Using Machine Learning for Modeling Human Behavior and Analyzing Friction in Generalized Second Price Auctions

Karthik Kannan *
kkarthik@purdue.edu

Vandith Pamuru *
vandith@purdue.edu

Yaroslav Rosokha *
yrosokha@purdue.edu

Abstract

Recent advances in technology have reduced frictions in various markets. In this research, we specifically investigate the role of frictions in determining the efficiency and bidding behavior in a generalized second price auction (GSP) – the most preferred mechanism for sponsored search advertisements. First, we simulate computational agents in the GSP setting and obtain predictions for the metrics of interest. Second, we test these predictions by conducting a human-subject experiment. We find that, contrary to the theoretical prediction, the lower-valued advertisers (who do not win the auction) substantially overbid. Moreover, we find that the presence of market frictions moderates this phenomenon and results in higher allocative efficiency. These results have implications for policymakers and auction platform managers in designing incentives for more efficient auctions.

*Krannert School of Management, Purdue University. Authors listed in alphabetical order of last name.

1 Introduction

The online ad market has seen spectacular growth in recent years. For example, the revenue generated by platforms such as Google, Bing, and Yahoo exceeded $92 billion in 2017 (Statista, 2017). Generally, these platforms are believed to have reduced the cost of participation and allow more advertisers to enter these markets. Furthermore, advances in automation and artificial-intelligence tools (e.g., auto bidders) have contributed to a substantial reduction in the cost of bid adjustments – i.e., friction – once the advertiser has entered the market. Even then, prior literature has not studied the role that such friction costs have on the sponsored-search-auction outcomes.

In this paper, we study the role of frictions in the outcome of the generalized second price auction (henceforth GSP). Seminal theoretical works such as Edelman, Ostrovsky, and Schwarz (2007) and Varian (2007) have provided a structured approach of analyzing this context. A few papers have focused on deviations from the basic theoretical model (e.g., Simonov, Nosko, and Rao, 2018; Jerath, Ma, Park, and Srinivasan, 2011). Yet, none have analyzed the role of frictions in explaining commonly observed deviations such as overbidding (e.g., Noti, Nisan, and Yaniv, 2014; McLaughlin and Friedman, 2016; Sheremeta, 2010; Cooper and Fang, 2008; Kamijo, 2013). In particular, our focus is on the friction associated with the bidding process, which has been decreasing with availability of advanced tools like auto-bidders, etc. Specifically, we study the following research questions: Do frictions lead to an increase or decrease in overbidding? What are the consequences for the auctioneer, for the advertiser, and for the allocative efficiency of the GSP?

These questions cannot be studied empirically. The secondary data do not have the advertisers' private valuations, so the identification can become a problem. Therefore, we study the above questions through a combination of computational and experimental approaches. Specifically, we first investigate bidding behaviors and allocative efficiency using a computational model involving reinforcement-learning agents. The computational model provides several testable predictions about the effect of friction costs. We subsequently validate these hypotheses using a human-subject experiment in a controlled laboratory setting. Finally, after establishing the validity of our computational model, we use simulations to provide additional insights into the role that frictions play in the markets that we cannot feasibly (or practically) investigate with human-subject experiments.

For our computational approach, we use a well-established model from the machine-learning literature that dates back to the 1980s. In particular, our computational agents implement a version of Q-learning (Sutton and Barto, 1998) in an environment akin to Edelman, Ostrovsky, and Schwarz (2007).[1] The market frictions are modeled as additional costs that the participants incur during their bid-adjustment process. We find the following key results through the machine-learning model: (a) Contrary to the theoretical models, the lowest-valued advertisers submit bids higher than their private valuation; (b) this overbidding phenomenon leads to allocative inefficiencies in the market; and (c) the allocative efficiency may increase as the frictions increase. We confirm these results experimentally. Thus, in some regards, our computational model may be perceived as a digital twin representation of advertisers in GSP.

The rest of the paper is organized as follows: In Section 2, we outline the previous literature that relates to our paper. In Section 3, we present details of the auction environment and develop the hypotheses using the agent-based model. In Section 4, we present the experimental design. In Section 5, we present the data and main results of the human-subject experiment. Finally, in Section 6, we summarize our research findings and discuss future research directions.

2 Literature

Our study contributes to three streams of literature. First, we contribute to the literature on sponsored search advertisement and auction mechanisms. Second, we contribute to a growing body

[1]Q-learning is extensively used in machine-learning and deep-learning applications, including Google's Deepmind solving AlphaGo.

of literature in Information Systems that uses economic experiments. And third, we contribute to the emerging literature that uses machine-learning models to study market outcomes. Next, we provide a brief review for each of the three streams of literature.

2.1 Sponsored Search Keyword Auctions

The sponsored search auctions have attracted tremendous interest in the IS literature. Several empirical studies have focused on the evolution of bidding strategies and the resulting impact on sponsored search metrics (e.g., Ghose and Yang, 2009; Animesh, Viswanathan, and Agarwal, 2011; Animesh, Ramachandran, and Viswanathan, 2010). The bidding strategies have also been the focus of Zhang and Feng (2011), who introduce a dynamic model to study the cyclic bidding by the advertisers. A number of theoretical papers have expanded on the works of Edelman, Ostrovsky, and Schwarz (2007) and Varian (2007) in order to improve the auction outcomes (e.g., Chen, Feng, and Whinston, 2010; Varian, 2009; Edelman and Schwarz, 2010; Amaldoss, Desai, and Shin, 2015) or evaluate alternative mechanisms (e.g., Feng, Bhargava, and Pennock, 2007). Though most of the works consider the auction for an individual keyword, some studies explore the bidding process for multiple keywords (e.g., Du, Su, Zhang, and Zheng, 2017) and the interaction between organic results and the sponsored search results as competing information sources (e.g., Agarwal, Hosanagar, and Smith, 2015; Xu, Chen, and Whinston, 2012). Furthermore, recent works have investigated more advanced variations of the auction environment including auctions with unknown click-through rates (Devanur and Kakade, 2009; Gatti, Lazaric, and Trovò, 2012), auctions with dependent click-through rates (Simonov, Nosko, and Rao, 2018; Kempe and Mahdian, 2008; Deng and Yu, 2009) and auctions with budget constraints (Arnon and Mansour, 2011; Zhou, Chakrabarty, and Lukose, 2008), etc. Qin, Chen, and Liu (2015) provide a comprehensive review of the sponsored search auction literature.

Our paper contributes to this vast literature by studying the GSP auction outcomes and bidding behavior in the presence of market frictions. In particular, the recent advances in communication technology have resulted in significantly reduced informational frictions in the sponsored search auctions. Further, the proliferation of AI tools (e.g., auto-bidders) have further contributed to reduction of frictions. Although frictions have been studied in different contexts including trade (Allen, 2014; Hou and Moskowitz, 2005), stock markets (Capasso, 2008), housing markets (Anenberg, 2016), and labor markets (Bassi and Nansamba, 2017), we aim to study the role of frictions in the context of sponsored search auctions. In particular, we incorporate these frictions into the model presented in Edelman, Ostrovsky, and Schwarz (2007). We then employ a twofold strategy of investigating the auction outcomes using machine-learning computational agents in combination with human-subject experiments. A brief review of literature on economic experiments and computational methods is next.

2.2 Economic Experiments in IS

The use of laboratory experiments to test theoretical insights and guide the design of systems has gained substantial traction in the IS field. Following the early work, which includes the development of the technology-acceptance model by Bagozzi, Davis, and Warshaw (1992) and the evaluation of the task-technology fit model by Goodhue and Thompson (1995), recent literature has expanded the use of experiments to a variety of IS applications, including privacy (e.g., Brandimarte, Acquisti, and Loewenstein, 2013; Tsai, Egelman, Cranor, and Acquisti, 2011), bundle pricing (e.g., Goh and Bockstedt, 2013), and recommender systems (e.g., Adomavicius, Bockstedt, Shawn, and Zhang, 2014; Adomavicius, Bockstedt, Curley, and Zhang, 2013). In a recent review of the experimental literature in IS, Gupta, Kannan, and Sanyal (2018) argue that experiments can yield meaningful results that overcome the limitations faced by empirical and theoretical studies. In the context of our study, to make any conclusions about allocative efficiency or bidding behavior, we must observe the private valuations of the auction advertisers, which is not possible using real-world data.

The experimental approach has been particularly fruitful in studying auctions (e.g., Sanyal, 2016; Adomavicius, Curley, Gupta, and Sanyal, 2013, 2012; Cason, Kannan, and Siebert, 2011; Adomavicius, Gupta, and Sanyal, 2006; Bapna, Barua, Mani, and Mehra, 2010). Several early experiments on behavior in auctions report that subjects rarely choose the dominant strategy (Kagel, Harstad, and Levin, 1987; Kagel, Levin, and Harstad, 1995; Kagel and Levin, 2001). In particular, the robust finding in this literature is that human subjects regularly bid higher than their value. Closest to our paper are four recent studies that investigate GSP auctions experimentally: Fukuda, Kamijo, Takeuchi, Masui, and Funaki (2013); Noti, Nisan, and Yaniv (2014); McLaughlin and Friedman (2016); Che, Choi, and Kim (2017). With the exception of McLaughlin and Friedman (2016), these studies find significant overbidding behavior by the participants. The primary question in all of these studies, however, is different from ours. Whereas prior studies have focused on the comparison of GSP to VCG mechanisms, we investigate the role of frictions in improving outcomes of the GSP auction. In particular, we focus on how the presence of market frictions may mitigate overbidding behavior and lead to higher allocative efficiency. In addition, the use of machine-learning agents in combination with human-subject experiments is a distinctive feature of this paper.

2.3 Agent-based Computational Models

In addition to theoretical, empirical, and experimental approaches, agent-based simulations have been successfully used to provide insights in the context of allocation problems. For example, Guo, Koehler, and Whinston (2012) analyze bundle trading markets for distributed resource allocations. Ketter, Collins, Gini, Gupta, and Schrater (2012) study trading agent competition in a supply-chain context, in which a need exists to make product-pricing and inventory-resource-allocation decisions in real time. In the context of auctions, Bichler, Shabalin, and Ziegler (2013) study the efficiency of combinatorial clock auctions, Kiose and Voudouris (2015) study power auctions, and Guerci, Kirman, and Moulet (2014) investigate sequential Dutch auctions.

To be able to replicate (imitate) human behavior using computer agents, the agents need to adopt a learning process. For our study, we employ Q-learning (Watkins and Dayan, 1992) to model the behavior of advertisers in the GSP environment. We chose this approach for several reasons. First, Q-learning has been used to investigate learning in multiple-agent environments (e.g., Littman, 1994; Sandholm and Crites, 1996; Bowling and Veloso, 2001; Greenwald, Hall, and Serrano, 2003). Second, Q-learning has been used to match behavioral regularities observed in human subject experiments (e.g., Rosokha and Younge, forthcoming). Third, Q-learning algorithms have been successfully applied to investigate the efficacy of information revelation and structural properties in a variety of auctions (Greenwald, Kannan, and Krishnan, 2010). Finally, the reinforcement-learning approach is not new to the GSP. Chen, Liu, and Yang (2016), establish a connection between machine-learned models and the game-theoretic properties of a system using real data from a sponsored-search-advertising platform.

3 Computational Analysis and Hypotheses Development

In this section, we first present the environment (Section 3.1); second, we present details of our implementation of the agent-based model (Section 3.2); and finally, we provide the results of our simulations and state the main hypotheses (Section 3.3).

3.1 Core GSP Model

To study the problem in a structured manner, it makes sense to build from a theoretical model. However, we are not aware of any universally accepted theoretical model in GSP as certain assumptions have been shown to be violated. Therefore, we consider a specialization of, arguably, the most well-known paper — Edelman, Ostrovsky, and Schwarz (2007) — and incorporate elements related to frictions. In particular, the GSP environment contains $J = 3$ advertisers, each with a unit demand, participating in an auction for $K = 2$ ad slots. Each advertiser submits only one bid. The advertisers may be placed in either of the slots (note that one advertiser will fail to appear in any of the slots). Let α_k represent the click-through rate for the k-th ad slot. Without the loss of generality, we assume $\alpha_1 = 1 > \alpha_2 = \alpha > \alpha_3 = 0$. The first slot is more desired and has a higher click-through rate than the second slot.[2] The third slot, which does not exist, is assumed to have a zero click-through rate for ease of representation. Thus, a higher value of α means that the two ad slots are more similar.

