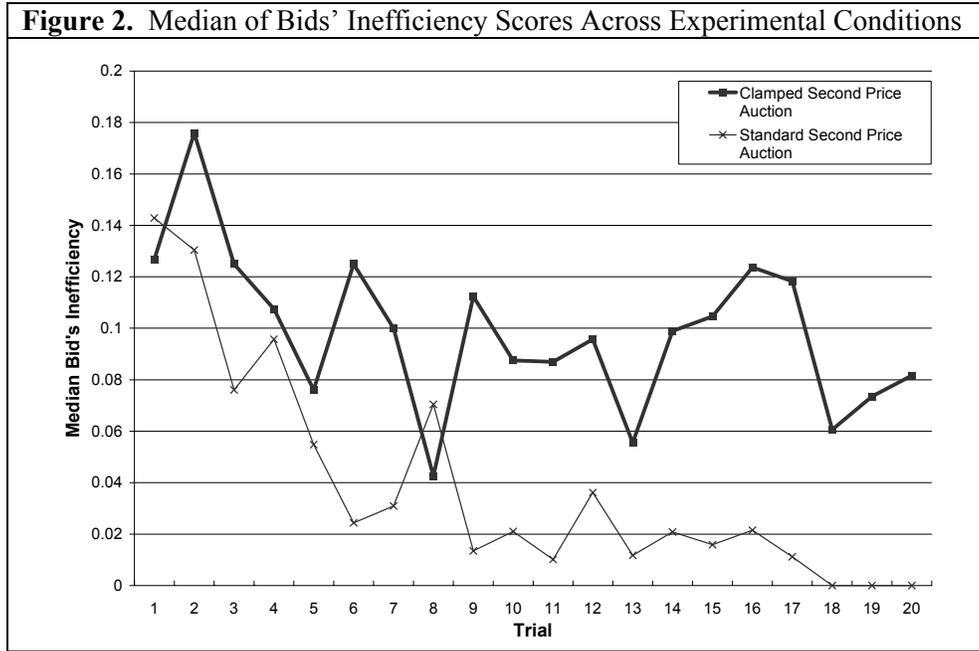
so what they had paid for the item being auctioned off), learned their payoff for the round, and learned their cumulative earnings for the series of auctions they had participated in (see Appendix F for game screenshots). Payoffs were determined in each auction according to the relevant auction rule, and each point earned was converted into \$0.04 of incentive pay. All players in both experimental conditions participated in 75 sequential auctions.

B. Results

As in study one, an initial examination of the bidding data from this study revealed that the vast majority of learning took place during the first 20 auctions, so we will only present an analysis of learning during the first 20 rounds of auction games in which subjects participated. We again evaluate the inefficiency (defined in Section III.B, Formula (1)) of bids in each round of each treatment condition. Our small sample of subjects (9 in each condition) again led us to examine the median inefficiency scores of the bids placed by subjects in each auction round in order to prevent outliers from exerting undue influence on our results.

Figure 2 plots the inefficiency score of the median bid in each of the first 20 rounds of auctions in our two experimental conditions. As in study one, the learning trajectories plotted on this graph suggest that players learn to bid optimally considerably faster in auctions with a standard second price mechanism than in auctions with a clamped second price mechanism. Also as in study one, we run a number of statistical tests whose results offer strong support for this conclusion. We again find that the proportion of rounds in which the inefficiency of the median bid in the standard condition exceeds that in the clamped condition (18/20) is significantly greater than the proportion of rounds in which the inefficiency of the median bid in the clamped condition exceeds that in the standard condition (2/20) (applying a binomial probability test, $p\text{-value} < 0.001$). In addition, we find that the clamped second price auction yields significantly ($p\text{-value} < 0.05$) more inefficient median bids than the standard second price auction in 9 of the 20 rounds (applying a non-parametric K-sample test on the equality of medians), while the standard second price auction never yields a significantly more inefficient median bid than the clamped second price auction. Finally, we find that the clamped second price condition yields more inefficient

median bids on average than the standard second price auction condition (applying a nonparametric K-sample test on the equality of medians yields significance at greater than the 0.1% level).



Also as in study one, we find that subjects in the clamped second price auction condition spend significantly longer (7.6 seconds on average) deciding what bid to place during each of the first 20 rounds of auction games than subjects in the standard second price auction condition (5.0 seconds on average) (applying a t-test, p-value < 0.001).

C. Discussion

This study provides additional evidence supporting our contention, based on the results of study one, that a standard second price auction mechanism helps players learn to play their optimal strategies in auction games faster than a clamped second price auction mechanism. In this study, we take steps to eliminate the possible confound in study one that subjects are more likely to be familiar with the rules and associated optimal strategy in a standard second price auction than in a clamped second price auction.⁶ In addition, by providing all subjects in this study with essentially identical instructions we eliminate the

⁶ It should be noted that the software players viewed in this study did label the auction game differently in the two conditions, calling it a “second price auction game” in one condition and a “clamped second price auction game” in the other. However, we argue that this did not provide subjects with enough information to infer anything about how their payoffs were being calculated.

