

## 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.



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



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



learnability as a design criterion.<sup>3</sup> 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

---

<sup>3</sup> 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.



### ***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:



$\frac{1}{2}$ , 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.



## 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



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.

	<b>Clamped Second Price Auction</b>	<b>Standard Second Price Auction</b>
<b>Mean Inefficiency Score of Bids</b>	0.12	0.10
<b>Median Bid's Inefficiency Score</b>	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.



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





