Conditional on the click-through, advertiser j realizes a value v^j and, without loss of generality, we assume $v^1 > v^2 > v^3$. In our implementation, we assume that the valuations are private information and drawn from a uniform distribution. We define a rank function $j \rightarrow (k)$ that maps an advertiser j to an ad-slot k, based on the descending order of bids.[3] Therefore, we represent the valuation realized by the k-th highest advertiser as $v^{(k)}$ and her corresponding bid as $b^{(k)}$. The key element of the GSP is that each winning advertiser pays an amount equal to the next highest bid. Therefore, the payoff for the k-th highest-bidding advertiser (allotted to the k-th slot) is given by $\alpha_k(v^{(k)} - b^{(k+1)})$.[4]

In this paper, we focus on two metrics of interest. The first metric of interest is the *allocative efficiency* of the auction. This metric captures the amount of realized social welfare relative to the

[2] Anecdotally, ads in the higher slots tend to receive more clicks, making them more attractive to the advertisers.

[3] Note that search engines have evolved to calculate a quality score based on the previous performance of the ads and the bid amount placed by the advertiser. These quality scores are being used to determine the slot to be assigned to the advertiser. For our work, we consider a simplified case in which the slot allocation is based only on bids.

[4] For ease of notation, we assume that $b^{(4)} = 0$.

maximum possible value. In particular, given the setup above, we define the allocative efficiency as:

$$\Psi = \frac{v^{(1)} + \alpha v^{(2)}}{v^1 + \alpha v^2}. \tag{1}$$

The second metric of interest is the *bid-to-value ratio* for each of the advertisers. This metric captures behavior at the individual level. Specifically, we define the bid-to-value ratio for rank k as:

$$\Omega^{(k)} = \frac{b^{(k)}}{v^{(k)}}. \tag{2}$$

Note that whereas Ω^j is calculated at the individual level, Ψ is calculated at the market level.

We model *frictions* as an additional cost C^j incurred by advertiser j from revising their bids. That is, the payoff for advertiser j is

$$\pi^{(k)} = \alpha_k(v^{(k)} - b^{(k+1)}) - C^{(k)},$$

where C^j could, for example, correspond to implicit costs, such as the efforts taken by the advertisers to repeatedly change their bids, or explicit costs, such as fees charged by the platforms. Notice that in the latter case, the allocative efficiency is equivalent to the overall efficiency of the market.

The equilibrium analysis from Edelman, Ostrovsky, and Schwarz (2007) provides several theoretical predictions for the case in which $C^j = 0$. In particular, regarding the allocative efficiency, the theory predicts that the outcome of the auction will be fully efficient (i.e., $\Psi = 1$). Regarding the bid-to-value ratio, the theory does not make a precise prediction, because infinitely many equilibria are possible. Nevertheless, the assumption that the lowest-valued advertisers will bid their true value (i.e., $\Omega^{(3)} = 1$), which is a weakly dominant strategy for those players, is common. This assumption in turn leads to a prediction regarding the bid-to-value ratio and the slot-similarity parameter for the medium-valued advertisers (i.e., $j = 2$). Specifically, it is straightforward to derive that $\Omega^{(2)}$ is decreasing in α.[5] Crucially, theory makes no predictions regarding the behavior of the highest-valued advertisers (i.e., $j = 1$), and is at odds with experimental evidence on overbidding mentioned above. Furthermore, the existing theory does not incorporate friction costs for the case of $C^j > 0$. Therefore, we turn to a computational model to develop hypotheses for the outcomes of the GSP auction.

3.2 Agent-based Implementation of the GSP Environment

Consistent with the theoretical underpinning, we allow a group of $J = 3$ agents to compete for $K = 2$ ad-slots. The agents learn how to bid using the Q-learning model, described next.

3.2.1 Q-Learning Model Details

The objective of a Q-learning agent is to learn an optimal policy that maximizes the expected reward. Fundamental to the algorithm are the Q-values, denoted as $Q(s, a)$, which represent the value of taking an action a in state s. The Q-values are learned over time using a reinforcement-learning process. Specifically, suppose at time t, the agent selects an action a_t, observes a reward π_t, and enters a state s_{t+1}. Then, the Q-value is updated as follows:

$$Q^{new}(s_t, a_t) \leftarrow (1 - \delta)\; Q(s_t, a_t) + \delta\; (\pi_t + \gamma\; max_a\; Q(s_{t+1}, a)), \tag{3}$$

where $0 \leq \delta \leq 1$ is the learning rate and $0 \leq \gamma \leq 1$ is the discount factor.

Given Q-values, the agent chooses an action (i.e., bid) using a policy function. We use the *softmax* policy function, which is common in the literature. Specifically, the action is determined

[5] In the constructed equilibrium, the second advertiser submits a bid satisfying the following condition: $\alpha(v^2 - b^3) = v^2 - b^2$. By assuming $b^3 = v^3$ (as is done for their equilibrium derivation), we obtain: $\frac{b^2}{v^2} = 1 + \alpha(\frac{v^3}{v^2} - 1)$, where $\frac{v^3}{v^2} < 1$ by construction.

using the Boltzmann probability distribution: $\frac{e^{\lambda Q(s,a)}}{\sum_{a_i} e^{\lambda Q(s,a_i)}}$ $\forall a_i \in A(s)$, where $A(s)$ is the set of actions available in state s, and λ captures the amount of exploration. In this way, Q-learning is a type of stochastic learning model that selects more profitable actions more often.

We implement the Q-learning algorithm in our GSP context as follows. A group of $J = 3$ agents compete for $K = 2$ ad-slots over $M = 2,000$ matches. Each match lasts $T = 100$ time periods, indexed by t. At the beginning of each match ($t = 0$), private values v^j are randomly drawn from $\mathcal{U}\{1, 10\}$. The values are retained for the duration of the match. For every t, the agent chooses to place a bid b_t^j from $\mathcal{U}\{0, 10\}$. As mentioned earlier, the chosen bid depends on the state s that the agent is in. In our implementation of the GSP environment, the state s_t is a tuple of the agent's private valuation and the current bid: (v^j, b_{t-1}^j). That is, the agent j decides on a bid b_t based on the private value, v^j, and own previous bid, b_{t-1}^j. Importantly, we assume that the friction cost of C is incurred every time the agent changes the bid from period $t-1$ to t (i.e, if $b_t \neq b_{t-1}$), and that this cost is a constant and the same across all agents.

Table 1: Summary of Q-learning variables and parameters for GSP

States and actions	
State: $s_t^j \rightarrow$	(v^j, b_{t-1}^j)
Action: $a_t^j \rightarrow$	b_t^j
Reward: $\pi_t^j \rightarrow$	$\pi_t^{(k)} = \alpha_k(v^{(k)} - b_t^{(k+1)}) - C^{(k)}$
Environment parameters	
$\delta \rightarrow$	0.1
$\gamma \rightarrow$	0.99
$\lambda \rightarrow$	1
Treatment variables	
$C \rightarrow$	$\{0.0, 0.5, 1.0\}$
$\alpha \rightarrow$	$\{0.2, 0.5, 0.8\}$

Notes: Recall that a rank function maps agent $j \rightarrow$ slot (k).

At any given t, after all bids are submitted, the slots are allocated in the order of the bids (with ties broken randomly). Each agent is assumed to gain $\alpha_k v^{(k)}$ but incurs $\alpha_k b^k$ as payment to the intermediary and $C^{(k)}$ as the bid-adjustment cost. At the end of the match ($t = 100$), the bids and the outcomes (bid-to-value ratio and the allocative efficiency) are recorded. To derive comparative static predictions, we execute all these steps for various combinations of C and α, specifically, $C \in \{0.0, 0.5, 1.0\}$ and $\alpha \in \{0.2, 0.5, 0.8\}$. Further details on our implementation of the Q-learning model as well as robustness checks for different parameter values are presented in Appendix A.

3.3 Computational Predictions

In this section, we present results of our learning-model simulations. In particular, in Section 3.3.1, we present the market-level outcomes, whereas in Sections 3.3.2 and 3.3.3, we consider individual-level outcomes.

3.3.1 Allocative Efficiency

Panel (a) of Figure 1 presents the evolution of the allocative efficiency, Ω, over the learning horizon, whereas panel (b) of the figure presents a more detailed breakdown of the converged outcomes.

Figure 1: Allocative Efficiency

(a)

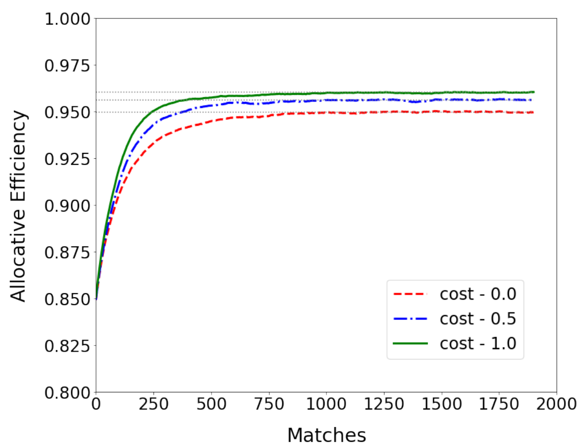

(b)

	$C = 0.0$	$C = 0.5$	$C = 1.0$
$\alpha = 0.2$	0.944	0.952	0.960
$\alpha = 0.5$	0.95	0.956	0.960
$\alpha = 0.8$	0.951	0.957	0.961
Average	0.948	0.955	0.960

Notes: (a) Evolution of allocative efficiency throughout the learning horizon (for $\alpha = 0.5$). (b) Allocative efficiency in converged markets (matches 1,800-2,000).

Figure 1a shows that, initially, allocative efficiency is the same across the three cost treatments (around 0.85). However, as agents learn, the efficiency increases and the differences among the three cost scenarios appear. In particular, the main observation from the figure is that allocative efficiency is higher with higher costs. Figure 1b shows that this finding is true regardless of α, although when alpha is low, the increase is higher. We summarize these observations with Hypothesis 1.

Hypothesis 1: *Allocative efficiency of the market increases as friction costs increase.*

3.3.2 Exploratory Behavior

Table 2 presents the number of bid adjustments that agents make, on average, across $N = 10,000$ simulations. The three panels of Table 2 present the number of bid adjustments for each of the three agents, $j \in \{1, 2, 3\}$, sorted based on their private values and labeled as highest-valued, medium-valued, and lowest-valued agents, respectively.[6] The rows correspond to different values of α, whereas the columns within each panel correspond to different values of C.

Table 2: Number of Bid Adjustments

	(a) Highest-valued agent			(b) Medium-valued agent			(c) Lowest-valued agent		
	$C = 0.0$	$C = 0.5$	$C = 1.0$	$C = 0.0$	$C = 0.5$	$C = 1.0$	$C = 0.0$	$C = 0.5$	$C = 1.0$
$\alpha = 0.2$	48	19	7	49	18	8	67	45	29
$\alpha = 0.5$	51	25	11	45	15	6	57	32	20
$\alpha = 0.8$	53	32	18	45	18	8	52	26	15
Average	51	25	12	46	17	7	59	34	21

Notes: Results are rounded to nearest integer. The maximum number of adjustments can be 100.

Notice that the number of bid adjustments is a simple measure of exploratory behavior by the

[6] For example, if the three participants were assigned private values of 7,4,2, then the participant with the private value of 7 would be considered highest-valued, the participant with the private value of 4 would be considered medium-valued, and the participant with the private value of 2 would be considered lowest-valued.

agents. Thus, the key takeaway from Table 2 is that the number of bid adjustments decreases as the friction costs increase. This finding is intuitive – when the exploration becomes costlier, the agents explore less. We expect to observe a similar result with our human subjects and summarize this prediction with Hypothesis 2.

Hypothesis 2: *Lower costs lead to more exploration.*

3.3.3 Bidding Behavior

Table 3 presents outcomes in terms of the bid-to-value ratios. Similarly to Table 2, the three panels of Table 3 present the average bid-to-value ratios for each of the three agents, $j \in \{1, 2, 3\}$, sorted based on their private values. Again, the rows corresponds to different values of α, whereas the columns within each panel correspond to different values of C.

Table 3: Bid-to-Value Ratios

	(a) Highest-valued agent			(b) Medium-valued agent			(c) Lowest-valued agent		
	$C = 0.0$	$C = 0.5$	$C = 1.0$	$C = 0.0$	$C = 0.5$	$C = 1.0$	$C = 0.0$	$C = 0.5$	$C = 1.0$
$\alpha = 0.2$	0.842	0.852	0.852	0.913	0.918	0.901	1.269	1.232	1.11
$\alpha = 0.5$	0.753	0.756	0.759	0.826	0.829	0.826	1.083	1.056	0.996
$\alpha = 0.8$	0.686	0.688	0.688	0.774	0.777	0.771	0.982	0.969	0.910
Average	0.760	0.765	0.766	0.838	0.842	0.833	1.111	1.086	1.005

Notes: $\Omega^j > 1$ would mean that agent j is overbidding.

