6-22-2023

Discrete Rule Learning in First Price Auctions

Jason Shachat

Lijia Wei

Discrete Rule Learning in First Price Auctions

Comments

ESI Working Paper 23-07

Discrete Rule Learning in First Price Auctions*

Jason Shachat[†] Lijia Wei[‡]

June 22, 2023

Abstract: We present a hidden Markov model of discrete strategic heterogeneity and learning in first price independent private values auctions. The model includes three latent bidding rules: constant absolute mark-up, constant percentage mark-up, and strategic best response. Rule switching probabilities depend upon a bidder's past auction outcomes and a myopic reinforcement learning dynamic. We apply this model to a new experiment that varies the number of bidders and the auction frame between forward and reverse. We find the proportion of bidders following constant absolute mark-up increases with experience, and is higher when the number of bidders is large. The primary driver here is subjects' increased propensity to switch strategies when they experience a loss (win) reinforcement when following a strategic (heuristic) rule.

Keywords: private value auction; discrete heterogeneity; learning; hidden Markov model; laboratory experiment

JEL Classification Numbers: D44; C72; C92; D87; C15

*This paper supercedes the working paper Shachat and Wei (2016),"Discrete Rule Learning and the Bidding of the Sexes."

[†]e-mail:*jason.shachat@durham.ac.uk*. Durham University Business School

[‡]e-mail:*ljwei.whu@gmail.com*. School of Economics and Management, Wuhan University

1 Introduction

Auctions are an enduring topic in economics because of their common usage and amenability to various modes of inquiry. There is a particularly extensive literature documenting and modeling how bidding behavior deviates from the predictions of standard theory in first price sealed bid auctions experiments.[1] A latter strand in this literature explores mixture models in which a bidder may follow one of several alternative bidding rules of varying strategic sophistication. We introduce a dynamic mixture model that allows a bidder to change his rule in response to past auction outcomes. We also report on a new experiment that allows identification, previously not possible, of certain bidding rules. We the estimate the dynamic mixture model with this new data. This exercise generates insights into when and why bidders increase their use of simple bidding heuristics. Surprisingly, the proportion of bidders following simple heuristics grows over time when there are standard shocks to the environment, like an increase in the number of bidders.

Studies of symmetric independent private value (IPV, hereafter) first price sealed bid auctions by Crawford and Iriberri (2007), Shachat and Wei (2012) and Kirchkamp and Reiß (2019) introduce mixture bidding models with distributions of strategic and non-strategic bidders.[2] Crawford and Iriberri (2007) formulate a Level-k model, where k indicates the number of steps of iterated best response a bidder performs when selecting a strategy. They consider two non-strategic $k = 0$ rules: bid one's value or bid randomly according to a uniform distribution over the interval from the minimum allowable bid to value. A $k = 1$ type believes all other bidders follow a particular $k = 0$ strategy and best responds. Correspondingly, a $k = 2$ type believes all other bidders are $k = 1$ and best responds, and so on. Applying this model to the first five rounds of bidding in the IPV first price auction experiments of Goeree and

[1] See Kagel (1995) and Kagel and Levin (2011) for authoritative reviews of this literature.

[2] Isaac et al. (2012) also estimate a mixture model of reduced form linear bidding strategies for data from an IPV first price sealed bid auction experiment with an unknown number of bidders.

Holt (2002), they found approximately 4%, 76%, and 20% of the subjects followed the level $k = 0, 1$ and 2 rules, respectively. Kirchkamp and Reiß (2019) introduce a mixture model with two types. One type simply bids a fixed markdown of his valuation; the other is rational and best responds taking account of the proportions of bidders following the markdown rule and those also best responding. Kirchkamp and Reiß test their model in an experiment allowing bids less than the lowest possible valuation and find roughly 30% of the subjects follow the markdown rule.

Shachat and Wei (2012) extend the previous static approach by allowing rule switching according to a first order Markov process with exogenous transition probabilities. They estimate the model using data from an experiment on first price sealed bid reverse auctions (bidders attempt to sell, rather than purchase an object). The mixture model consists of two simple pricing rules of thumb suggested by Baumol and Quandt (1964): bidding a constant absolute mark-up of one's cost and bidding a constant percentage mark-up of one's cost. The third bidding rule is to best respond to the mixture probabilities and mark-up parameters. Like previous studies, they find initial high frequencies of strategic bidding in early auctions with approximately 75% of the subjects following the best response rule. Surprisingly this percentage quickly falls to a steady state of approximately 62% and the percentage of absolute mark-up bidding rises to over 30%. Skepticism of this result is natural; the model lacks a behavioural mechanism explaining this learning to bid irrationally.

Our study extends this approach in two ways. First, we endogenize the transition probabilities of rule switching as a function of how bidders react to *ex post* auction outcomes. We do this through myopic reinforcement learning dynamics similar to those introduced into strategic decision making by Erev and Roth (1998). Specifically, the relative attractiveness of the currently adopted rule adjusts when (1) the bidders wins the auction or (2) the bidder loses but could have profitably won with an alternative bid.[3] We call this extension the Hidden

[3] This post auction sub-optimality of a bid is often referred to as *ex post* regret. There is a literature which models the auction participants anticipating ex-post regret outcomes, incorportes these in their expected

Markov Bidding Model, or simply the HMBM. In our second extension, we add a fourth rule. Within the HMBM, the realization of a bidder's current rule realization is independent of his current realized value. The bidder has no discretion over his cognitive engagement with the bid formation process in the time between receiving his value and submitting his bid. Our fourth rule introduces this flexibility: a bidder who uses a simple heuristic for low quality values and bids more strategically for high quality high values.