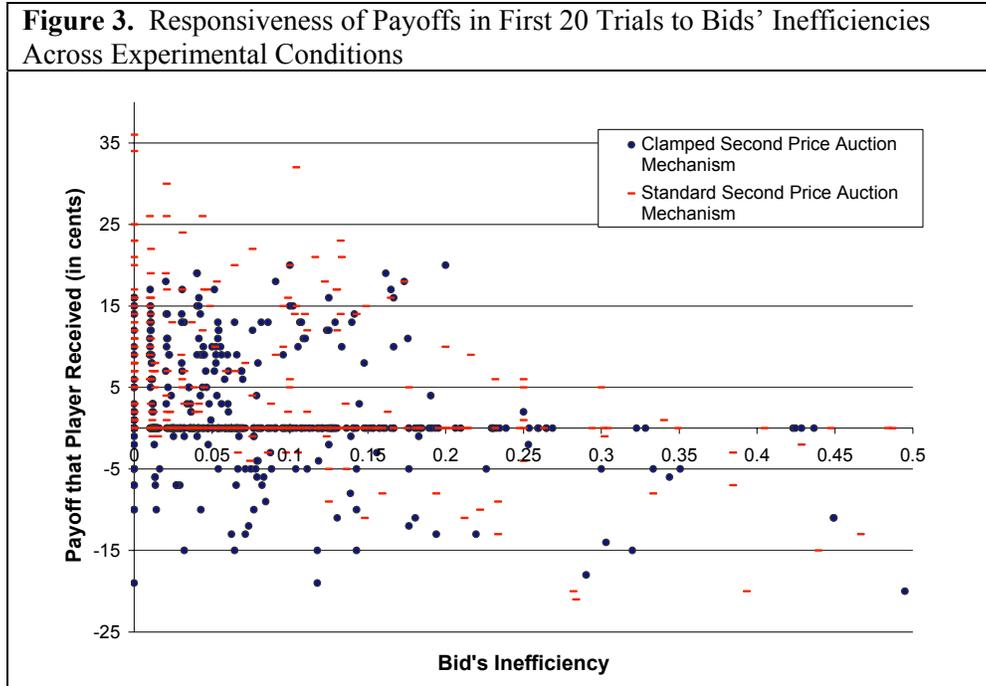
possibility that our study one results were caused by confusing instructions for the clamped second price auction game. Finally, this study demonstrates that even if subjects find it more difficult to calculate their optimal strategy in the clamped second price auction game than in the standard second price auction game based on the payoff rules described in study one, their slower learning in the clamped auction condition is a product of more than just this difficulty.

V. General Discussion and Conclusion

In this paper, we present evidence indicating that the clamped second price auction mechanism based on the theory of collectives actually induces *slower* learning in human subjects than the standard second price auction mechanism despite the fact that difference utilities provide faster learning with simulated reinforcement learning agents across many domains (Agogino and Tumer, 2005; Tumer and Agogino, 2007) including simulated auctions (Parkes 2004). This finding is surprising both in light of past simulation results and in light of the theoretical advantages that the theory of collectives suggests the clamped mechanism should provide in improving the responsiveness of payoffs to individuals' actions in multiagent learning settings. It seems important to discuss what the data we obtained from our experiments with human subjects can tell us about why both the theory of collectives and simulations with reinforcement learning agents made predictions about the performance of clamped versus second price auction mechanisms that contradicted our experimental results.

The theory of collectives suggests that the clamped second price auction mechanism should provide agents with payoffs that give them a more informative signal about the quality of their strategy choice than the standard second price auction mechanism, which should lead to faster learning. In fact, players in our experiments received payoffs that gave them a no more informative signal based on their strategy choices in the first twenty rounds of clamped auctions than in the first twenty rounds of standard auction games (see Figure 3). We measure "informativeness" here as the correlation between the inefficiency score of a bid and the payoff a player received.⁷

⁷ Across studies, examining the first twenty rounds of auction games, the correlation between the inefficiency score of a bid in a given auction and the payoff a player received in that auction was -0.17 (p-value < 0.0001) in the



Looking in more detail, we observe that the clamped auction provides more informative feedback to bidders when they *underbid* than the standard auction.⁸ However, the clamped auction provides considerably less informative feedback to bidders than the standard auction when they *overbid* in the first ten rounds of experimental auctions across conditions.⁹ In addition, as Table 3 illustrates, we observe markedly lower rates of underbidding and higher rates of overbidding in the standard auction than in the clamped auction across conditions. Overall, we see that the clamped auction provides better feedback than the standard auction when players underbid but worse feedback when they overbid, which coupled with higher rates of overbidding in the standard auction provides a potential explanation for our finding

standard second price auction and -0.17 ($p\text{-value} < 0.0001$) in the clamped second price auction. Examining only the first ten rounds of auction games, the correlation between the inefficiency score of a bid in a given auction and the payoff a player received in that auction was actually slightly higher in the standard second price auction (correlation = -0.21 , $p\text{-value} < 0.001$) than in the clamped second price auction (correlation = -0.20 , $p\text{-value} < 0.001$).