There are three takeaways from Table 3. The first takeaway is that the bid-to-value ratios are highest for the lowest-valued agents. Furthermore, for the lowest-valued agents, the bid-to-value ratios are greater than 1.0, on average (i.e., the agents submit bids more than their valuation).

The intuition for this result from our computational model is that the lowest-valued agents are *not* likely to win the auction even if they bid slightly more than their true value, and hence earn zero. Recall that, given that the agents implement the softmax action-selection policy, actions with similar payoffs are equally likely. Thus, bids above the true value would yield an expected payoff that is comparable to bidding the true value. We summarize this prediction with Hypothesis 3.

Hypothesis 3: *Bid-to-value ratio is higher for lower-valued agents.*

The second takeaway from Table 3 is that costs play a role in the bid-to-value ratios *only* for lowest-valued agents. Specifically, the bid-to-value ratios of the highest-valued and medium-valued agents stay remarkably consistent independent of the costs, whereas the bid-to-value ratios for the lowest-valued agents drop by approximately 10% for each of the three values of α. These results suggest that friction moderates the overbidding, particularly for the lowest-valued advertisers. We summarize this prediction with Hypothesis 4.

Hypothesis 4: *For the lowest-valued agents, the bid-to-value ratio increases as friction costs decrease.*

The last takeaway from Table 3 is regarding the the role of α. In particular, the table shows that as α increases, the bid-to-value ratio decreases for all agents, regardless on their respective private value ranks. We summarize this prediction with Hypothesis 5.

Hypothesis 5: *The bid-to-value ratio decreases as α increases.*

Several points are important to reiterate. First, with the exception of Hypothesis 4, the theory of Edelman, Ostrovsky, and Schwarz (2007) does not provide predictions that correspond to those obtained with our agent-based model. Second, the predictions obtained in this section are not intended as point predictions; instead, the takeaways from the tables should be qualitative in nature.

Finally, the predictions are based on a relatively simple learning framework that is independent of the other behavioral factors that may also play a role when humans participate in the auction. Thus, although we expect the general trends observed in the computational model to hold in a human-subject experiment, psychological factors such as auction fever (Adam, Krämer, Jähnig, Seifert, and Weinhardt, 2011), spite (Cooper and Fang, 2008), and joy-of-winning (Kamijo, 2013) among non-winners also can further strengthen or weaken these results.

4 Experimental Design and Administration

This section describes the experimental design of the auction described in Section 3.2.1. The nomenclature used below is consistent with our usage when conducting the experiments and is somewhat different from the description above. We have used different nomenclature in the experiment so that it is intuitive for the subjects. In particular, we refer to the advertisers as *participants* in the experiment, the ad slots as *goods* that the participants bid for, and the auction itself as a *match* that participants bid in. As is the norm with economic experiments, the amount of money that participants make at the end of the experiment depends on their performance in the experiment.

4.1 Treatments

Our primary objective is to investigate the role of friction cost on the outcomes of the GSP. Therefore, the two main treatments of the experiment are with respect to the costs of the bid adjustments imposed on the participants. Specifically, the experiment consisted of two *between-subject* treatments with respect to the cost of adjustment, C. We also set out to vary α – the correlation between the value of the top two slots. However, we varied α within treatment. That is, during the experiment, each participant was likely to experience multiple α's, but the same C. Table 4 presents a summary of the two treatment dimensions.

Table 4: Treatments Summary

Treatments	Parameter Varied	Description
$Between - Subjects$	$C \in \{0.0, 0.1\}$	*Costless* or *Costly* bid adjustment within matches
$Within - Subjects$	$\alpha \in \{0.2, 0.5, 0.8\}$	*Value of the second good* as a fraction of the first good

The between-subject treatments were: *costless* bid adjustment ($C = 0$) and *costly* bid adjustment ($C = 0.1$). In the costless treatment, subjects incurred a cost of $C = 0.0$ for changing their bid during the match. In the costly treatment, subjects incurred a cost of $C = 0.1$ for changing their bid during the auction. In both cases, however, subjects could place the initial bid at no cost. Because subjects incurred the cost every time they changed their bid, they could incur multiple costs in the same match. In particular, a subject could make as many adjustments as she wanted until the time for the match expired. The within-subject treatments were *low* correlation ($\alpha = 0.2$), *medium* correlation ($\alpha = 0.5$), and *high* correlation ($\alpha = 0.8$). All participants in a group had the same $\alpha \in \{0.2, 0.5, 0.8\}$ for the duration of each.

4.2 Matches

Each session consisted of $M = 10$ matches. At the beginning of each match, participants were randomly split into groups of three ($J = 3$) and remained so until the end of the match. The regrouping for the next match was random to avoid any systematic learning about participant behaviors. Earnings for the experiment were the sum of payoffs across all 10 matches.

To avoid the end-of-match effects associated with the fixed duration of a match, we opted for random termination. Specifically, each match lasted at least 20 seconds, after which, there was a 1% chance that the match would terminate each second. Thus, the expected duration of each match was two minutes. To ensure a valid comparison across sessions, the same sequence of seconds across matches was used in every session.[7] We summarize this design choice with Design Remark 1.

Design Remark 1: *Random termination protocol.*

For each match, the participants were provided with randomly drawn private values for good 1. We then obtained the value of good 2 by multiplying the value of good 1 and α. Parameters for the initial four matches were drawn at random *without any restriction*. However, in matches 5 through 10, we aimed to provide a clean comparison among the treatments. Therefore, we used common seeds to generate the same random values across the two treatment dimensions. This approach ensured that any learning that took place in the early matches was not systematically biased and allowed us to compare across different values of α using the later matches. We summarize this design choice with Design Remark 2.

Design Remark 2: *We used common random numbers in matches 5 through 10 such that for*

- *Each match, $m \in \{5, ..., 10\}$, all groups had the same three valuations $\{v_m^1, v_m^2, v_m^3\}$.*
- *Each match, $m \in \{5, ..., 10\}$, there was at least one group for each value of α.*

4.3 Auction Details

At the beginning of a match, each participant j submits a bid b^j at no cost. Participants can then revise their bids in continuous time during the match. Participants can lock the bids to see the associated outcome and payoff with the current combination of bids. Specifically, if her b^j bid is the highest, then she would get the first good and pay the amount equal to the second-highest bid (Recall $\alpha_1 = 1$) minus the friction costs incurred during the match ($b^{(2)} - C^{(1)}$). If her bid is the second highest, then she would get the second good and pay the amount equal to the third-highest

[7] Table C.1 in the Appendix presents the match duration for each of the ten matches.

bid minus friction costs incurred times α, that is, $\alpha(b^{(3)} - C^{(2)})$. If her bid is the lowest, she wouldn't receive any good and would pay nothing. We announced that the α would be the same for all participants in the group during the instructions to ensure common knowledge of this fact.

4.4 Experimental Interface

The experiment was conducted using an interface that was programmed by the authors. The interface implements the continuous-time feature of the auction with participants being able to make bid adjustments in real time. Figure 2 presents the screenshot of the interface. The participant's screen summarized information provided for that match (#1 in Figure 2), as well as current action and the outcome associated with this action (#2 in Figure 2). To place or revise the bid, participants had to use the scale displayed on the left side of the screen (#3 in Figure 2).

We anticipated a potential problem involving mouse clicks. Specifically, mouse clicks could be used as a source of additional information about the behavior of other participants in this experiment. For example, a bid adjustment by subject i, if heard by subject j, could lead to subject j trying to check whether new profitable adjustments were available, by making one or more adjustments herself, which, in turn, could lead to many more clicks. This issue is particularly relevant because we ran multiple groups per session. Furthermore, as described earlier, the number of adjustments was one of the key differences between the treatments that we were looking for; therefore, comparing sessions that contained a large number of clicks (which could be clearly heard in the room) with sessions that didn't could be problematic. To resolve this issue, we implemented a *silent protocol*. Specifically, instead of clicking to select a new bid, subjects placed and adjusted their bid by moving the mouse back and forth across the scale border. This approach resulted in a quiet room throughout the experiment. Thus, subjects could not detect bid changes other than through the information provided on the screen. We summarize this design choice with Design

Remark 3:

Design Remark 3: *Silent protocol.*

Figure 2: Experimental Interface

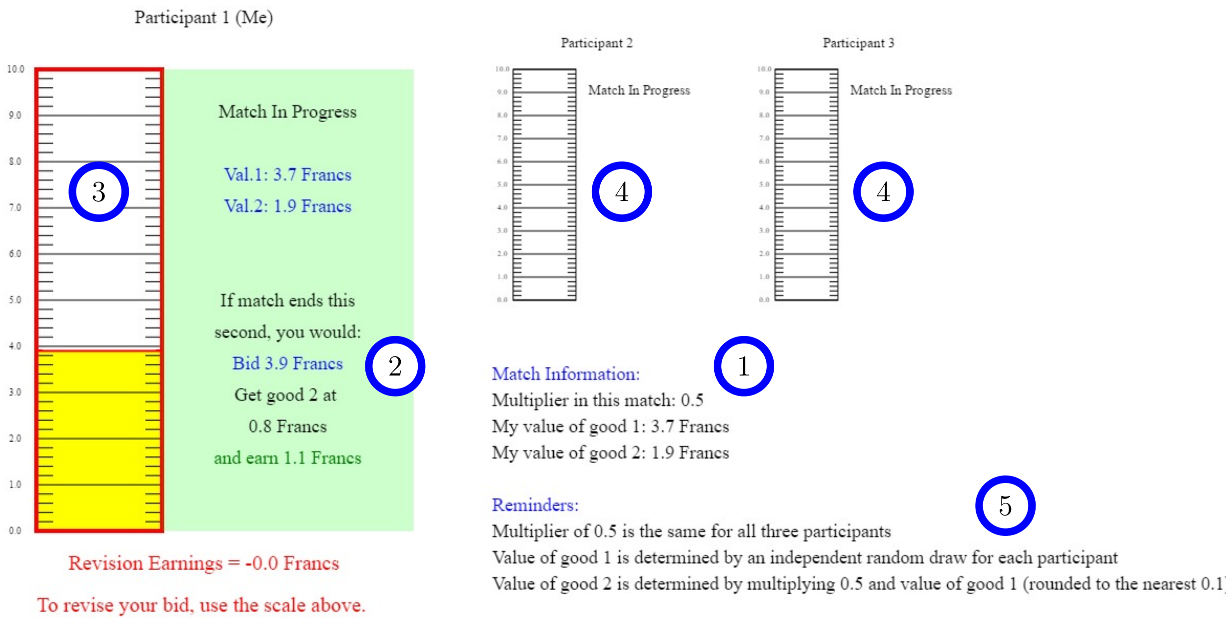

Notes: The screenshot shows: **(1)** Match information. This information is provided prior to the beginning of the match. **(2)** Action and the outcome associated with this action. The outcomes is updated live and depends on the actions of all three participants in the group. **(3)** Scale that is used to place and revise bids. **(4)** Scales for the other two participants. These scales remain blank until the match is over, at which point the actions of the other participants are revealed. **(5)** Reminders about the rules of the experiment.

4.5 Experiment Administration

For the experiment, we recruited 138 participants using ORSEE software (Greiner, 2015) on the campus of Purdue University. We administered eight sessions of the experiment, with the number of participants in each session varying between 15 and 18.[8] Upon entering the laboratory, participants

[8] After running the first four sessions (two for $C = 0.0$ and two for $C = 1.0$), we discovered an error in the way random seeds were generated by the software (recall that the seeds were used in the generation of common values across the groups in matches 5-10). The error was that for matches 5-10, the seed was incremented by 1 rather than 3. So for each group in period $t+1$, there were two values that were the same as in period t, and one new value (instead of three new values). The values were then randomly reassigned within the group. This means that there was an approximately 20% chance for a given subject to have the same value in two consecutive matches and an approximately 10% chance for a subject to have the same value in matches t and $t+2$. The bug was the same across the treatments, so in terms of comparison $C = 0.0$ versus $C = 1.0$ or in terms of comparison across α's, there should be no systematic effect between treatments. In terms of implications for the results, however, this means that subjects

were assigned to a computer terminal. All terminals were separated by physical barriers such that participants could not see choices made by other participants in the room. Participants remained anonymous throughout the experiment.

To ensure that subjects understood the interface and the bid-adjustment process, we took several steps. First, subjects were given a handout containing the instructions (Online Appendix B). An experimenter read them out loud to ensure common knowledge of the environment. Second, subjects had to complete six practice tasks that dealt with placing and modifying the bid. Subjects could proceed to the next task only after correctly completing the previous task. Third, the subjects had to go through five examples, which, to eliminate any bias, were generated at random. In the examples, subjects could practice placing and revising their bids for hypothetical actions by the opponents. Finally, subjects were provided with a calculator, pen, and paper for the duration of the instructions and the experiment. Thus, they were able to verify calculations in the instructions and practices tasks. Furthermore, subjects could make any necessary calculations during the experiment.