Chen and Plott (1998) study the use of linear bidding heuristics in first price IPV auctions. While predating the mixture models literature, it remains relevant. In their setting, valuations are not uniformly distributed and the symmetric Bayesian Nash equilibrium strategy is non-linear, providing discrimination between strategic and heuristic behaviour. Their statistical methodology conducts "horse races" between several models: percentage markdown, an absolute and percentage markdown hybrid, the symmetric equilibrium strategy and a piecewise linear rule that uses a markdown rule for values below the knot and a more aggressive markdown above. Hypothesis tests always select the equilibrium strategy over linear markdown strategies. However, the hybrid markdown strategy is always selected over the equilibrium one. Chen and Plott's interpret this as the subjects behaving strategically as the slopes the bidding strategy changes with the slope of the valuation probability density function. This is consistent to maximizing expected utility against a belief regarding the distribution of bids - just not the equilibrium belief. When we extend the HMBM to include a fourth rule, it has a similar structure to the piecewise linear rule of Chen and Plott, but with a different interpretation. When a bidder receives a high quality values they bid strategically, otherwise he uses a simple absolute markup rule.

There is a modest literature studying how individuals learn to bid over time in first price auctions. The most common approach is the directional learning framework of Selten and

utility calculations (often called anticipated regret utility) and derives new optimal behaviour and Nash equilibria (Filiz-Ozbay and Ozbay, 2007; Engelbrecht-Wiggans and Katok, 2007, 2008; Ratan and Wen, 2016; Shachat and Tan, 2023).

Buchta (1998). Directional learning is a behavioral principle where individuals adjust their strategies toward those offering ex-post higher payoffs conditional upon available information. Studies such as Selten and Buchta (1998); Guth et al. (2003); Neugebauer and Selten (2006) estimate how bidding rules adjust according to the *ex post* information provided on auction outcomes. These models, unlike ours, have individuals calculating the counterfactual payoffs for strategies not played.[4]

A few studies take a structural approach to estimate learning in auctions. Bajari and Hortacsu (2005) find that econometric models which assume subjects follow an adaptive learning process do not outperform equilibrium models with risk aversion in recovering structural parameters. Studies such as James and Reagle (2009) and Dittrich et al. (2012) attempt to model learning, but find estimates unreliable, using Camerer and Ho's (1999) experienced weighted attraction model (EWA). Grundl and Zhu (2022) approach the overbidding in experimental first price auctions as an identifying restriction for the set various heterogeneous bidding rules participants to estimate bounds on revenue and optimal reserve prices.

The HMBM is an example of a growing literature that not only evolves from a lens of horse races to identify a single representative behavioural rules to ones that accept rule heterogeneity and identify their distribution among a population, but further to identifying how this distribution of rules adjusts dynamically. The structure is typically a set of choice rules that are only followed stochastically, rendering them unobservable, coupled with a Markov process governing the switching of rule assignment to individuals. This structure is called a hidden Markov model (HMM). In behavioural economics, the earliest uses of HMM were the discussed Shachat and Wei (2012) study, and Ansari et al. (2012) who study the evolution of adopted learning rules in repeatedly played games. Shachat et al. (2015) use a HMM as a method to detect when individuals are playing pure or strictly mixed strategies. In a study of the asymmetric coordination Hawk-Dove game, Maruotti et al. (2022) use a

[4]More recently, studies have examined how emotional responses to how first price auction feedback moderates bidding behaviour (Breaban et al., 2022; Karmeliuk et al., 2022).

HMM to study the dynamic impact of individual adaptive dynamics and the nature of game role endowment on equilibrium selection. In an investigation of play in Hide-Seek games played on a computer using eye-tracking, Li and Camerer (2021) use a HMM to estimate the transitions between more and less cognitively sophisticated modes of decision-making. In an innovative study investigating how individuals update subjective judgements after the arrival of relevant information, Alós-Ferrer and Garagnani (2023) formulate a HMM in which individuals switch among rules including Bayes rule and other heuristics.

We estimate the HMBM with data from a new experiment that involves two within-subject treatments. Each experimental session consists of 18 subjects participating in a sequence of 100 n-bidder first price IPV auctions. The first treatment variable is the auction frame. Fifty of the auctions are forward auctions, the auctioneer sells an object, and the other fifty are reverse auctions, the auctioneer purchases an object. This treatment identifies differences in the bidding rules of thumb used when in the role of a buyer versus when in the role of the seller. While the forward and reverse framing are analogous in standard game theoretic models, the percentage mark-up rule generates quite different outcomes and earnings in the two frames. The second treatment variable is the number of bidders n. Half of the time $n = 3$, and the other half of the time $n = 6$. This variation creates identification between the best response rule and the percentage mark-up rule in the forward auction framing.

The estimated HMBM provides the following insights.

1. When the number of bidders is larger, there is a corresponding reduction in the proportion of best response bidders and increase in the proportion of absolute mark-up bidders. This moderates the convergence to competitive pricing.
2. The probability that a subject switches from the constant absolute mark-up rule to the best response rule spikes after a bidder wins the auction. This happens more frequently in the three-bidder auctions than the six bidder auctions.
3. The probability that a subject switches from the best response rule to a mark-up rule spikes after the subject loses an auction but could have profitably won at a more competitive bid. This happens more often in six bidder auctions.
4. The size of absolute mark-ups is quite small -0.56 in forward auctions and even sta-

tistically smaller in magnitude, 0.48, in reverse auctions. The confidence intervals for these estimates are narrow.

5. The size of the percentage mark-up is approximately 20% in both forward and reverse auction frames, but both have wide confidence intervals.
6. Incorporating a fourth bidding rule that allows the rationality of the rule that bids with an absolute mark-up for low quality types and with a best response for high quality types does not improve measurements of fit over the three-rule HMBM.