⁸ In both the first ten and the first twenty auction games, examining only bids that were below a bidder's true value, we find that the correlation between the inefficiency score of a bid in a given auction and the payoff a player received in that auction is higher in the clamped second price auction (correlation_{20rounds} = -0.20 , $p\text{-value} < 0.001$; correlation_{10rounds} = -0.25 , $p\text{-value} < 0.001$) than in the standard second price auction (correlation_{20rounds} = -0.11 , $p\text{-value} < 0.1$; correlation_{10rounds} = -0.07 , $p\text{-value} > 0.1$).

⁹ Examining only bids that were above a bidder's true value, we find that the correlation between the inefficiency score of a bid in a given auction and the payoff a player received is higher in the standard auction (correlation_{10rounds} = -0.35 , $p\text{-value} < 0.01$) than in the clamped auction (correlation_{10rounds} = -0.15 , $p\text{-value} > 0.1$). Looking at the first 20 rounds, we find in comparison that the feedback for overbids seems equally informative in the clamped auction (correlation_{20rounds} = -0.26 , $p\text{-value} < 0.01$) and in the standard auction (correlation_{20rounds} = -0.26 , $p\text{-value} < 0.001$).

that players learn to play their optimal strategies faster in the standard auction than in the clamped auction.

Table 3
PERCENT OF PLAYERS OVERBIDDING AND UNDERBIDDING BY CONDITION

	Instructions		No Instructions	
	Overbidding	Underbidding	Overbidding	Underbidding
Standard Second Price Auction	34%	19%	22%	47%
Clamped Second Price Auction	27%	50%	14%	83%

Like the theory of collectives, past simulations with reinforcement learning agents made the prediction that the clamped second price auction mechanism would lead people to learn their optimal strategies faster than the standard second price auction mechanism (Parkes, 2004). However, reinforcement agents in these simulations were assumed to exhibit certain behaviors in the course of learning that real agents in our experiments did not exhibit, which may explain the discrepancy between these results and the results of our laboratory experiments. First, the initial propensities of agents in the laboratory to select different strategies were distributed differently than the initial propensities assigned to agents in simulations.¹⁰ Second, a negative payoff led subjects across conditions in our experiment to lower their bid-to-value ratio, regardless of whether the negative payoff resulted from overbidding.¹¹ Simulated agents, however, are designed to adjust their bid-to-value ratio upwards if a negative payoff results from underbidding and downwards if a negative payoff results from overbidding. These are not the only differences between the learning behaviors we observed in human subjects and the learning behaviors exhibited by simulated agents, but they are the differences we consider to be particularly likely culprits for driving a wedge between simulated agent behavior and observed human behavior.

¹⁰ Among human players across conditions, underbidding was the initial strategy adopted by the overwhelming majority of subjects ($\text{underbid_percentage}_{1st_bid,clamped} = 70\%$, $\text{underbid_percentage}_{1st_bid,standard} = 77\%$). However, reinforcement learners are assigned equal initial propensities to adopt an overbidding or underbidding strategy (see Parkes, 2004).

¹¹ In the clamped second price auction, 58% of bids that resulted in negative payoffs were associated with bid-to-value ratios of one or less, and yet 58% of the adjustments in bid-to-value ratio resulting from these negative payoffs were downward adjustments. In the standard second price auction, 100% of the bids that resulted in negative payoffs were associated with bid-to-value ratios of one or more (by design), and 80% of the adjustments in bid-to-value ratio resulting from these negative payoffs were downward adjustments. Thus, the apparent human tendency to adjust downward in response to a loss is helpful to players in the standard second price auction (and rational) but harmful to players in the clamped second price auction (and apparently irrational).

Turning back to the implications of the results of our two studies, we conclude that while creating more learnable mechanisms is an important goal for academics to pursue, the approach taken by Tumer et al. (2002) in the context of learning in multiagent systems does not appear to offer useful mechanism design insights in the domain of small auctions intended for human participants. However, it has been demonstrated that the new type of mechanism tested in this paper has desirable learning properties for artificial agents, and it may be the case that mechanisms based on the theory of collectives (Wolpert and Tumer, 2001) have desirable learning properties for human agents in non-auction settings or for human agents competing in auction games with different properties than those tested (e.g., more than three participants or with a different clamping value). Given our results, however, and given the design difficulties associated with implementing mechanisms based on the theory of collectives (Wolpert and Tumer, 2001), such as the need for a means of imposing fines on mechanism participants, we remain somewhat pessimistic about finding useful applications in settings with human participants. Our results also illustrate the importance of testing economic theories with human subjects rather than relying on the results of simulations with reinforcement learners to predict human behavior, as human behavior and the behavior of simulated reinforcement learners can differ.¹² We hope our study will encourage future research on mechanisms designed to increase the speed at which human subjects learn to play their optimal strategies and on how simulated learners can be designed to better mimic the behaviors of human learners.