The above steps took approximately 20 minutes. Then, prior to each match, subjects were given time to review information for that match. Only after they were ready, did they placed their initial bid. Once everyone had placed the initial bid, the match began. The 10 matches took approximately 30 minutes to complete. After the 10 matches, subjects were paid in cash.

5 Results

The results section is organized as follows: First, in Sections 5.1-5.1.2, we test the hypotheses developed using the computational model. Second, having validated the computational model, we conduct several exercises to provide additional insights into the role that frictions play in the

had more learning opportunities about the same values, which makes our findings about excessive experimentation and over-bidding conservative. We present the data broken down by the first four and last four sessions in Online Appendix C. As expected, we find the results obtained using data from sessions 1-4 to be consistent with results obtained using data from sessions 5-8.

outcome of the GSP auction.

5.1 Allocative Efficiency

Table 5 presents the allocative efficiency in our experiment. Recall that we focus on the outcomes in matches 5-10 to eliminate concerns about initial learning about the environment by the human subjects. In addition, for each match $m \in \{5, ..., 10\}$, private values across groups in session 1-4 and session 5-8 were the same, providing for a clean comparison across treatments.

Table 5: GSP Efficiency

	$C = 0.0$		$C = 0.1$
$\alpha = 0.2$	0.926 (0.021)	$\overset{0.113}{\sim}$	0.967 (0.014)
	$\wr$ 0.894		$\wr$ 0.196
$\alpha = 0.5$	0.93 (0.023)	$\overset{0.016}{\ll}$	0.987 (0.005)
	$\wr$ 0.204		$\wr$ 0.912
$\alpha = 0.8$	0.967 (0.014)	$\overset{0.184}{\sim}$	0.988 (0.007)
	$\wr$ 0.132		$\wr$ 0.259
$\alpha = 0.2$	0.926 (0.021)		0.967 (0.015)
Average:	0.94 (0.012)	$\overset{0.001}{\lll}$	0.98 (0.006)

Notes: Average allocative efficiency for the last bid that subjects placed in each match. Unit of observation is a group of three subjects. Bootstrapped standard errors are in parentheses. $>$, $\gg$, and $\ggg$ denote significance at 0.10, 0.05, and 0.01 levels, respectively. p-values are determined using two-tailed permutation tests (Good, 2013).

Table 5 shows that allocative efficiency is significantly higher when the costs are $C = 0.1$. This result provides support to Hypothesis 1 and leads to the conclusion that frictions in the GSP market can help improve efficiency and therefore the overall social welfare.

5.1.1 Exploratory Behavior

Table 6 presents the number of bid adjustments observed across the treatments of our experiment in matches 5-10. The table is split into three panels based on the rank of the private value of the participant similar to the simulation results in Section 3.3. That is, the participant with the highest private value among the three is labeled as the highest-valued, the participant with the second-highest private value among the three is labeled as medium-valued, and the participant with the lowest private value among the three is labeled as lowest-valued. The columns within each panel vary the costs of bid adjustments. The rows vary the similarity of the two slots that are auctioned off.

Table 6: Subject Bid Adjustments

	(a) Highest-valued			(b) Medium-valued			(c) Lowest-valued		
	$C = 0.0$		$C = 0.1$	$C = 0.0$		$C = 0.1$	$C = 0.0$		$C = 0.1$
$\alpha = 0.2$	31.125 (5.002)	$\ggg^{0.0}$	3.438 (0.76)	51.292 (5.324)	$\ggg^{0.0}$	3.458 (0.454)	65.583 (8.988)	$\ggg^{0.0}$	3.083 (0.525)
	$\sim^{0.657}$		$\sim^{0.45}$	$>^{0.084}$		$\sim^{0.599}$	$\sim^{0.204}$		$\sim^{0.583}$
$\alpha = 0.5$	34.167 (4.238)	$\ggg^{0.0}$	2.667 (0.279)	38.646 (4.77)	$\ggg^{0.0}$	3.125 (0.359)	51.583 (6.065)	$\ggg^{0.0}$	2.708 (0.349)
	$\sim^{0.74}$		$\gg^{0.044}$	$>^{0.053}$		$\sim^{0.287}$	$\sim^{0.938}$		$\sim^{0.346}$
$\alpha = 0.8$	36.857 (6.744)	$\ggg^{0.0}$	1.929 (0.221)	26.69 (3.372)	$\ggg^{0.0}$	2.571 (0.329)	52.333 (6.996)	$\ggg^{0.0}$	4.024 (1.222)
	$\sim^{0.503}$		$\ll^{0.029}$	$\lll^{0.0}$		$\sim^{0.135}$	$\sim^{0.266}$		$\sim^{0.536}$
$\alpha = 0.2$	31.125 (5.074)		3.438 (0.745)	51.292 (5.299)		3.458 (0.447)	65.583 (8.985)		3.083 (0.52)
Average:	33.928 (3.076)	$\ggg^{0.0}$	2.71 (0.29)	39.406 (2.793)	$\ggg^{0.0}$	3.072 (0.224)	56.681 (4.37)	$\ggg^{0.0}$	3.239 (0.435)

Notes: **Panel (a)** presents the average number of bid adjustments made by the highest-valued participants in each group. **Panel (b)** presents the average number of bid adjustments made by the medium-valued participants in each group. **Panel (c)** presents the average number of bid adjustments made by the lowest-valued participants in each group. The unit of observation is a subject per match. Bootstrapped standard errors are in parentheses. $>$, $\gg$, and $\ggg$ denote significance at 0.10, 0.05, and 0.01 levels, respectively. p-values are determined using two-tailed permutation tests (Good, 2013).

We find that the number of bid adjustments in the $C = 0.0$ treatment is an order of magnitude larger than in the $C = 0.1$ treatment. Thus, we find support for Hypothesis 2, namely, that lower costs lead to more exploration in the GSP auction.

5.1.2 Bid-to-Value Ratios

Table 7 presents the bid-to-value ratios observed in the experiment. We find three main results from Table 7. First, comparing panels (a), (b), and (c), we find that the lowest-valued participants substantially overbid compared with the others. This finding validates Hypothesis 3. Second, by comparing columns $C = 0.0$ and $C = 0.1$ from panel (c) we find that the bid-to-value ratio for the lowest-valued participants is higher when frictions are absent. This finding validates Hypothesis 4. Finally, for the bid-to-value ratios across the different values of α, we find partial support for Hypothesis 5. In particular, for two (out of six) cases, the bid-to-value ratios for $\alpha = 0.2$ are significantly higher than for $\alpha = 0.8$, which is consistent with Hypothesis 5; however, for the other four cases, the differences are not significant.

Table 7: Subject Bid-to-Value Ratios

	(a) Highest-valued			(b) Medium-valued			(c) Lowest-valued		
	$C = 0.0$		$C = 0.1$	$C = 0.0$		$C = 0.1$	$C = 0.0$		$C = 0.1$
$\alpha = 0.2$	0.952 (0.032)	$\overset{0.209}{\sim}$	0.901 (0.025)	0.951 (0.037)	$\overset{0.453}{\sim}$	0.913 (0.031)	2.826 (0.377)	$\overset{0.003}{\ggg}$	1.668 (0.191)
	$\gg$ 0.018		$\gg$ 0.011	$\sim$ 0.605		$\sim$ 0.285	$\sim$ 0.372		$\sim$ 0.194
$\alpha = 0.5$	0.839 (0.033)	$\overset{0.316}{\sim}$	0.793 (0.032)	0.92 (0.044)	$\overset{0.329}{\sim}$	0.863 (0.034)	3.815 (0.889)	$\overset{0.0}{\ggg}$	1.341 (0.155)
	$\gg$ 0.041		$\sim$ 0.209	$\ggg$ 0.003		$\sim$ 0.192	$\sim$ 0.307		$\sim$ 0.802
$\alpha = 0.8$	0.725 (0.044)	$\overset{0.942}{\sim}$	0.729 (0.039)	0.734 (0.045)	$\overset{0.39}{\sim}$	0.789 (0.044)	2.646 (0.61)	$\overset{0.006}{\ggg}$	1.284 (0.153)
	$\lll$ 0.0		$\lll$ 0.0	$\lll$ 0.0		$<$ 0.023	$\sim$ 0.816		$\sim$ 0.13
$\alpha = 0.2$	0.952 (0.032)		0.901 (0.024)	0.951 (0.037)		0.913 (0.031)	2.826 (0.373)		1.668 (0.185)
Average:	0.844 (0.023)	$\overset{0.283}{\sim}$	0.811 (0.019)	0.874 (0.026)	$\overset{0.635}{\sim}$	0.858 (0.021)	3.115 (0.387)	$\overset{0.0}{\ggg}$	1.438 (0.097)

Notes: **Panel (a)** presents the average bid-to-value ratio per match across subjects with the highest private value in each group. **Panel (b)** presents the average bid-to-value ratio per match across subjects with the second-highest private value in each group. **Panel (c)** presents the average bid-to-value ratio per match across subjects with the lowest private value in each group. The unit of observation is a subject per match. Bootstrapped standard errors are in parentheses. $>$, $\gg$, and $\ggg$ denote significance at 0.10, 0.05, and 0.01 levels, respectively. p-values are determined using two-tailed permutation tests (Good, 2013).

To summarize the computational and experimental results – we find that behavior by the lowest-valued participant (as captured by the number of bid adjustments and bid-to-value ratios) is key to the outcomes of the GSP. As presented in Table 6, in the presence of frictions, the agents seldom explore and experiment with higher bidding strategies. However, in the absence of friction costs, the exploration increases substantially. The exploratory behavior, in turn, is associated with substantial overbidding for the lowest-valued participants (as presented in Table 7). Such overbidding further cascades into the bids placed by the higher-valued agents, by either pushing them to increase their bids or stay put and be less likely to win the auction, which in turn may lead to an inefficient allocation (as presented in Table 5).

5.2 Additional Insights

Given that our computational model was successful at qualitatively predicting the outcomes from the human-subject experiments, we now extend it to scenarios that are not practical (e.g., costly) to carry out in the laboratory. In other words, we treat the computational agents as being the digital twins to the experimental agents to study several scenarios. In particular, in Section 5.2.1, we consider the effect of increased market demand on bid-to-value ratios and allocative efficiency. Then, in Section 5.2.2, we consider the effect of increased market supply on the two measures of interest.

5.2.1 Impact of Increasing Market Demand

What happens if the number of advertisers in the market increases? In particular, what if the entry cost is lowered and a new set of low-valued advertisers enters the market? According to Edelman, Ostrovsky, and Schwarz (2007), such an increase in low-valued advertisers should not affect the bidding behavior or efficiency of the auctions, because the non-winning participants submit their values and would not win. However, as we have shown empirically, the lowest-valued advertisers are most likely to overbid. Hence, an increase in competition among the lowest-valued advertisers might further impact our auction metrics.

Figure 3 presents the efficiency and the bid-to-value ratios when we vary the number of advertisers from three to six. Specifically, we hold $\alpha = 0.5$, $C = 0.5$, $K = 2$, and vary $J \in \{3, 4, 5, 6\}$. Crucially, in the simulations, we have retained the same valuation for the top two advertisers. That is, additional agents had valuations at most as the second-highest advertiser. In the figure, the bid-to-value ratio of agents is grouped by the rank of the advertiser. The colors across the groups corresponds to a specific scenario; for example, the red bar corresponds to having three agents in the marketplace, and the blue bars correspond to having six agents in the marketplace.