We present the HMBM in the next section. Then we describe the experimental design. In the penultimate section, we present the estimated bidding model and the bulk of the empirical results. We offer discussion of the scope of our results and extensions in the conclusion. There are four appendices: the experimental instructions, one presenting some of the detailed theoretical analysis, another detailing the Gibbs Sampler and Markov Chain Monte Carlo method used to conduct the Bayesian estimation of the HMBM and finally estimation results for the four-model HMBM.

2 The Hidden Markov Bidding Model

We consider the setting in which a bidder participates in a sequence of single object auctions, indexed $t = 1, \dots, T$. In each auction there are n_t bidders, indexed by i. The auction frame, f_{it}, denotes whether the auction is a forward (F) or reverse (R) one. While the number of bidders and the frame may vary within the auction sequence, these values are always common knowledge. A bidder's type in an auction period, denoted v_{it}, is private information. In a forward action v_{it} is bidder i's value for the object, and in a reverse auction it is his cost. Each v_{it} is an independent draw from the uniform distribution on $[L, H]$. Bidders simultaneously submit bids in a forward (reverse) auction; the one submitting the highest (lowest) bid purchases (sells) the object and pays (receives) the amount of his bid. The winning bidder's payoff is the amount of realized consumer (producer) surplus, and losing bidders' payoffs are zero. Bidders are myopic - only concerning themselves with the current auction payoff - and types are drawn anew each auction.

The HMBM consists of three components. First, there is a finite set of latent linear bid rules mapping from the auction frame and bidder type to bid amount. This set may contain rules both derived strategically and those representing simple heuristics. Rules are latent because they are followed imprecisely, prescribed bid amounts are submitted with some random perturbations. The second component is an exogenous multinomial distribution governing the initial assignment of bidders to bidding rules. The third component is a first order Markov matrix of transition probabilities governing the switching of rules. These transition probabilities are functions of a bidder's previous auction participation outcome.

2.1 The set of latent bidding rules

We assume the set of latent bidding rules contains three elements, $\{AM, PM, BR\}$ with generic element s, each reflecting a distinct behavioral heuristic. A constant absolute mark-up (AM) bidder always demands a fixed surplus independent of his value. The AM bidding rule is the affine function with slope of one,

$$b_{AM}(v_{it}|f_{it}) = \kappa_{f_{it}} + v_{it} + \epsilon_{AM_{it}}.$$

We specify that the markup parameter $\kappa_{f_{it}}$, sometimes expressed more compactly as κ, should be negative for forward frames and positive for reverse frames. It is a linear combination of the socioeconomic effects conditional on the frame. We don't impose that $\kappa_{f_{it}}$ has the same magnitude for both two auction frames. Also notice that the AM bidding rule does not depend up the number of bidders nor the distribution of private types. Finally, $\epsilon_{AM_{it}}$ is a heteroscedastic independent random perturbation following the normal distribution with mean of zero and variance of σ^2_{AM}.

A constant percentage mark-up (PM) bidder always demands surplus that is a fixed percentage of his realized type. Thus, the PM bidding rule, conditional on the frame is,

$$b_{PM}(v_{it}|f_{it}) = \big(1 + \rho_{f_{it}}\big)v_{it} + \epsilon_{PM_{it}}$$

Again we allow the possibility that percentage mark-up $\rho_{f_{it}}$, sometimes compactly denoted ρ, may differ in sign and magnitude in forward and reverse auctions. The PM rule, like the AM rule, does not depend up the number of bidders nor the distribution of private types. Finally note that this rule is adopted imperfectly with the heteroscedastic independent normally distributed perturbation $\epsilon_{PM_{it}}$ with a mean of zero and variance of σ^2_{PM}.

The final latent bidding rule is the strategic best response or BR. A bidder adopting the BR rule maximizes his expected utility conditional upon his realized type, the mark-up parameters κ and ρ, the number of bidders n, and his beliefs regarding the rules each of the bidders is currently adopting. Regarding these beliefs, let π_{AM} be the probability any bidder is a AM bidder, π_{PM} be the probability any bidder is a PM bidder, and that $(1-\pi_{AM}-\pi_{PM})$ is the probability any bidder is a BR bidder.[5] Proceeding with the approach of Cox et al. (1982), we assume each bidder has the von Neumann-Morgenstern expected utility function $V(y) = \eta y^{\frac{1}{\eta}}$, where y is a non-negative change in wealth and η is his constant coefficient of relative risk aversion. The BR bidding rule in the forward framing, specifies an expected utility maximizing bid for each possible realization of her value:

$$b_{BR}(v_{it}|f_{it}) = \arg\max_{b} \int_{L}^{H} V(b - v_{it})dF(b; \pi_{AM}, \pi_{PM}, \kappa, \rho)^{n-1},$$

where $F(x, ; \pi_{AM}, \pi_{PM}, \kappa, \rho)^{n-1}$ is the culmative distribution function of the maximum order statistic of n bids given the HMBM parameters. When the distribution of bidders' types is uniform, there is a linear function that is a best response, $\bar{b}(v)$, to any mixture over a finite set of linear increasing bidding functions B such that the $\min\{b(H)|b \in B\} \geq \bar{b}(H)$. When

[5] In the repeated auction context we can think of these beliefs as a state variable. While the myopia assumption allows us think of a bidder as only concerned about his current auction payoff, the formation of these beliefs can involve a complicated inference problem depending upon the information feedback provided in the auction. In our experiment, subjects are randomly rematched into new bidding cohorts every period which eliminates this conditional inference problem.