¹² For those who are familiar with the RELACS learning model (Erev and Barron, 2005), it is worth noting that we conducted simulations to compare the speed at which RELACS learners converge to optimal bidding behavior in a standard second price auction versus a clamped second price auction and found that the clamped auction mechanism induces faster learning in RELACS agents than the standard second price auction mechanism. These simulation results, like those discussed by Parkes (2004), point to the continued need for experimentation with human subjects when testing theories about how people will learn when faced with a given set of stimuli.

Bibliography

- Agogino, A. and K. Tumer (2005). "Multiagent Reward Analysis for Learning in Noisy Domains," in *Proc. AAMAS'05*.
- Ariely, D., A. Ockenfels, and A. E. Roth (2005). "An Experimental Analysis of Ending Rules in Internet Auctions," *Rand Journal of Economics*, 36(4), 891-908.
- Armstrong, M. (2000). "Optimal Multi-Object Auctions," *Review of Economic Studies*, 67(3), 455-481.
- Ausubel, L. (2004) "An Efficient Ascending-Bid Auction for Multiple Objects," *American Economic Review*, 94(5), 1452-1475.
- Borgs, C., J. Chayes, O. Etesami, N. Immorlica, K. Jain, and M. Mahdian (2007). "Dynamics of Bid Optimization in Online Advertisement Auctions," in *Proc. WWW2007*, 531 – 540.
- Cox, J., B. Robertson and V. Smith (1982) "Theory and Behavior in Single Object Auctions." in *Research in Experimental Auctions*, V. Smith (ed.), JAI Press: Greenwich, CT.
- Dasgupta, P. and E. Maskin (2000). "Efficient Auctions," *Quarterly Journal of Economics*, 115(2), 341-389.
- Edelman, B. and M. Ostrovsky (2007). "Strategic Bidder Behavior in Sponsored Search Auctions," *Decision Support Systems*, 43(1), 192-198.
- Edelman, B., M. Ostrovsky, and M. Schwarz (2007). "Internet Advertising and the Generalized Second Price Auction: Selling Billions of Dollars Worth of Keywords," *American Economic Review*, 97(1), 242-259.
- Erev, I. and G. Barron (2005). "On Adaptation, Maximization, and Reinforcement Learning Among Cognitive Strategies," *Psychological Review*, 112(4), 912-931.
- Feng, J. and X. Zhang (2007). "Dynamic Price Competition on the Internet: Advertising Auctions," in *Proc. ACM Conference on Electronic Commerce '07*, 57-58.
- Fischbacher, U. (2007). "z-Tree: Zurich Toolbox for Ready-made Economic Experiments," *Experimental Economics*, 10(2), 171-178.
- Groves, T. (1973). "Incentives in Teams," *Econometrica*, 41, 617-631.
- Kagel, J.H., R. M. Harstad, and D. Levin. (1987). "Information Impact and Allocation Rules in Auctions with Affiliated Private Values: A Laboratory Study," *Econometrica*, 55, 1275-1304.
- Kagel, J.H. and D. Levin (1993). "Independent Private Value Auctions: Bidder Behavior in First-, Second-, and Third-Price Auctions with Varying Numbers of Bidders," *Economic Journal*, 103, 868-879.
- Kagel, J.H. (1995) "Auctions: A Survey of Experimental Research," in *Handbook of Experimental Economics*, J. Kagel and A.E. Roth (eds.), Princeton University Press: Princeton, NJ.
- Mathvet, L. (2007). "Supermodular Bayesian Implementation: Learning and Incentive Design," *Caltech Working Paper*.