Figure 3: Increasing Number of Advertisers

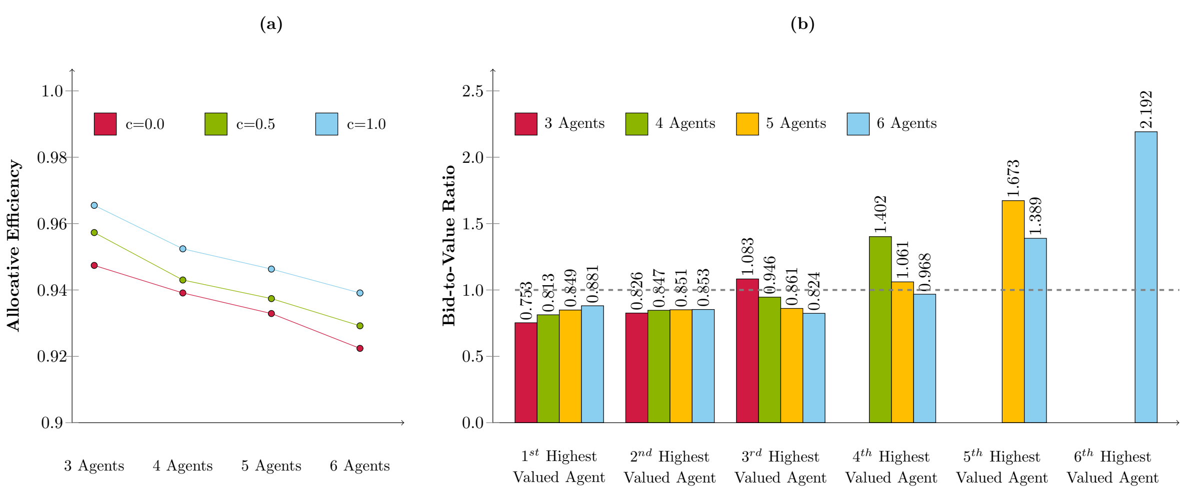

Note: (a) Allocative efficiency. (b) Bid-to-value ratios. The simulation results are for $\alpha = 0.5$,and $K = 2$ ad-slots. Private values of the first- and second highest-valued agents are held the same. Private values for the remaining agents are restricted to be at most the value of the second highest-valued agent. For the bid-to-value ratio simulation, the bid-adjustment cost is set to $C = 0.5$

Several observations from Figure 3 are worth noting. First, panel (a) shows that as the number of advertisers increases, the efficiency of the auction decreases. Second, panel (b) shows that the lowest-valued agents overbid. Furthermore, as the number of agents increases, the amount of overbidding by the lower-valued players also increases. Finally, panel (b) shows that there are cascading price increases across the entire cohort; that is, the lowest-valued advertiser increases the bid by the second-lowest valued advertiser and so on. So, if we focus on the "1st highest-valued agent," we find that the bid-to-value ratio increases from 0.753 to 0.881, when the number of advertisers increases. Importantly, these changes are not consistent with the theory in Edelman, Ostrovsky, and Schwarz (2007). Not surprisingly, the main result of our paper —allocative efficiency increasing with an increase in friction costs— holds across different market sizes.

5.2.2 Impact of Increasing Market Supply

What happens if the number of ad-slots available increases? To answer this question, we fix the number of agents at $N = 5$ and vary the number of slots $K \in \{2, 3, 4\}$. To proceed with this exercise, we need to make an additional assumption regarding the similarity of new ad-slots. Specifically, we assume that the value of α changes exponentially with the number of slots. For example, when 3 slots are available, we have $\alpha_1 = 1$, $\alpha_2 = \alpha$, and $\alpha_3 = \alpha^2$. Figure 4 presents the results. Each color in the figure corresponds to the number of auctioned ad slots. Figure 4 shows that as the number of slots increases, efficiency goes up, and the bid-to-value ratio of the lowest-valued advertizers goes down.

Figure 4: Increasing Number of Ad Slots

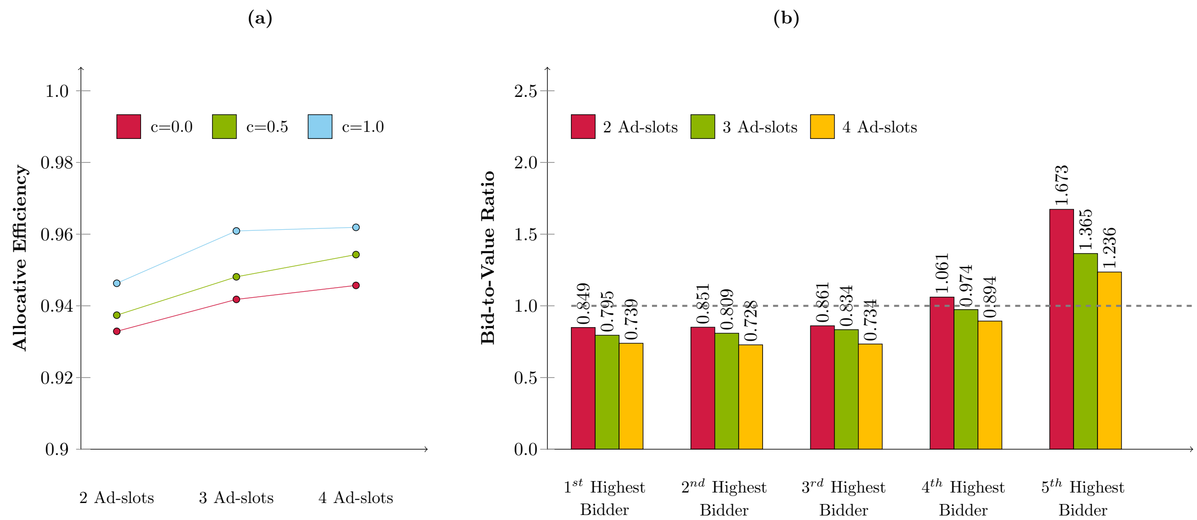

Note: (a) Allocative efficiency, (b) Bid-to-value ratios. The simulation results correspond to $\alpha = 0.5$ and $N = 5$ agents. The simulation results correspond to $\alpha = 0.5$. For the bid-to-value ratio simulation, the bid-adjustment cost is set to $C = 0.5$

To summarize, the additional analyses provide valuable insights regarding the allocative efficiency of the GSP auction across a number of market scenarios. In particular, we find that an increase in market demand (i.e., increase in the number of advertisers) exacerbates the problem of excessive exploration by the lowest-valued players, leading the market to less efficient outcomes. Therefore, increasing the supply of ad slots by the platform possibly resolves this issue, even if the extra slots are not of high value. However, across all the scenarios, we find that friction costs play a consistent role in determining allocative efficiency of the GSP auction. Specifically, allocative efficiency with costs is generally higher than without costs.

6 Conclusion

In this research, we investigate the role frictions for GSP outcomes. First, we computationally replicate the GSP environment and obtain predictions pertaining to bidding behavior and auction efficiency. We further test these predictions using an economic experiment with human subjects.

We find significant overbidding by the lowest-valued advertisers, and that the presence of friction costs moderates this overbidding. In particular, subjects with the lowest private valuations are the most likely to explore their bidding strategies and learn about the behavior of other agents. This exploration leads the lowest-valued player to overbid, contradicting the assumption made in the theory of GSP. We find that the absence of friction costs leads to excessive experimentation, which hinders the market's ability to discover the optimal allocation.

From a slightly broader perspective, we make three key contributions. First, we demonstrate systematically using both the computational and the experimental model, that the problem of reducing frictions through algorithms does not necessarily translate into improving social welfare. As machine learning and AI become more prevalent, we should be cognizant of the impact of reducing frictions. The second key contribution that we wish to highlight is the use of machine-learning agents to create digital twins – a concept quite prevalent in the industry – to develop some actionable insights. We are unaware of any prior work in the IS area that has demonstrated the similarity in behaviors between computational and experimental agents. Third, we wish to highlight that the use of computational agents to develop hypotheses as another distinctive feature of our paper.

Our research is not without limitations. In particular, we considered a simplified version of the GSP in which the rank is determined solely by the advertiser's bid. In recent years, however, sponsored-search-advertising platforms have started to include other factors (e.g., ad quality, advertiser's history, etc.) to determine rank of the bid. Future research could incorporate these factors and ranking methods to understand the properties of the new auction mechanisms. The second limitation is that in this paper we have assumed all advertisers face the same cost for bid adjustments. In practice, however, there is vast heterogeneity among advertisers in terms of costs they incur, both for participating in the market and for making bid adjustments. Finally, in our research we focused on the scenario in which agents bid directly. Future research could extend our work to scenarios in which agents choose among a set of AI tools that would make bids for them. In particular, it will be important to develop mechanisms that are robust to the presence of both types of bidders in the market.

References

Adam, M. T., J. Krämer, C. Jähnig, S. Seifert, and C. Weinhardt (2011): "Understanding auction fever: A framework for emotional bidding," *Electronic Markets*, 21(3), 197.

Adomavicius, G., J. Bockstedt, C. Shawn, and J. Zhang (2014): "De-biasing user preference ratings in recommender systems," in *Joint Workshop on Interfaces and Human Decision Making for Recommender Systems, IntRS 2014, Co-located with ACM Conference on Recommender Systems, RecSys 2014*. CEUR-WS.

Adomavicius, G., J. C. Bockstedt, S. P. Curley, and J. Zhang (2013): "Do Recommender Systems Manipulate Consumer Preferences? A Study of Anchoring Effects," *Information Systems Research*, 24(4), 956–975.

Adomavicius, G., S. P. Curley, A. Gupta, and P. Sanyal (2012): "Effect of information feedback on bidder behavior in continuous combinatorial auctions," *Management Science*, 58(4), 811–830.

——— (2013): "Impact of information feedback in continuous combinatorial auctions: An experimental study of economic performance," *MIS Quarterly: Management Information Systems*, 37(1), 55–76.

Adomavicius, G., A. Gupta, and P. Sanyal (2006): "Computational feedback mechanisms for iterative multiunit multiattribute auctions," in *16th Workshop on Information Technologies and Systems, WITS 2006*, pp. 205–210. Social Science Research Network.

Agarwal, A., K. Hosanagar, and M. D. Smith (2015): "Do Organic Results Help or Hurt Sponsored Search Performance?," *Information Systems Research*, 26(4), 695–713.

Allen, T. (2014): "Information frictions in trade," *Econometrica*, 82(6), 2041–2083.

Amaldoss, W., P. S. Desai, and W. Shin (2015): "Keyword search advertising and first-page bid estimates: A strategic analysis," *Management Science*, 61(3), 507–519.

Anenberg, E. (2016): "Information frictions and housing market dynamics," *International Economic Review*, 57(4), 1449–1479.

Animesh, A., V. Ramachandran, and S. Viswanathan (2010): "Research noteâĂŤQuality uncertainty and the performance of online sponsored search markets: An empirical investigation," *Information Systems Research*, 21(1), 190–201.

Animesh, A., S. Viswanathan, and R. Agarwal (2011): "Competing âĂIJcreativelyâĂİ in sponsored search markets: The effect of rank, differentiation strategy, and competition on performance," *Information Systems Research*, 22(1), 153–169.

Arnon, A., and Y. Mansour (2011): “Repeated budgeted second price ad auction,” in *International Symposium on Algorithmic Game Theory*, pp. 7–18. Springer.

Bagozzi, R. P., F. D. Davis, and P. R. Warshaw (1992): “Development and test of a theory of technological learning and usage,” *Human relations*, 45(7), 659–686.

Bapna, R., A. Barua, D. Mani, and A. Mehra (2010): “Research commentaryâĂŤCooperation, coordination, and governance in multisourcing: An agenda for analytical and empirical research,” *Information Systems Research*, 21(4), 785–795.

Bassi, V., and A. Nansamba (2017): “Information Frictions in the Labor Market: Evidence from a Field Experiment in Uganda,” *University College London. Processed.*

Bichler, M., P. Shabalin, and G. Ziegler (2013): “Efficiency with linear prices? A game-theoretical and computational analysis of the combinatorial clock auction,” *Information Systems Research*, 24(2), 394–417.

Bowling, M., and M. Veloso (2001): “Rational and convergent learning in stochastic games,” in *International joint conference on artificial intelligence*, vol. 17, pp. 1021–1026. LAWRENCE ERLBAUM ASSOCIATES LTD.

Brandimarte, L., A. Acquisti, and G. Loewenstein (2013): “Misplaced confidences: Privacy and the control paradox,” *Social Psychological and Personality Science*, 4(3), 340–347.

Capasso, S. (2008): “Endogenous information frictions, stock market development and economic growth,” *The Manchester School*, 76(2), 204–222.

Cason, T. N., K. N. Kannan, and R. Siebert (2011): “An experimental study of information revelation policies in sequential auctions,” *Management Science*, 57(4), 667–688.

Che, Y.-K., S. Choi, and J. Kim (2017): “An experimental study of sponsored-search auctions,” *Games and Economic Behavior*, 102, 20–43.

Chen, J., J. Feng, and A. B. Whinston (2010): “Keyword Auctions, UnitâĂŘPrice Contracts, and the Role of Commitment,” *Production and Operations Management*, 19(3), 305–321.

Chen, W., T.-Y. Liu, and X. Yang (2016): “Reinforcement learning behaviors in sponsored search,” *Applied Stochastic Models in Business and Industry*, 32(3), 358–367.

Cooper, D. J., and H. Fang (2008): “Understanding overbidding in second price auctions: An experimental study,” *The Economic Journal*, 118(532), 1572–1595.