$b' = \text{argmin}\{b(H)|b \in B\}$ is used with positive probability and $b'(H) < \bar{b}(H)$, then $\bar{b}(v)$ is no longer a best response to this mixture for the range of values $(\bar{b}^{-1}(b'(H)), H)$. Intuitively, for valuations in the this range there is less competition for the object and it is optimal to bid more aggressively. This introduces a kink point in the best response function at the valuation, $v^* = \bar{b}^{-1}(b'(H))$. For the HMBM there are six potential orderings of $b_{BR}(H), b_{AM}(H)$ and $b_{PM}(H)$, but in our data analysis we only observe the ordering $b_{PM}(H) < b_{BR}(H) < b_{AM}(H)$. For this ordering we find, and provide details on its derivation in Appendix B, the BR bidding rule for the forward frame is:

$$b_{BR}(v_{it}|f_{it} = F) = \begin{cases} \dfrac{(L + \pi_{AM}\kappa)(1+\rho)}{\big[1 + \rho(1-\pi_{PM})\big]M} + \dfrac{M-1}{M}v_{it} + \epsilon_{BR_{it}} & \text{if } v_{it} \leq v^*_{Fn} \\ \dfrac{G_F(v^*)[H(1+\rho) - v^*]}{G_F(v_i)} + v_{it} - \dfrac{\int_{v^*}^{v_i} G_F(x)\mathrm{d}x}{G_F(v_i)} + \epsilon_{BR_{it}} & \text{if } v_{it} > v^*_{fn} \end{cases},$$

where $G_F(x) = \left(\pi_{PM} + (1-\pi_{PM})\frac{x-L}{H-L}\right)^{M-1}$, $v^*_{Fn} = \frac{H(1+\rho)M}{(M-1)} - \frac{(L+\pi_{AM}\kappa)(1+\rho)}{(1+\rho(1-\pi_{PM}))(M-1)}$ and $M = \eta(n-1)+1$.

The BR bidding rule for reverse frame has a symmetric structure expect the kink occurs when $b_{PM}(L) > b_{BR}(L)$ at the valuation $v^* = \bar{b}^{-1}(b_{PM}(L))$:

$$b_{BR}(v_{it}|f_{it} = R) = \begin{cases} \dfrac{(H + \pi_{AM}\cdot\kappa)(1+\rho)}{(\big[1 + \rho(1-\pi_{PM})\big]M} + \dfrac{M-1}{M}v_{it} + \epsilon_{BR_{it}} & \text{if } v_{it} \geq v^*_{Rn} \\ \dfrac{G_R(v^*)[L(1+\rho) - v^*]}{G_R(v_i)} + v_{it} - \dfrac{\int_{v_*}^{v_i} G_R(x)\mathrm{d}x}{G_R(v_i)} + \epsilon_{BR_{it}} & \text{if } v_{it} < v^*_{Rn} \end{cases}.$$

where $G_R(x) = \left(\pi_{PM} + (1-\pi_{PM})\frac{H-x}{H-L}\right)^{M-1}$ and $v^*_{Rn} = \frac{L(1+\rho)M}{(M-1)} - \frac{(H+\pi_{AM}\kappa)(1+\rho)}{(1+\rho(1-\pi_{PM}))(M-1)}$. The heteroscedastic random perturbation $\epsilon_{BR_{it}}$ is independently normally distributed with a mean of zero and variance of σ^2_{BR}.[6]

[6]It is important to note we assume the BR bidder does not consider the noise terms $\epsilon_{r_{it}}$ when calculating his optimal bid as he would in statistical equilibrium concepts like the Quantal Response Equilibrium (McKelvey and Palfrey, 1995).

The linear portion of the BR rule has two notable features. First, the mark-up rule parameters κ and ρ, and the beliefs about the distribution of bidding types only affect the intercept term. Second, when the number of bidders and the risk attitude are consistent between the forward and reverse frames, the slope terms are the same.

At this point, we have a fully formulated static model, from which we can generate several alternative models with appropriate restrictions on parameters and beliefs. The Bayes-Nash equilibrium model (Vickrey, 1961) occurs when beliefs are $1-\pi_{AM}-\pi_{PM}=1$ and we restrict bidders to be risk neutral, i.e. $\eta=1$. If we instead restrict the constant coefficient of risk aversion to be the same for all bidders and in the open unit interval, we obtain the risk averse Bayes-Nash equilibrium model of Holt (1980). We recover the model of Kirchkamp and Reiß (2019) by setting $\pi_{PM}=0$. Finally, we can obtain a version of the Level-k model of Crawford and Iriberri (2007) by setting $\pi_{AM}=1$ and the absolute mark-up $\kappa=0$.

2.2 Markovian rule switching and auction feedback

In the HMBM, the evolution of strategy adoption begins with an initial assignment of bidding rules according to multinomial distribution Π_1. The variable s_{it} indicates the rule bidder i uses in auction t. From the second auction onward, we assume s_{it} depends upon s_{it-1} and i's outcome from participation in auction $t-1$. We summarize the transition probabilities with the Markov matrix P,

$$P=\begin{pmatrix} \Pr_{BR,BR}(o_{it-1}) & \Pr_{BR,AM}(o_{it-1}) & \Pr_{BR,PM}(o_{it-1}) \\ \Pr_{AM,BR}(o_{it-1}) & \Pr_{AM,AM}(o_{it-1}) & \Pr_{AM,PM}(o_{it-1}) \\ \Pr_{PM,BR}(o_{it-1}) & \Pr_{PM,AM}(o_{it-1}) & \Pr_{PM,PM}(o_{it-1}) \end{pmatrix},$$

where $\Pr_{jk}(o_{it-1})=\Pr(s_{it}=j|s_{it-1}=k,o_{it-1})$ is the transition probability of moving from bidding rule j to bidding rule k, and o_{it-1} is the outcome bidder i receives from his participation in auction $t-1$.

We classify the outcome o_{it}, as one of three possible types; NR, LR, and WR. The NR outcome is *neutral reinforcement*; the bidder loses the auction, but there was no other bid at which he could have won and earned positive surplus. The second outcome is *loss reinforcement* (LR) in which the bidder loses the auction but there was an alternative bid at which he could have won and earned positive surplus. The final potential outcome is *win reinforcement* (WR); the bidder wins the auction.