- Milgrom P. and J. Roberts (1990). "Rationalizability Learning, and Equilibrium in Games with Strategic Complementarities," *Econometrica*, 58, 1255-1277.
- Milgrom, P. and R. Weber (1982). "A Theory of Auctions and Competitive Bidding," *Econometrica*, 50, 1089-1122.
- Myerson, R. (1981). "Optimal Auction Design." *Mathematics of Operations Research*, 6, 58-73.
- Parkes, D. C. (2004). "On Learnable Mechanism Design," in *Collectives and the Design of Complex Systems*, K. Tumer and D.H. Wolpert (eds.), Springer-Verlag: New York, NY.
- Tumer, K., Agogino, A.K., and D. Wolpert (2002). "Learning Sequences of Actions in Collectives of Autonomous Agents," in *Proc. AAMAS'02*.
- Tumer, K. and A.K. Agogino (2005). "Coordinating Multi-Rover Systems: Evaluation Functions for Dynamic and Noisy Environments," in *Proc. GECCO 2005*.
- Tumer, K. and A.K. Agogino (2007). "Distributed Agent-Based Air Traffic Flow Management," in *Proc. Sixth International Joint Conference on Autonomous Agents and Multi-Agent Systems '07*.
- Tumer, K. and D. Wolpert (2004). *Collectives and the Design of Complex Systems*. Springer-Verlag: New York, NY.
- Varian, H. (in press). "Position Auctions," *International Journal of Industrial Organization*.
- Vickrey, W. (1961). "Counterspeculation, Auctions, and Competitive Sealed Tenders," *Journal of Finance*, 16, 8-37.
- Wolpert, D. and K. Tumer (2001). "Optimal Payoff Functions for Members of Collectives," *Advances in Complex Systems*, 4, 265-279.
- Wolpert, D. and K. Tumer (2002) "Collective Intelligence, Data Routing and Braess' Paradox," *Journal of Artificial Intelligence Research*, 16, 359-387.

Appendix A: Instructions Received By Subjects in Study One, Standard Condition

A Study of Decision Making in Auction Games

The purpose of this session is to study decision making in an auction market. In this study you will participate in a series of auctions. In each auction, there will be exactly three bidders, you and two others. You will not be paired with the same bidder from one auction to the next. You will be randomly re-matched to bid against two other bidders in each auction. All bidders receive exactly the same instructions. All bidders will have to wait until everyone else has placed their bid in each round before they will be able to proceed to the next auction. Your earnings will be tallied in points, which can be traded in for dollars at the end of the experiment at an exchange rate of: 100 points = \$3.00.

What is auctioned off and what are my payoffs?

The objects for which you will be bidding in these auctions will be imaginary. You will be informed at the outset of the auction how much the object for sale is worth to you. This amount, called your personal value, v , will be randomly selected from a uniform distribution between 60 and 100 points. At the end of the auction, if you won the auction and paid a price, p , (more on how this price will be determined later), you will have earned $v - p$ points.

If you do not win the auction, you will earn nothing for that auction.

If your earnings are negative, they represent a loss. Gains are added to and losses are subtracted from the show-up fee and paid to you in cash at the end of the experiment.

Here are three examples of auction outcomes. Each row contains a separate game.

Example	You win the auction	Your personal value	Amount paid	Your cumulative earnings
1	Yes	90 points	80 points	10 points
2	No	55 points	0 points	10 points
3	Yes	85 points	82 points	13 points

How can I bid?

The rules of the auctions that you will play today are as follows: The player with the highest bid wins the auction and pays a price equal to the second highest bid. If multiple bidders submit the same exact winning bid, one of the bids is chosen randomly to be the winning bid.

The price paid by losing bidders is 0 points.

Here are four examples of auction outcomes. Each row contains a separate game.

Example	Bidder 1's bid	Bidder 2's bid	Bidder 3's bid	Highest bidder	Price paid by winning bidder	Price paid by losing bidders
1	95 points	78 points	70 points	1	78 points	0 points
2	67 points	88 points	99 points	3	88 points	0 points
3	60 points	89 points	89 points	2 or 3	89 points	0 points
4	60 points	63 points	71 points	3	63 points	0 points

So, your task is to submit bids for the item in competition with the other two bidders. The price you will pay if you win the auction will not necessarily be the same as your bid. However, the price will never exceed your bid. Once a bid is submitted, the software will determine the winning bidder and your earnings for the round. While your bid and earnings are kept secret, the price of the good will be revealed to the winning bidder.

In each auction you will be given one and only one opportunity to submit your bid for each object. Once all participants submit their bids, the highest bidder will win the auction for the price set by the rules described above. The winner will earn a number of points for that auction equal to her value for the object minus its price. At that point, the next auction will start.

Comprehension Check:

Please fill in the blanks on the following chart:

Case	Bidder 1's bid	Bidder 2's bid	Bidder 3's bid	Price paid by winning bidder	Price paid by losing bidders
1	85 points	95 points	70 points		
2	68 points	72 points	77 points		
3	95 points	64 points	75 points		

Please Enter Your Computer Terminal Number Here: _____

Appendix B: Instructions Received By Subjects in Study One, Clamped Condition

A Study of Decision Making in Auction Games

The purpose of this session is to study decision making in an auction market. In this study you will participate in a series of auctions. In each auction, there will be exactly three bidders, you and two others. You will not be paired with the same bidder from one auction to the next. You will be randomly re-matched to bid against two other bidders in each auction. All bidders receive exactly the same instructions. All bidders will have to wait until everyone else has placed their bid in each round before they will be able to proceed to the next auction. Your earnings will be tallied in points, which can be traded in for dollars at the end of the experiment at an exchange rate of: 100 points = \$3.00.