Deng, X., and J. Yu (2009): “A new ranking scheme of the GSP mechanism with Markovian users,” in *International Workshop on Internet and Network Economics*, pp. 583–590. Springer.

Devanur, N. R., and S. M. Kakade (2009): “The price of truthfulness for pay-per-click auctions,” in *Proceedings of the 10th ACM conference on Electronic commerce*, pp. 99–106. ACM.

Du, X., M. Su, X. M. Zhang, and X. Zheng (2017): “Bidding for Multiple Keywords in Sponsored Search Advertising: Keyword Categories and Match Types,” *Information Systems Research*, 28(4), 711–722.

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.

Edelman, B., and M. Schwarz (2010): "Optimal Auction Design and Equilibrium Selection in Sponsored Search Auctions," *American Economic Review*, 100(2), 597–602.

Feng, J., H. K. Bhargava, and D. M. Pennock (2007): "Implementing sponsored search in web search engines: Computational evaluation of alternative mechanisms," *INFORMS Journal on Computing*, 19(1), 137–148.

Fukuda, E., Y. Kamijo, A. Takeuchi, M. Masui, and Y. Funaki (2013): "Theoretical and experimental investigations of the performance of keyword auction mechanisms," *The RAND Journal of Economics*, 44(3), 438–461.

Gatti, N., A. Lazaric, and F. Trovò (2012): "A truthful learning mechanism for contextual multi-slot sponsored search auctions with externalities," in *Proceedings of the 13th ACM Conference on Electronic Commerce*, pp. 605–622. ACM.

Ghose, A., and S. Yang (2009): "An Empirical Analysis of Search Engine Advertising: Sponsored Search in Electronic Markets," *Management Science*, 55(10), 1605–1622.

Goh, K. H., and J. C. Bockstedt (2013): "The Framing Effects of Multipart Pricing on Consumer Purchasing Behavior of Customized Information Good Bundles," *Information Systems Research*, 24(2), 334–351.

Good, P. (2013): *Permutation tests: a practical guide to resampling methods for testing hypotheses*. Springer Science & Business Media.

Goodhue, D. L., and R. L. Thompson (1995): "Task-technology fit and individual performance," *MIS quarterly*, pp. 213–236.

Greenwald, A., K. Hall, and R. Serrano (2003): "Correlated Q-learning," in *ICML*, vol. 3, pp. 242–249.

Greenwald, A., K. Kannan, and R. Krishnan (2010): "On Evaluating Information Revelation Policies in Procurement Auctions: A Markov Decision Process Approach," *Information Systems Research*, 21(1), 15–36.

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

Guerci, E., A. Kirman, and S. Moulet (2014): "Learning to bid in sequential Dutch auctions," *Journal of Economic Dynamics and Control*, 48, 374–393.

Guo, Z., G. J. Koehler, and A. B. Whinston (2012): "A Computational Analysis of Bundle Trading Markets Design for Distributed Resource Allocation," *Information Systems Research*, 23(3-part-1), 823–843.

Gupta, A., K. Kannan, and P. Sanyal (2018): "Economic Experiments in Information Systems," *Management Information Systems Quarterly*, 42(2), 595–606.

Hou, K., and T. J. Moskowitz (2005): "Market frictions, price delay, and the cross-section of expected returns," *The Review of Financial Studies*, 18(3), 981–1020.

Jerath, K., L. Ma, Y.-H. Park, and K. Srinivasan (2011): "A âĂIJposition paradoxâĂİ in sponsored search auctions," *Marketing Science*, 30(4), 612–627.

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: Journal of the Econometric Society*, pp. 1275–1304.

Kagel, J. H., and D. Levin (2001): "Behavior in multi-unit demand auctions: experiments with uniform price and dynamic Vickrey auctions," *Econometrica*, 69(2), 413–454.

Kagel, J. H., D. Levin, and R. M. Harstad (1995): "Comparative static effects of number of bidders and public information on behavior in second-price common value auctions," *International Journal of Game Theory*, 24(3), 293–319.

Kamijo, Y. (2013): "Bidding behaviors for a keyword auction in a sealed-bid environment," *Decision Support Systems*, 56, 371–378.

Kempe, D., and M. Mahdian (2008): "A cascade model for externalities in sponsored search," in *International Workshop on Internet and Network Economics*, pp. 585–596. Springer.

Ketter, W., J. Collins, M. Gini, A. Gupta, and P. Schrater (2012): "Real-Time Tactical and Strategic Sales Management for Intelligent Agents Guided by Economic Regimes," *Information Systems Research*, 23(4), 1263–1283.

Kiose, D., and V. Voudouris (2015): "The ACEWEM framework: An integrated agent-based and statistical modelling laboratory for repeated power auctions," *Expert Systems with Applications*, 42(5), 2731–2748.

Littman, M. L. (1994): "Markov games as a framework for multi-agent reinforcement learning," in *Proceedings of the eleventh international conference on machine learning*, vol. 157, pp. 157–163.

McLaughlin, K., and D. Friedman (2016): "Online Ad Auctions: An Experiment," Working Papers 16-05, Chapman University, Economic Science Institute.

Noti, G., N. Nisan, and I. Yaniv (2014): "An experimental evaluation of bidders' behavior in ad auctions," in *Proceedings of the 23rd international conference on World wide web*, pp. 619–630. ACM.

Qin, T., W. Chen, and T.-Y. Liu (2015): "Sponsored search auctions: Recent advances and future directions," *ACM Transactions on Intelligent Systems and Technology (TIST)*, 5(4), 60.

Rosokha, Y., and K. Younge (forthcoming): "Motivating Innovation: The Effect of Loss Aversion on the Willingness to Persist," *Review of Economics and Statistics*, pp. 1–45.

Sandholm, T. W., and R. H. Crites (1996): "Multiagent reinforcement learning in the iterated prisoner's dilemma," *Biosystems*, 37(1-2), 147–166.

Sanyal, P. (2016): "Characteristics and Economic Consequences of Jump Bids in Combinatorial Auctions," *Information Systems Research*, 27(2), 347–364.

Sheremeta, R. M. (2010): "Experimental comparison of multi-stage and one-stage contests," *Games and Economic Behavior*, 68(2), 731–747.

Simonov, A., C. Nosko, and J. M. Rao (2018): "Competition and crowd-out for brand keywords in sponsored search," *Marketing Science*, 37(2), 200–215.

Statista (2017): "Paid Search Advertising Expenditure Worldwide," https://www.statista.com/statistics/267056/paid-search-advertising-expenditure-worldwide/.

Su, X. (2008): "Bounded rationality in newsvendor models," *Manufacturing & Service Operations Management*, 10(4), 566–589.

Sutton, R. S., and A. G. Barto (1998): *Reinforcement learning: An introduction*, vol. 1. MIT press Cambridge.

Tsai, J. Y., S. Egelman, L. Cranor, and A. Acquisti (2011): "The effect of online privacy information on purchasing behavior: An experimental study," *Information Systems Research*, 22(2), 254–268.

Varian, H. R. (2007): "Position auctions," *International Journal of Industrial Organization*, 25(6), 1163 – 1178.

——— (2009): "Online Ad Auctions," *American Economic Review*, 99(2), 430–34.

Watkins, C. J., and P. Dayan (1992): "Q-learning," *Machine learning*, 8(3-4), 279–292.

Xu, L., J. Chen, and A. Whinston (2012): "Effects of the Presence of Organic Listing in Search Advertising," *Information Systems Research*, 23(4), 1284–1302.

Zhang, X. M., and J. Feng (2011): "Cyclical Bid Adjustments in Search-Engine Advertising," *Management Science*, 57(9), 1703–1719.

Zhou, Y., D. Chakrabarty, and R. Lukose (2008): "Budget constrained bidding in keyword auctions and online knapsack problems," in *International Workshop on Internet and Network Economics*, pp. 566–576. Springer.

Online Appendix

A Additional Details about Q-learning for GSP

In this section, we elaborate the details of Q-learning implementation for GSP. In particular, Q-learning (Watkins and Dayan, 1992) is a Markov Decision Process (MDP), where agents learn an optimal (action selection) policy that maximizes the expected value of their overall reward. In other words, it is an advanced stochastic learning process. The learning happens through repeated interactions with the environment over time. For example, suppose that at time t agent is in state s_t and selects an action $a \in A(s_t)$. Furthermore, suppose that an agent observes a reward π_t and transitions to the state s_{t+1}, then the agent updates their value ($Q(s_t, a)$) of being in state s_t and taking an action a through a simple convex combination of the old value and the new value (i.e., value of the reward plus the value of being in the new state), as follows:

$$Q^{new}(s_t, a) \leftarrow (1 - \delta)\, Q(s_t, a) + \delta\, (\pi_t + \gamma\, max_{a' \in A(s_{t+1})}\, Q(s_{t+1}, a')), \tag{4}$$

where $0 \leq \delta \leq 1$ is the learning rate and $0 \leq \gamma \leq 1$ is the discount factor.

The actions are chosen stochastically in such a way that better actions (i.e., those that have higher Q-values) are chosen more often. A common implementation of the action-section policy is through the use of the Boltzmann function (also known as the *softmax* policy). Algorithm 1 presents the pseudo code for implementing Q-learning in the context of GSP.

Initialize $Q(s, a)$ for all state-action pairs (s,a) for $J = 3$ agents. The state, s, is a tuple of private value and previous bid (v^j, b^j_{t-1}), and the action, a, is a possible bid $b \in \{0.0, 1.0, ..., 9.0, 10.0\}$

for $N \leq 1000$ *iterations* **do**
- **for** $M \leq 2000$ *matches* **do**
  - Assign private values $V^j \in$ U$\{1, 10\}$ for all J agents
  - **for** $t \leq 100$ *time periods* **do**
    - For each agent j
      1. Draw a bid $a^j_t \in U\{0, 10\}$ using *Boltzmann softmax* policy over $Q(s^j_t)$
      2. Observe reward, $\pi^{(k)}_t = \alpha_k (v^{(k)}_t - b^{(k+1)}_t) - C^{(k)}$, and next state, s^j_{t+1}
      3. Update $Q^{new}(s_t, a_t) \leftarrow (1 - \delta)\, Q(s_t, a_t) + \delta\, (\pi_t + \gamma\, max_a\, Q(s_{t+1}, a))$
  - **end**
  - Record the auction outcomes
- **end**
- Repeat with the same values of for V

end

Algorithm 1: Pseudo code for the Q-learning algorithm

The learning simulations are run for $M = 2000$ matches, each with $t = 100$ periods. Once the agents submit their bids at period t, the auctioneer ranks them and reports the allotted slot as well as the clearing price for each agent. Since there are $K = 2$ slots, only two agents receive allocations. The agents update their Q-value function at the end of each period t and auction observations are recorded at $t = 100$ (i.e., at the end of each match). For each match, we record private value V^j, final period bid b^j_{100} and number of bid updates made by each player during the match. Using these observations, we calculate allocative efficiency (Ψ) and bid-to-value ratios (Ω^j). We execute the entire learning model for $N = 1000$ iterations. To make these iterations comparable, we retain v^j to be the same for a given match m across the iterations.

The main research questions of the paper are regarding the effect of friction costs (C) and ad slot similarity (α), therefore we varied $C \in \{0.0, 0.5, 1.0\}$ and $\alpha \in \{0.2, 0.5, 0.8\}$ in the main body of the paper. Nevertheless, there are three additional parameters (δ, γ and λ) that are part of the learning model. In particular, for the results carried out in the main body of the paper, we set $\delta = 0.1$, $\gamma = 0.99$ and $\lambda = 1$. Next, we describe these parameters in more detail and provide robustness checks on the extent to which our conclusions depend on this choice.

A.1 Learning rate, δ

The learning rate, δ, determines the extent to which the reward and value of being in the new state override the previously learned value. It is common practice is to use $\delta = 0.1$. To check whether predictions obtained in Section 3.3 of the paper are robust to the choice of δ, we vary $\delta \in \{0.05, 0.1, 0.2\}$. Figure A.1 presents the results for allocative efficiency (panel (a)) and bid-to-value ratios (panel(b)). The figure shows that the key takeaways on allocative efficiency and bid-to-value ratios are qualitatively the same, regardless of δ.