We quantify these auction outcome effects through a state dependent index indicating the attractiveness of each rule. We assume the auction outcome adjusts the index of a bidder's adopted rule prior to the determination his subsequent rule. Bidder i's rule indices for period t, conditional on $s_{it} = j$, are

$$\Psi_{it}^{jk} = \begin{cases} \gamma_{jk} + \xi_{it}^{jk} & \text{if } j \neq k \\ \gamma_{jk} + I_{LR}(\gamma_{j1} + \gamma_{j2} L_{it}) + I_{WR}(\gamma_{j3} + \gamma_{j4} W_{it}) + \xi_{it}^{jk} & \text{if } j = k \end{cases}, \tag{1}$$

where ξ_{it}^{jk} is an independent standard normal innovation.[7] The dummy variable I_x takes the value of one when $o_{it} = x$ and zero otherwise. Loss or win reinforcements have affine impacts on the index of the adopted strategy. The variables L_{it} and W_{it} are the surplus amounts associated with the respective reinforcements calculated as follows:

$$L_{it} = \begin{cases} v_{it} - p_{it} & \text{if } f_{it} = F \\ p_{it} - v_{it} & \text{if } f_{it} = R \end{cases},$$

where p_{it} is the winning bidding in auction t, and

$$W_{it} = \begin{cases} v_{it} - b_{it} & \text{if } f_{it} = F \\ b_{it} - v_{it} & \text{if } f_{it} = R \end{cases}.$$

[7]Setting the variance term σ_ξ^2 equal to one is without loss of generality, and it allows identification of the transition probability parameters.

We assume bidder i transitions to the bidding rule with the largest index Ψ_{it}^{jk}. This implies that each row of the matrix P is a multinomial probit choice model.

3 Experiment design

We ran all experiments at the Finance and Economics Experimental Laboratory (FEEL) of Xiamen University. We used the ORSEE system (Greiner, 2004) to recruit subjects. All subjects were either undergraduate or master level students from a cross section of schools in the university. Monetary values presented in this section are expressed in experimental currency units (ECU), and final experimental balances were converted to cash payments at an exchange rate of 1 RMB = 10 ECU. The average total earnings per subject, including a 10 RMB show-up fee, was 63.18 RMB. The study consisted of 10 sessions each with 18 subjects. Of the 180 subjects 99 were females. Each session took approximately 2 hours to complete a common sequence of four tasks[8]:

Task 1: collection of a saliva sample which includes reading task specific instructions[9],

Task 2: reading instructions for and participation in 100 auction periods,

Task 3: completion of a survey, and

Task 4: collection of a second saliva sample.

After completing the four tasks, we paid subjects privately as they exited one-by-one.

[8]We provide English translations of the instructions for the four tasks in Appendix A.

[9]We collected saliva samples in order to measure certain sex hormone levels of the participants. This data is part of an investigation of gender differences using sex hormone data extracted from saliva samples is reported in our original working paper Shachat and Wei (2016) and will be reported in separately in a future paper.

3.1 The auction task

Within the 100 auction periods of each session there are two within subject treatments: the auction frame (forward and reverse), and the number of bidders (3 and 6). After period 50, we switch the auction frame and make a public announcement to remind subjects of this. In one-half of the sessions subjects participate in the forward frame first, and the other sessions start with the reverse frame. Within the first 50 auction periods, we vary the numbers of bidders between the first and second blocks of 25 periods. This order is switched in the second half of the session. Table 1 presents the sequences of these treatments for each session. Subjects are randomly matched in new groups each period to limit repeated game effects.

Table 1: The assigned sequence of within subject treatments by experimental session

Session	Periods 1-25	Periods 26-50	Periods 51-75	Periods 76-100
1	Forward; $n=6$	Forward; $n=3$	Reverse; $n=3$	Reverse; $n=6$
2	Forward; $n=3$	Forward; $n=6$	Reverse; $n=6$	Reverse; $n=3$
3	Forward; $n=3$	Forward; $n=6$	Reverse; $n=6$	Reverse; $n=3$
4	Forward; $n=6$	Forward; $n=3$	Reverse; $n=3$	Reverse; $n=6$
5	Forward; $n=3$	Forward; $n=6$	Reverse; $n=6$	Reverse; $n=3$
6	Reverse; $n=6$	Reverse; $n=3$	Forward; $n=3$	Forward; $n=6$
7	Reverse; $n=6$	Reverse; $n=3$	Forward; $n=3$	Forward; $n=6$
8	Reverse; $n=3$	Reverse; $n=6$	Forward; $n=6$	Forward; $n=3$
9	Reverse; $n=3$	Reverse; $n=6$	Forward; $n=6$	Forward; $n=3$
10	Reverse; $n=6$	Reverse; $n=3$	Forward; $n=3$	Forward; $n=6$

Each auction is a first price sealed bid auction. In each period, the computer program[10] informs a subject of the auction frame, the number of bidders, and his value/cost. Each subject is asked to submit a bid in the range of 0 to 60. The subject who bids the highest (lowest) price in each group of bidders wins the forward (reverse) auction and pays (receives) the amount of his bid.[11] After the auction concludes, the computer program informs each

[10] Developed with the Z-tree programming language (Fischbacher, 2007).

[11] In the case of multiple subjects submitting the winning bid amount within a group, one of them is randomly selected to be the auction winner.

subject whether or not he won, the winning price and his payoff in the period. Note during the auction periods, a subject can view his entire past auction experience.