What is auctioned off and what are my payoffs?

The objects for which you will be bidding in these auctions will be imaginary. You will be informed at the outset of the auction how much the object for sale is worth to you. This amount, called your personal value, v , will be randomly selected from a uniform distribution between 60 and 100 points. At the end of the auction, if you won the auction and paid a price, p , (more on how this price will be determined later), you will have earned $v - p$ points.

If you do not win the auction and if the winning bid, w , is greater than 80, you will earn nothing for that auction. However, if the winning bid is less than 80, you will have lost the difference between the winning bid and 80 points.

If your earnings are negative, they represent a loss. Gains are added to and losses are subtracted from the show-up fee and paid to you in cash at the end of the experiment.

Here are three examples of auction outcomes. Each row contains a separate game.

Example	You win the auction	Your personal value	Amount paid	Your cumulative earnings
1	Yes	90 points	80 points	10 points
2	No	55 points	2 points	8 points
3	Yes	85 points	82 points	11 points

How can I bid?

The rules of the auctions that you will play today are as follows: The player with the highest bid wins the auction. The winner pays a price equal to the maximum of the second highest bid and 80 points. In other words, if you are the highest bidder for one of the objects being auctioned, the amount you will pay will be equal to the second highest bid if it is greater than 80 points. Otherwise, the amount you pay will be 80 points. If multiple bidders submit the same exact winning bid, one of the bids is chosen randomly to be the winning bid.

The price paid by losing bidders is equal to 0 points if the winning bidder placed a bid, w , of more than 80 points. Otherwise, the price paid by losing bidders is equal to the difference between 80 points and the amount bid by the winning bidder, or $80 - w$ points.

Here are four examples of auction outcomes. Each row contains a separate game.

Example	Bidder 1's bid	Bidder 2's bid	Bidder 3's bid	Highest bidder	Price paid by winning bidder	Price paid by losing bidders
1	95 points	78 points	70 points	1	80 points	0 points
2	67 points	88 points	99 points	3	88 points	0 points
3	60 points	89 points	89 points	2 or 3	89 points	0 points
4	60 points	63 points	71 points	3	80 points	9 points

So, your task is to submit bids for the item in competition with the other two bidders. The price you will pay if you win the auction will not necessarily be the same as your bid. Once a bid is submitted, the software will determine the winning bidder and your earnings for the round. While your bid and earnings are kept secret, the price of the good will be revealed to the winning bidder.

In each auction you will be given one and only one opportunity to submit your bid for each object. Once all participants submit their bids, the highest bidder will win the auction for the price set by the rules described above. The winner will earn a number of points for that auction equal to her value for the object minus its price. At that point, the next auction will start.

Comprehension Check:

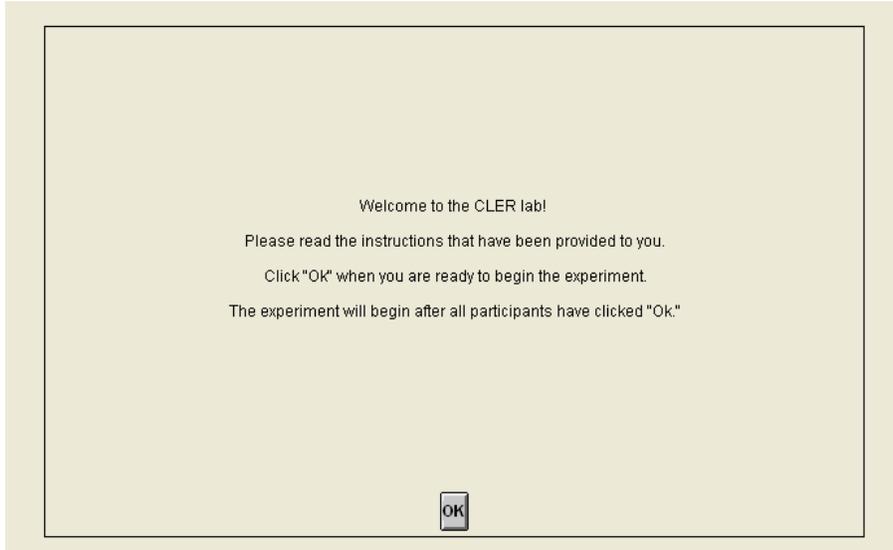
Please fill in the blanks on the following chart:

Case	Bidder 1's bid	Bidder 2's bid	Bidder 3's bid	Price paid by winning bidder	Price paid by losing bidders
1	85 points	95 points	70 points		
2	68 points	72 points	77 points		
3	95 points	64 points	75 points		