Figure A.1: Robustness of δ

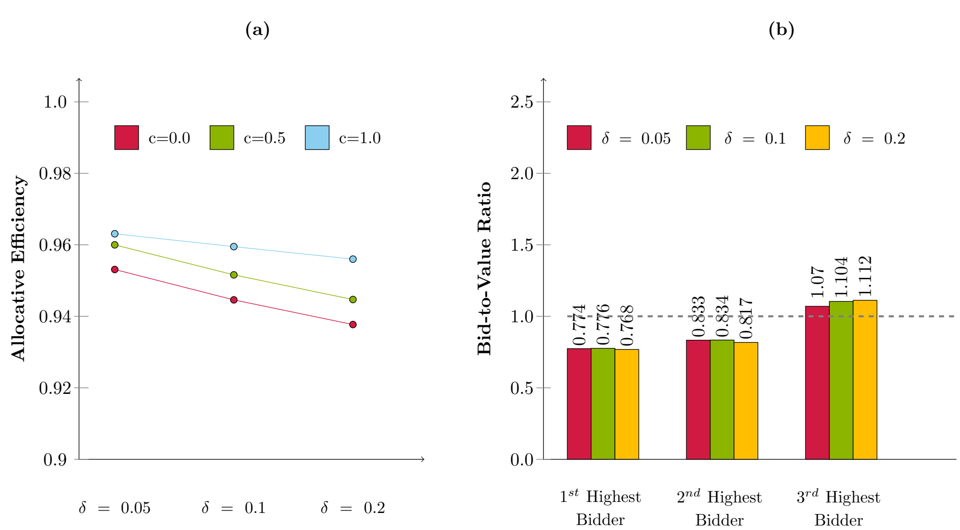

Note: (a) Allocative efficiency. (b) Bid-to-value ratios. The simulation results are for $\alpha = 0.5$, $\lambda = 1$, $\gamma = 0.99$ and $K = 2$ ad slots. For the bid-to-value ratio simulation, the bid-adjustment cost is set to $C = 0.5$.

A.2 Discounting factor, γ

The discounting factor, γ, determines the extent to which the agent values the immediate reward relative to the future rewards. For example, when $\gamma \to 0$, the agent is myopic, because she only considers current reward. For the simulations provided in the main body of the paper, we used $\gamma = 0.99$. This was done in order to keep the value γ the same as for human-subject experiments.[9] To further clarify that our results are *not* driven by the choice of γ, we present the robustness results for $\gamma \in \{0.1, 0.5, 0.99\}$ in Figure A.2. Notice that the impact of costs on allocative efficiency and bidding behaviour remain qualitatively consistent with our main result. The figure also shows that the problem identified in the main body of the paper is amplified as agents become more myopic.

Figure A.2: Robustness of γ

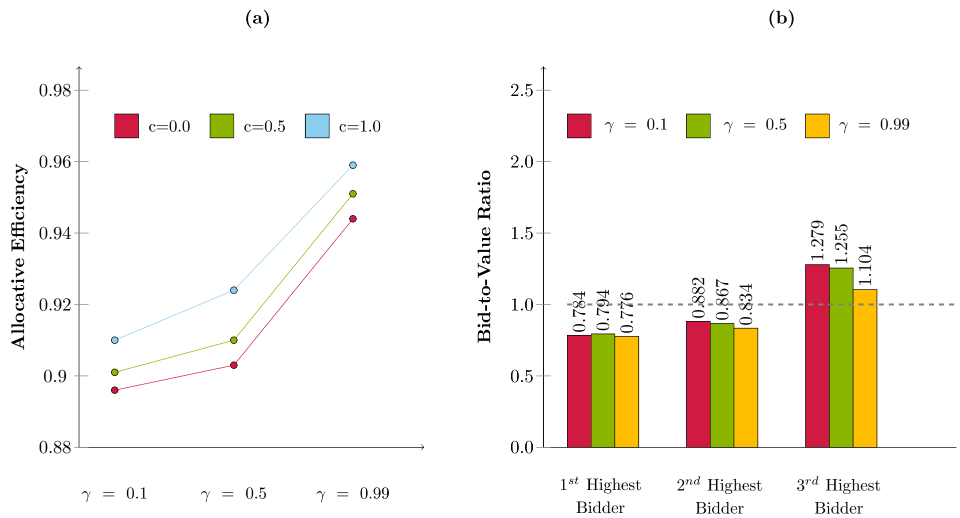

Note: (a) Allocative efficiency. (b) Bid-to-value ratios. The simulation results are for $\alpha = 0.5$, $\lambda = 1$, $\delta = 0.1$ and $K = 2$ ad-slots. For the bid-to-value ratio simulation, the bid-adjustment cost is set to $C = 0.5$.

A.3 Temperature/exploration parameter, λ

The temperature parameters, λ, determines the extent to which better actions are chosen over poor actions. For example, as $\lambda \to 0$, all the actions will have equal probability of being chosen, while as $\lambda \to \infty$ the probability of the choosing the best action tends to 1. For this reason, the temperature parameters, λ, is often interpreted as the rationality parameters (e.g., Su, 2008). In the main body of the paper, we set $\lambda = 1$. Robustness results for varying λ are presented in Figure A.3. In particular, we find the results are qualitatively consistent with our main result. However, we do find that overbidding by the third highest-valued agents is not as prevalent when λ is lower.

[9]In the experiment there was a 99% chance that the match will continue to the next period and 1% probability that the match will terminate in the current period.

Figure A.3: Robustness of λ

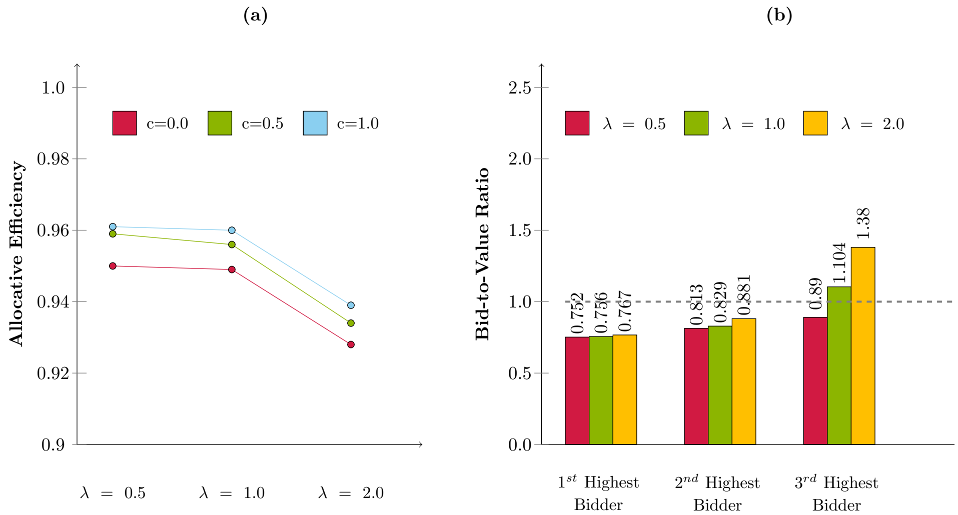

Note: (a) Allocative efficiency. (b) Bid-to-value ratios. The simulation results are for $\alpha = 0.5$, $\gamma = 0.99$, $\delta = 0.1$ and $K = 2$ ad-slots. For the bid-to-value ratio simulation, the bid-adjustment cost is set to $C = 0.5$.

B Experimental Instructions

Experiment Overview

You are about to participate in an experiment in the economics of decision-making. If you listen carefully, you could earn a large amount of money that will be paid to you in cash in private at the end of the experiment.

It is important that you remain silent and do not look at other people's work. If you have any questions, or need any assistance of any kind, please raise your hand and an experimenter will come to you. During the experiment, **do not talk, laugh or exclaim out loud**, and be sure to keep your **eyes on your screen only**. In addition, please **turn off your cell phones, etc.** and put them away. Anybody that violates these rules will be asked to leave and will **not** be paid. We expect and appreciate your cooperation.

Agenda

1. We will first go over the printed instructions.
2. After the printed instructions, there will be a set of interactive instructions on the computer that will guide you through elements of the interface.
3. After all the instructions, the experiment will begin. In the experiment, you will be working with a fictitious currency called Francs. You will be paid in US Dollars at the end of the experiment. **The exchange rate today is: 1 Franc = 0.75 USD.**

Experiment Details

- This experiment consists of **ten** matches.
- At the beginning of each match you will be **randomly matched with two other participants**. You will remain matched with these same participants until the **end of the match**, but then you will be re-matched with another two randomly selected participants for the following match.
- Each match will have the same structure, but may take different amount of time.
- You will remain **anonymous** throughout the experiment. You will not know the identity of the participant that you are matched with, and they will not know your identity.
- Your earnings in a given match is based solely on the choices made by you and the participants with whom you are matched. The choices made by you and by the participants with whom you are matched will have no effect on the earnings of participants in other groups and vice versa.

Specific Instructions for Each Match

- At the beginning of each match you will be assigned a value for good 1 and a value for good 2 as follows:
    - Your value for good 1 will be drawn at random from $\{0.1, 0.2, 0.3, \ldots, 9.8, 9.9\}$ with each number equally likely.

  - Your value for good 2 will be determined by multiplying your value of good 1 and a "multiplier" which is a fraction between 0 and 1. This **multiplier will be common for all participants in your group** and fixed for the duration of each match.
  - Thus, the value of good 1 is greater than the value of good 2.

- For example, if your value for good 1 is randomly drawn to be 7.0 Francs and the multiplier is 0.5, your value for good 2 is 3.5 (=7.0*0.5). Note that all numbers are rounded to the nearest 0.1.

- **Remember, all participants in your group have the same multiplier but independently drawn values of good 1.**

- Your decision screen will be displayed like this:

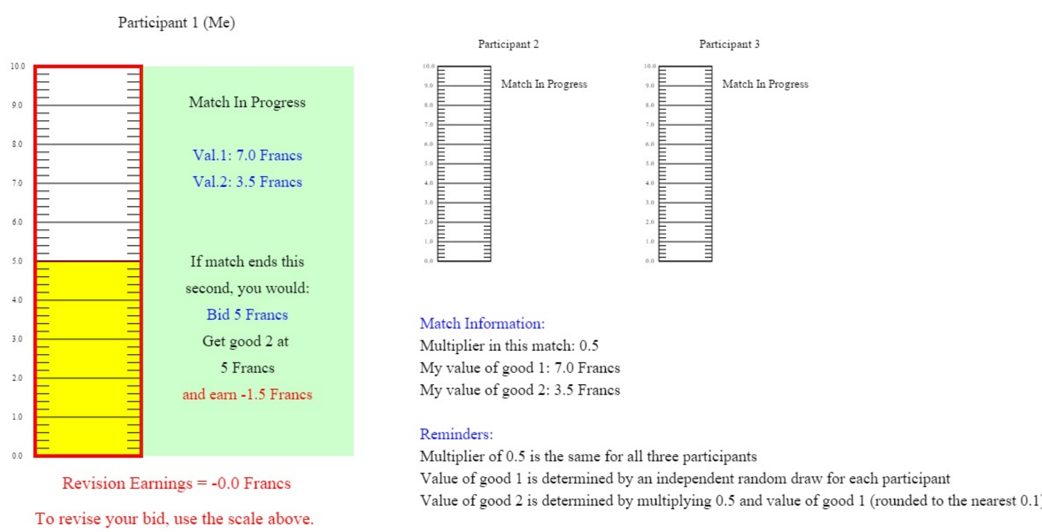

- Your task will be to bid for a good using a slider on the left side of your screen.

- **You may buy only one good** - good 1 or good 2. The outcome of which good you buy depends on your bid and the bids of the participants that you are matched with as follows:
  - **If your bid is the highest** among the three bids, then
    - You will get good 1.
    - You will pay the amount equal to the second highest bid.
  - **If your bid is the second highest** among the three bids, then
    - You will get good 2.
    - You will pay the amount equal to the third highest bid.
  - **If your bid is the third highest** among the three bids, then
    - You will get neither good 1 nor good 2.
    - You will not have to pay anything.

- Note: in case of a tie ranks are determined randomly.

Additional Information about Matches

- Before every match, you will have time to review the information provided about that match.
- **When you are ready to begin** move the mouse inside the scale and wait for the match to begin.
- Mouse clicks are disabled. **You place your bid by moving the mouse** in and out of the scale rectangle on either the left of the right side.
- You will be able to **place the initial bid at no cost**.
- For the duration of the match, you will be able to make as many revisions to your bid as you want. However, the **revisions are costly**. Specifically, you will incur a **cost of 0.1 Francs per revision**.
- Current outcomes will be displayed continuously, but the outcome "that counts" is the one selected when the match ends.
- The duration of each match will not be known ahead of time. The duration of each match was determined randomly using the following procedure.
  - Each match will last at least 20 seconds. After the first 20 seconds are up, each second, a number will be chosen randomly from the set of numbers $\{1, 2, 3, \ldots, 98, 99, 100\}$, where **each number is equally likely**.
  - If the number is 1, then the match will end.
  - If the number not 1, then the match will last an additional second.
  - The number will always be placed back into the set after it is drawn.
  - Thus, after the first 20 seconds, any additional second there is a 1% CHANCE that the match will end and a 99% CHANCE that the match will continue.
  - Therefore, the expected number of seconds in each match will be 120, which means that the expected length of each match is **two minutes**.
  - You will not see the number selected from $\{1, 2, 3, \ldots, 98, 99, 100\}$.
  - To ensure that the length of the match is not dependent on your play, the number of seconds for each match has been written on the board before the experiment, and will be uncovered at the end of the experiment.
- After every match, you will have 30 seconds to review the summary for the match.