To directly compare the impact of auction frame, we use the symmetrical setting between the forward and reverse formats. For each of the 50 forward auctions, subject i's values are independently drawn from a uniform distribution on the range 20 to 40. Suppose V_i is the vector of these realized draws. For the 50 reverse auctions, subject i's cost vector is generated by the formula $C_i = 40 - (V_i - 20) = 60 - V_i$. So there are 50 pairs of value and cost for each subject. The orders of first 25 elements of V_i and C_i are randomly sorted to form the actual sequence of types in the 3 bidder auctions, and the last 25 are randomly sorted to form the actual sequence of types in the 6 bidder auctions.

In the third task, subjects complete a computerized survey. The survey contains questions about individual characteristics for all subjects and, for female subjects, additional questions about their menstrual cycle.

4 Estimation Results

In this section we present the estimates of the HMBM. We use a Bayesian statistical approach to estimate the unknown parameters of the HMBM. The approach starts by specifying independent and diffuse marginal priors on the parameters. Then we use an iterative Monte Carlo Markov Chain (MCMC) with a Gibbs sampler to generate estimates of the marginal posterior distributions of the parameters.

A key property of this procedure is that the empirical distribution of random parameter draws converges to the true marginal posteriors. We conduct 3000 iterations and establish that the empirical distributions of the drawn parameters has converged (Geweke, 1991). Then we run the procedure an additional 2000 iterations, from which we construct empirical density functions. These empirical density functions are used to calculate the posterior means and confidence intervals reported in this section. In Appendix C, we provide a full description of this statistical procedure.

4.1 Estimated bidding rules

We first consider the estimated posterior means of the AM and PM bidding rule parameters, presented in Table 2 along with respective 95% confidence intervals. Note we use the dummy variables I_F and I_R to indicate Forward and Reverse auction frames respectively. In the forward auction, we see the estimated mark-down for the AM rule is only fifty-six cents. In the reverse auction, the absolute mark-ups are even smaller at forty-eight cents. The average percentage mark-up/mark-down under the PM rule is statistically the same in forward and reverse auction frames, roughly 20%. However, we do note there is a large variance in the distribution of the percentage marks-ups particularly in comparison to the very low estimated variance of absolute mark-ups under AM.

Table 2: Mark-up bidding rule parameters: posterior means with 95% confidence intervals

Rule	Variable	Coef.	95% Confidence interval
AM	$I_F \cdot \kappa$	-0.56	[-0.58,-0.54]
	$I_R \cdot \kappa$	0.48	[0.46,0.50]
	σ^2_{AM}	0.16	[0.15,0.17]
PM	$I_F \cdot \rho$	-0.20	[-0.22,-0.19]
	$I_R \cdot \rho$	0.21	[0.20,0.23]
	σ^2_{PM}	47.76	[43.96,51.79]

In Table 7 we report the estimated BR bidding rule parameters. As we noted when deriving the Best Response rules, the presence of a kink in the BR rule depends on the rank of the bids submitted under the three rules at the most competitive value/cost. We find distinct kinks in the Forward auction frame for the three and six bidder settings, both estimated kinks are well below the highest possible value of forty. This provides a large number of observations to estimate a quadratic approximation of the concave portion of the BR rule beyond the kink. However, in the Reverse auction frame, the estimated kinks for both the three and six bidder settings are close to the lowest possible cost of twenty. The narrow gaps between the lowest cost and the kinks only includes a small number of observations, leading to imprecise estimates of the convex portion of the BR rule. Accordingly we only report estimates of simple linear rules in these two cases.

We first consider the estimated slope terms of the linear portions found in the third row of Table 7. We see that the estimated slopes for the six bidder auctions exceeds those of the three bidder ones. In both auction frames, the change in the estimated slope going from three to six bidders is in the predicted direction; however, the magnitude does not reflect a constant η.[12] Estimated slopes are steeper in Forward frame than in the Reverse frame. Within the HMBM, this implies bidders act "as if" they are more risk averse in forward auctions. Finally, in Forward auctions, we estimate for values above the kink the BR quadratic approximation is close to linear. However, the slope drops significantly in the three bidder auctions but not in six bidder auctions.

4.2 Estimated endogenous rule switching probabilities

We first report the estimates the parameters of attractiveness indices underlying the multinomial probit models of rule transition probabilities in Table 4. A LR event, not its size, reduces the attractiveness of the BR and PM strategies. On the other hand, the size of a WR reduces the attractiveness all three rules. However, for the PM rule this impact is opposed by a spike in the its attractiveness from simply experiencing a WR outcome.

The estimated values of the reinforcement coefficients alone doesn't adequately convey how these outcomes affect the transition probability matrix P. To better convey how the auction

[12] Not finding constant coefficients of relative risk aversion when the number of bidders changes is a leading critique against risk aversion based explanations of overbidding in first price auctions; for example, see Kagel and Levin (1993).

Table 3: Strategic Best Response (BR) bidding rule parameters: posterior means with 95% confidence intervals

	Forward Auction		Reverse Auction	
Variable	n=3	n=6	n=3	n=6
Intercept	1.81	0.42	7.04	5.12
	[1.56,2.06]	[0.16,0.67]	[6.86,7.22]	[4.94,5.30]
v_i	0.88	0.94	0.82	0.89
	[0.87,0.89]	[0.93,0.96]	[0.82,0.83]	[0.88,0.89]
$(v_i - v^*) \cdot I(v > v^*)$	-0.23	-0.05	-	-
	[-0.34,-0.13]	[-0.15,0.04]	-	-
$(v_i - v^*)^2 \cdot I(v > v^*)$	-0.01	-0.03	-	-
	[-0.02, 0.01]	[-0.04,-0.02]	-	-
v^*	33.92	33.28	20.83	21.79
	[33.07,34.52]	[32.51,33.84]	[20.38,21.26]	[21.37,22.19]
σ^2_{BR}	0.77			
	[0.75,0.81]			

outcome impacts P, we calculate the estimated matrix P for each of the three auction outcomes. We report these estimated P's in Table 5.