Please Enter Your Computer Terminal Number Here: _____

Appendix C: Screenshots from Study One

Welcome Screen

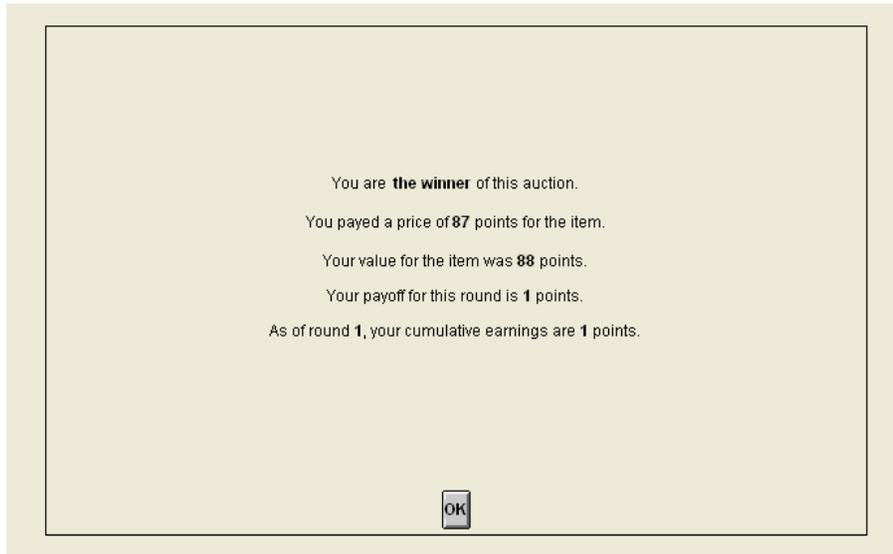


Screen Displayed at Outset of Each Auction Game¹³



¹³ Note that players in the clamped second price condition saw text saying they were playing in a “clamped second price auction game” instead of a “second price auction game.”

Screen Displayed to Winner at Conclusion of Each Auction Game



Screen Displayed to Loser at Conclusion of Each Auction Game



Screen Displayed After 150 Auction Games When the Experiment is Over



Appendix D: Instructions Received By Subjects in Study Two, Standard Condition

A Study of Decision Making in Auction Games

The purpose of this session is to study decision making in an auction market. In this study you will participate in a series of auctions. In each auction, there will be exactly three bidders, you and two others. You will not be paired with the same bidder from one auction to the next. You will be randomly re-matched to bid against two other bidders in each auction. All bidders receive exactly the same instructions. All bidders will have to wait until everyone else has placed their bid in each round before they will be able to proceed to the next auction. Your earnings will be tallied in points, which can be traded in for dollars at the end of the experiment at an exchange rate of: 100 points = \$4.00.

What is auctioned off and what are my payoffs?

The objects for which you will be bidding in these auctions will be imaginary. You will be informed at the outset of the auction how much the object for sale is worth to you. This amount, called your personal value, v , will be randomly selected from a uniform distribution between 60 and 100 points. Each bidder will have her own personal value. At the end of the auction, if you won the auction and paid a price, p , you will have earned $v - p$ points.

If your earnings are negative, they represent a loss. Gains are added to and losses are subtracted from the show-up fee and paid to you in cash at the end of the experiment.

Here are three examples of auction outcomes. Each row contains a separate game.

Example	You win the auction	Your personal value	Amount paid	Your earnings for this auction
1	Yes	90 points	80 points	10 points
2	No	55 points	0 points	0 points
3	Yes	85 points	82 points	3 points

How can I bid?

The rules of the auctions that you will play today are as follows: The player with the highest bid wins the auction and pays a price that depends on the bids of the other auction participants. If multiple bidders submit the same exact winning bid, one of the bids is chosen randomly to be the winning bid.

So, your task is to submit bids for the item in competition with the other two bidders. Once a bid is submitted, the software will determine the winning bidder and your earnings for the round. While your bid and earnings are kept secret, the price of the good will be revealed to all auction participants.

In each auction you will be given one and only one opportunity to submit a bid. Once all participants submit their bids, the auction outcome will be revealed. At that point, the next auction will start.

Please Enter Your Computer Terminal Number Here: _____

Appendix E: Instructions Received By Subjects in Study Two, Clamped Condition

A Study of Decision Making in Auction Games

The purpose of this session is to study decision making in an auction market. In this study you will participate in a series of auctions. In each auction, there will be exactly three bidders, you and two others. You will not be paired with the same bidder from one auction to the next. You will be randomly re-matched to bid against two other bidders in each auction. All bidders receive exactly the same instructions. All bidders will have to wait until everyone else has placed their bid in each round before they will be able to proceed to the next auction. Your earnings will be tallied in points, which can be traded in for dollars at the end of the experiment at an exchange rate of: 100 points = \$4.00.

What is auctioned off and what are my payoffs?

The objects for which you will be bidding in these auctions will be imaginary. You will be informed at the outset of the auction how much the object for sale is worth to you. This amount, called your personal value, v , will be randomly selected from a uniform distribution between 60 and 100 points. Each bidder will have her own personal value. At the end of the auction, if you won the auction and paid a price, p , you will have earned $v - p$ points.