Experiment Earnings

- Your earnings in each match is equal to the value of the good that you get minus the amount that you pay for that good and minus any revision costs.
- Your earnings in each round may be positive, negative, or zero depending on your bid and the bids by the participants with whom you are matched.
- Your earnings for the experiment will be the sum of the earnings for each match plus the starting 10 Francs.

- Your cumulative earnings for the experiment (including the 10 Francs) will be displayed at the top of your screen.
- At the end of the experiment, you will be paid in **cash**.

Interactive Instructions

- At the beginning of each match, to indicate that you are ready to begin, you will need to move the mouse over the scale. The match will start when all participants are ready.
- Remember, mouse clicks are disabled throughout this experiment.

Task 1 Move your mouse over the scale on the left side of the screen.

- To set your bid you need to HORIZONTALLY MOVE the mouse outside the scale at current bid. Each bid revision will cost 0.1 Francs.

Task 2 Set your bid to $5.0.

- Information pertaining to each match will be summarized in the middle of the screen (see example below). You can review this information before you choose to start the match.
- Once you have started, you will be able to revise your bids at any moment unitl the match ends.
- Match Information (EXAMPLE)
  - Multiplier in this match: 0.5
  - My value of good 1: 7.0 Francs
  - Ma value of good 2: 3.5 Francs
- Reminders (EXAMPLE)
  - Multiplier of 0.5 is the same for all three participants
  - Value of good 1 is determined by an independent random draw for each participant
  - Value of good 2 is determined by multiplying 0.5 and value of good 1 (rounded to the nearest 0.1)

Task 3 Revise your bid to $6.0.

- Remember, duration of each match is RANDOM. Specifically, each match will last at least 20 seconds, but after the first 20 seconds, each additional second, there will be a 1% chance that the match will end and 99% chance that the match will continue. Therefore, the expected length of each match is 20+100=120 seconds or 2 minutes. However, some matches will be shorter and some matches will be longer due to chance.

Task 4 Revise your bid to $4.0.

- The actual decision screen will include decision boxes for the participants with whom you are matched. If you change your bid or one of the participants that you are matched with changes his/her bid, your earnings for the match may change. Thus, even if you don't change your bid, revisions by other participants may result in different earnings. Note that presented outcome will be for the case of the match ending that second.

Task 5 Revise your bid to $5.0.

Examples

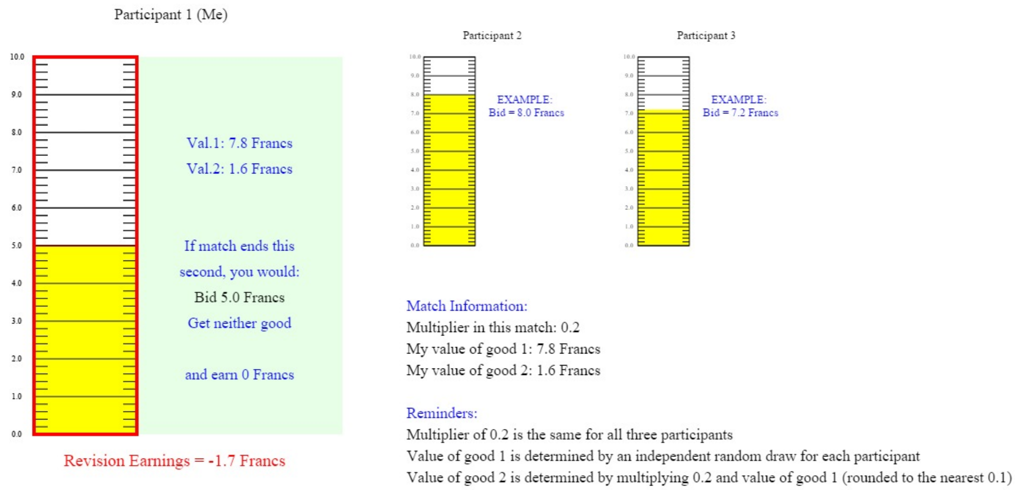

Example Explanation (Change your bid to see how explanation changes).

- Your value of good 1 was randomly drawn to be 7.8 Francs
- Your value of good 2 was obtained by multiplying 0.2 and 7.8
- Your current bid is 5.0 Francs
- Your current bid is the third highest among the three. Therefore, you get neither good 1 nor good 2.
- So if the match were to end this second, you would earn 0 Francs minus revision earnings.

Notes:

1. In the actual experiment, you will not see bids by the participants that you are matched with.
2. Everyone will be able to revise their bids at any moment in time.
3. You will see your earnings for the case if the match were to end this second.
4. Your earnings may be positive, negative, or zero depending on action by you and the participants that you are matched with.

[Click Here for Another Example]

C Additional Figures and Tables

Table C.1: Match Duration (Seconds).

Match:	1	2	3	4	5	6	7	8	9	10
Number of Seconds:	37	166	172	108	176	181	25	145	146	57

Table C.2: Subject Bid-to-Value Ratios in Sessions 1-4

	(a) Highest-valued			(b) Medium-valued			(c) Lowest-valued		
	$C = 0.0$		$C = 0.1$	$C = 0.0$		$C = 0.1$	$C = 0.0$		$C = 0.1$
$\alpha = 0.2$	0.964 (0.04)	$\overset{0.302}{\sim}$	0.912 (0.027)	0.955 (0.05)	$\overset{0.602}{\sim}$	0.916 (0.051)	2.359 (0.331)	$\overset{0.303}{\sim}$	1.887 (0.311)
	$\gg^{0.022}$		$\gg^{0.048}$	$\sim^{0.346}$		$\sim^{0.213}$	$\sim^{0.591}$		$\sim^{0.546}$
$\alpha = 0.5$	0.83 (0.039)	$\overset{0.577}{\sim}$	0.794 (0.05)	0.876 (0.063)	$\overset{0.591}{\sim}$	0.838 (0.033)	2.606 (0.298)	$\overset{0.019}{\gg}$	1.642 (0.247)
	$\sim^{0.433}$		$\sim^{0.605}$	$\sim^{0.667}$		$\sim^{0.963}$	$\sim^{0.163}$		$\sim^{0.444}$
$\alpha = 0.8$	0.765 (0.076)	$\overset{0.933}{\sim}$	0.757 (0.047)	0.835 (0.065)	$\overset{0.946}{\sim}$	0.841 (0.055)	1.991 (0.289)	$\overset{0.113}{\sim}$	1.363 (0.232)
	$\ll^{0.019}$		$\lll^{0.006}$	$\sim^{0.15}$		$\sim^{0.337}$	$\sim^{0.431}$		$\sim^{0.217}$
$\alpha = 0.2$	0.964 (0.04)		0.912 (0.027)	0.955 (0.05)		0.916 (0.051)	2.359 (0.328)		1.887 (0.307)
Average:	0.861 (0.031)	$\overset{0.397}{\sim}$	0.827 (0.026)	0.893 (0.035)	$\overset{0.562}{\sim}$	0.867 (0.027)	2.348 (0.181)	$\overset{0.006}{\ggg}$	1.655 (0.161)

Notes: **Panel (a)** presents average bid to value ratio per match across subjects with the highest private value in each group. **Panel (b)** presents average bid to value ratio per match across subjects with the second highest private value in each group. **Panel (c)** presents average bid to value ratio per match across subjects with the lowest private value in each group. Bootstrapped standard errors are in parentheses. $>$, $\gg$, and $\ggg$ denote significance at 0.10, 0.05, and 0.01 levels, respectively. P-values are determined using two-tailed permutation tests.

Table C.3: Subject Bid-to-Value Ratios in Sessions 5-8

	(a) Highest-valued			(b) Medium-valued			(c) Lowest-valued		
	$C = 0.0$		$C = 0.1$	$C = 0.0$		$C = 0.1$	$C = 0.0$		$C = 0.1$
$\alpha = 0.2$	0.94 (0.05)	$\overset{0.44}{\sim}$	0.889 (0.04)	0.946 (0.055)	$\overset{0.608}{\sim}$	0.91 (0.035)	3.294 (0.65)	$\overset{0.002}{\ggg}$	1.45 (0.209)
	$\overset{0.223}{\sim}$		$\overset{0.107}{\sim}$	$\overset{0.842}{\sim}$		$\overset{0.756}{\sim}$	$\overset{0.416}{\sim}$		$\overset{0.149}{\sim}$
$\alpha = 0.5$	0.848 (0.052)	$\overset{0.418}{\sim}$	0.793 (0.039)	0.963 (0.062)	$\overset{0.403}{\sim}$	0.888 (0.059)	5.023 (1.736)	$\overset{0.0}{\ggg}$	1.041 (0.173)
	$\overset{0.049}{\gg}$		$\overset{0.252}{\sim}$	$\overset{0.001}{\ggg}$		$\overset{0.129}{\sim}$	$\overset{0.361}{\sim}$		$\overset{0.524}{\sim}$
$\alpha = 0.8$	0.695 (0.052)	$\overset{0.864}{\sim}$	0.709 (0.058)	0.659 (0.057)	$\overset{0.298}{\sim}$	0.751 (0.065)	3.138 (1.04)	$\overset{0.031}{\gg}$	1.224 (0.208)
	$\overset{0.002}{\lll}$		$\overset{0.02}{\ll}$	$\overset{0.0}{\lll}$		$\overset{0.038}{\ll}$	$\overset{0.906}{\sim}$		$\overset{0.474}{\sim}$
$\alpha = 0.2$	0.94 (0.049)		0.889 (0.041)	0.946 (0.055)		0.91 (0.035)	3.294 (0.647)		1.45 (0.208)
Average:	0.828 (0.032)	$\overset{0.483}{\sim}$	0.797 (0.029)	0.856 (0.038)	$\overset{0.901}{\sim}$	0.85 (0.032)	3.818 (0.713)	$\overset{0.0}{\ggg}$	1.239 (0.114)

Notes: **Panel (a)** presents average bid to value ratio per match across subjects with the highest private value in each group. **Panel (b)** presents average bid to value ratio per match across subjects with the second highest private value in each group. **Panel (c)** presents average bid to value ratio per match across subjects with the lowest private value in each group. Bootstrapped standard errors are in parentheses. $>$, $\gg$, and $\ggg$ denote significance at 0.10, 0.05, and 0.01 levels, respectively. P-values are determined using two-tailed permutation tests.

Table C.4: GSP Efficiency

	Sessions 1–4				Sessions 5–8		
	$C = 0.0$		$C = 0.1$		$C = 0.0$		$C = 0.1$
$\alpha = 0.2$	0.936 (0.029)	$\overset{0.567}{\sim}$	0.958 (0.028)	$\alpha = 0.2$	0.916 (0.031)	$\overset{0.064}{<}$	0.976 (0.009)
	$\overset{0.939}{\sim}$		$\overset{0.52}{\sim}$		$\overset{0.78}{\sim}$		$\overset{0.188}{\sim}$
$\alpha = 0.5$	0.932 (0.031)	$\overset{0.15}{\sim}$	0.984 (0.009)	$\alpha = 0.5$	0.927 (0.033)	$\overset{0.025}{\ll}$	0.99 (0.004)
	$\overset{0.327}{\sim}$		$\overset{0.259}{\sim}$		$\overset{0.43}{\sim}$		$\overset{0.803}{\sim}$
$\alpha = 0.8$	0.972 (0.016)	$\overset{0.07}{<}$	0.997 (0.002)	$\alpha = 0.8$	0.963 (0.022)	$\overset{0.447}{\sim}$	0.981 (0.013)
	$\overset{0.379}{\sim}$		$\overset{0.204}{\sim}$		$\overset{0.239}{\sim}$		$\overset{0.779}{\sim}$
$\alpha = 0.2$	0.936 (0.028)		0.958 (0.027)	$\alpha = 0.2$	0.916 (0.03)		0.976 (0.01)
Average:	0.944 (0.016)	$\overset{0.083}{<}$	0.978 (0.011)	Average:	0.935 (0.017)	$\overset{0.007}{\lll}$	0.982 (0.005)

Notes: Unit of observation is a matched group of subjects. Bootstrapped standard errors are in parentheses. $>$, $\gg$, and $\ggg$ denote significance at 0.10, 0.05, and 0.01 levels, respectively. P-values are determined using two-tailed permutation tests.