Inspection of the posterior mean transition probabilities reveal three regularities. First, the continuation probabilities of following the same rule, given in the grey filled cells on the main diagonals, exceed seventy percent, except for the PM rule after a LR outcome. This suggests inertia in subjects' rule adoptions. Second, when subjects transition away from one of the non-strategic rules, AM and PM, it will almost always be to the strategic BR rule. In other words, the probability of switching from one non-strategic rule to another is very low, less than three percent in all reported cases. See the underlined transition probabilities in Table 5. Third, when subjects switch away from the BR rule, they are much more likely to adopt the AM rather than the PM rule.

Table 4: Estimated reinforcement effects on rule attractiveness indices: posterior means with 95% confidence intervals

Variable	BR	C.I.	AM	C.I.	PM	C.I.
I_{WR}	0.04	[-0.18, 0.29]	-0.14	[-0.39, 0.10]	0.58	[0.09, 1.09]
W_{it}	-0.17	[-0.25,-0.09]	-0.66	[-0.98, -0.33]	-0.12	[-0.22, -0.03]
I_{LR}	-0.74	[-0.97,-0.48]	0.11	[-0.31, 0.54]	-0.68	[-1.00, -0.32]
L_{it}	-0.08	[-0.18, 0.03]	0.00	[-0.59, 0.64]	-0.02	[-0.10, 0.06]

The impact of the different reinforcements are also evident in Table 5. Consider the differences in P between neutral versus loss reinforcement outcomes. This leads to an estimated reduction of fourteen and one-half percent in the continuation probability of the BR rule. It is as though a subject who follows the BR rule and loses the auction when he could have won with a lower surplus demanding bid probability and increases his likelihood of switching to one of the simple bidding heuristics. Loss reinforcement disrupts the continued use of the PM strategy, reducing the reported continuation probability from seventy-four to fifty-four percent. There is no statistically discernible LR impact on the rate of continued use of the AM; the estimated small surplus demands of this render LR outcomes rare.

Win reinforcement outcomes largely diminish the continued use of the AM rule, and to a lesser extent the BR rule. When following the AM rule and winning the auction, we estimate that on average the probability to continue following the AM rule is reduced by thirteen percent, falling from eighty-four to seventy-one percent. With respect to the BR rule, win reinforcement leads a decrease continuation probability from ninety-one to eighty-six percent.

4.3 Evolution of bidding rule adoption

We show the estimated HMBM leads to quickly leads to a steady assignment rule adoptions rates over the course of the one hundred auctions, when averaged over the balanced random assignment of treatment, but a striking rise in the use of bidding heuristics arises when the

Table 5: Transition matrices conditional upon auction outcome

	Neutral Reinforcement			**Loss Reinforcement**			**Win Reinforcement**		
	BR_t	AM_t	PM_t	BR_t	AM_t	PM_t	BR_t	AM_t	PM_t
BR_{t-1}	0.910	0.068	0.022	0.755	0.178	0.067	0.855	0.107	0.038
AM_{t-1}	0.155	0.838	0.007	0.138	0.856	0.006	0.270	0.713	0.017
PM_{t-1}	0.230	0.028	0.742	0.404	0.060	0.537	0.231	0.028	0.741

[A] Evaluated at the following average Win and Loss reinforcement levels: BR, $\overline{W} = 2.48$ and $\overline{L} = 1.59$; AM, $\overline{W} = 0.67$ and $\overline{L} = 0.42$; and PM, $\overline{W} = 4.68$ and $\overline{L} = 3.36$.

numbers of bidders is six versus three. One of the key latent variables in the HMBM is the sequence of bidding rule states a subject occupies over the 100 auction periods. The Gibbs sampler generates a realization of this random latent sequence for each subject in every iteration of the MCMC estimation. We calculate the mean posterior estimate of this sequence for each subject as follows. For auction period t, we calculate the proportion of the 2000 realized bidding rule sequences the subject adopts each bidding rule. Then we average element-by-element the appropriate sets of individual sequences to form different aggregate series of interest regarding auction frame and the number of bidders.

Figure 1 presents the estimated one hundred period sequences of auction rule adoption.At the start of the experiment, we can see subjects roughly 40:40:20 adoption ratio for the $BR : PM : AM$ rules. We see there is a precipitous drop-off in the PM adoption rate to ten percent or less within ten periods. We also observe a steady rise of the AM rule adoption rate to the 35-40% range. We also see, after an short initial adjustment, the average proportional use of the BR bidding strategy rises to steady rate around 60%. These initial conditions are quite different than those reported by Shachat and Wei (2012); their U.S. and Singaporean subject pools show much higher initial strategic behavior. But surprisingly the convergence of rule adoption proportions is quite similar. This suggests we have identified robust learning.