If your earnings are negative, they represent a loss. Gains are added to and losses are subtracted from the show-up fee and paid to you in cash at the end of the experiment.

Here are three examples of auction outcomes. Each row contains a separate game.

Example	You win the auction	Your personal value	Amount paid	Your earnings for this auction
1	Yes	90 points	80 points	10 points
2	No	55 points	2 points	-2 points
3	Yes	85 points	82 points	3 points

How can I bid?

The rules of the auctions that you will play today are as follows: The player with the highest bid wins the auction and pays a price that depends on the bids of the other auction participants. If multiple bidders submit the same exact winning bid, one of the bids is chosen randomly to be the winning bid.

The price paid, if any, by the losing bidders depends on the bids of the other auction participants.

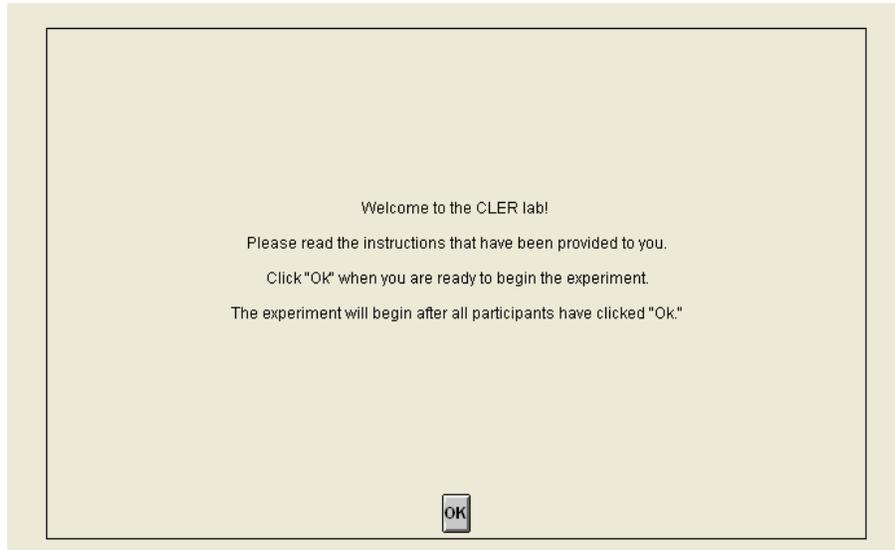
So, your task is to submit bids for the item in competition with the other two bidders. Once a bid is submitted, the software will determine the winning bidder and your earnings for the round. While your bid and earnings are kept secret, the price of the good will be revealed to all auction participants.

In each auction you will be given one and only one opportunity to submit a bid. Once all participants submit their bids, the auction outcome will be revealed. At that point, the next auction will start.

Please Enter Your Computer Terminal Number Here: _____

Appendix F: Screenshots from Study Two

Welcome Screen



Screen Displayed at Outset of Each Auction Game¹⁴

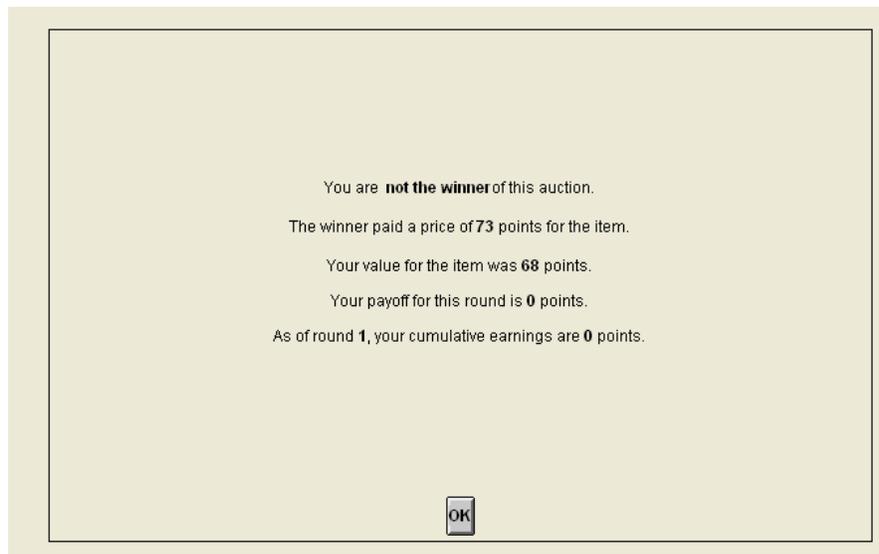


¹⁴ Note that players in the clamped second price condition saw text saying they were playing in a “clamped second price auction game” instead of a “second price auction game.”

Screen Displayed to Winner at Conclusion of Each Auction Game



Screen Displayed to Loser at Conclusion of Each Auction Game



Screen Displayed After 75 Auction Games When the Experiment is Over