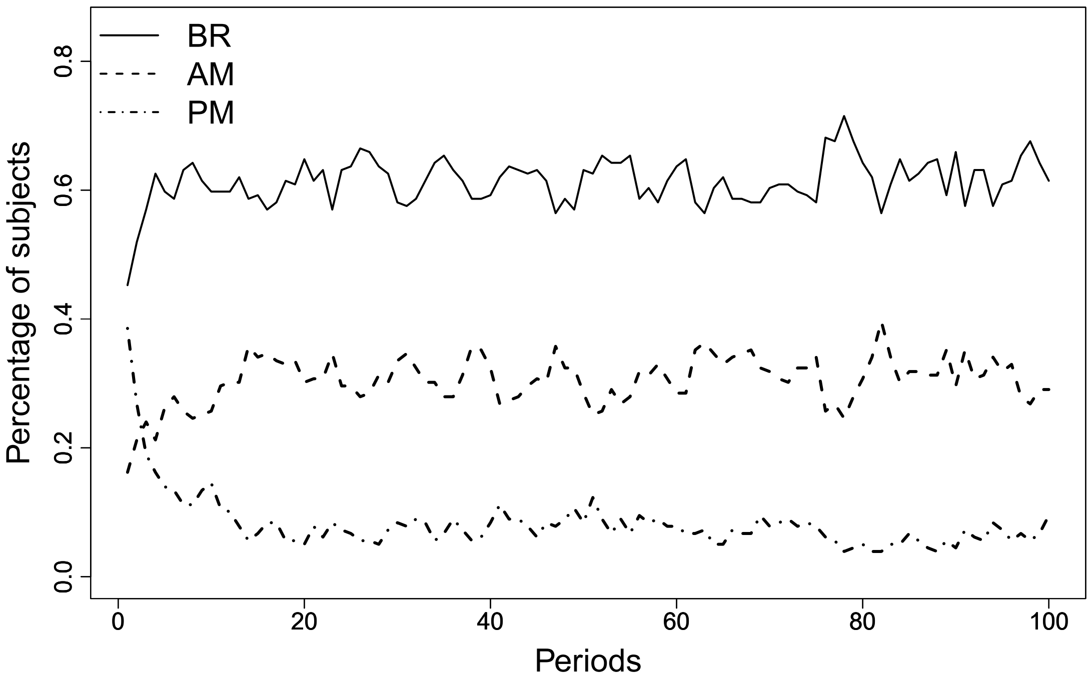

Figure 1: Estimated sequences of subject proportions using the AM, PM, and BR rules

One of this paper's innovations is introducing endogenous transition probabilities conditional upon the reinforcements of auction outcomes. Differences in the bidding rule parameter values and varying auction conditions such as the number of bidders can generate varying patterns of learning and rule adoption. Differing patterns are not possible when the matrix P is exogenous as in Shachat and Wei (2012).

Figure 2 is a 2×2 array of sequence plots that exhibits how the rule learning dynamics vary according to the number of bidders and the auction frame. The four rows in the array correspond to each of the 25 period blocks of the within crossing of the auction frame and the number of bidder treatment variables. In the three bidder auctions we observe fairly stable patterns previously noted when looking at the 100 period sequences. In contrast, the six bidder auctions all exhibit diminishing adoption of the BR and a corresponding increasing the use of the AM. These figures strongly suggest that purely increasing competition will not converge to competitive equilibrium outcomes at the rates suggested by purely strategic bidding models.

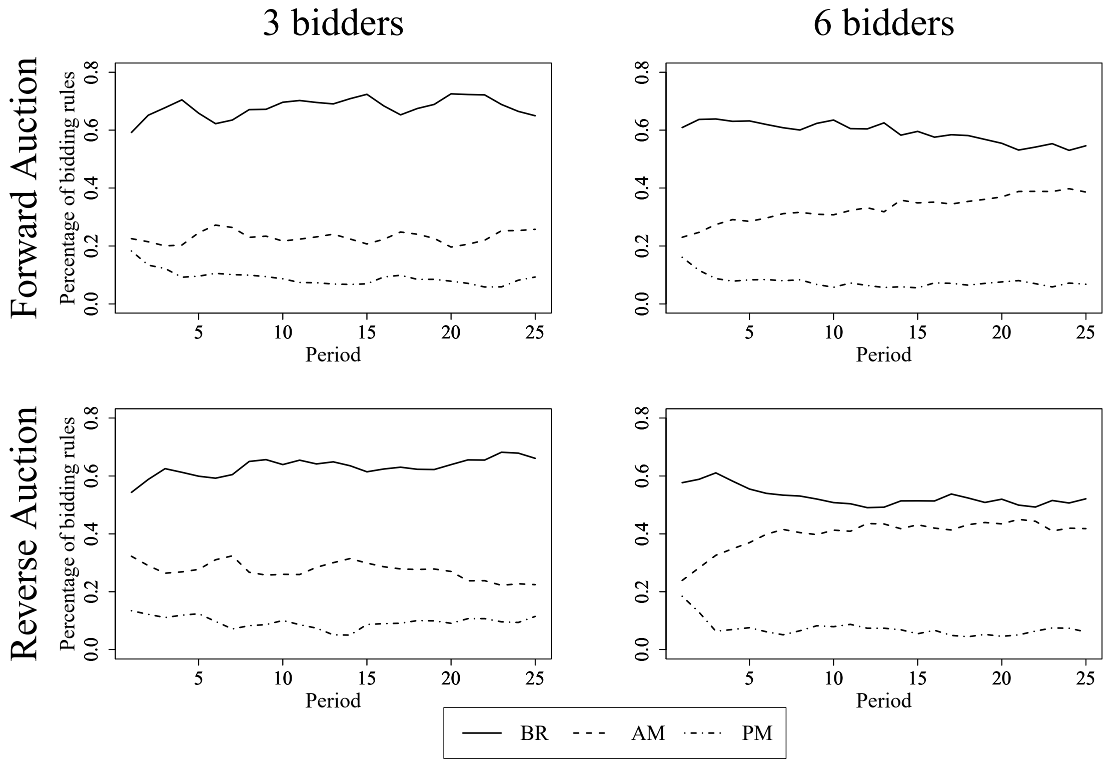

Figure 2: Estimated sequences of subject proportions using the AM, PM, and BR rules

4.4 The independence of bid rule adoption and realized value/cost

The rule one follows in an auction period solely depends upon the rule followed and auction outcome of the previous period is a key assumption of the HMBM. A reasonable challenge to this assumption is that an individual's current rule adoption depends upon his current realized value/cost. We evaluate this challenge by sorting the draws from the posterior of rule sequences, used in the previous subsection, according to the subjects' realized value in costs.

Figure 3 presents a stacked bar graph series, by auction format, for each one experimental dollar increment of the interval of potential value/cost. If one's rule is independent of the current realized private value, then the bar graphs should show show no trends. However,

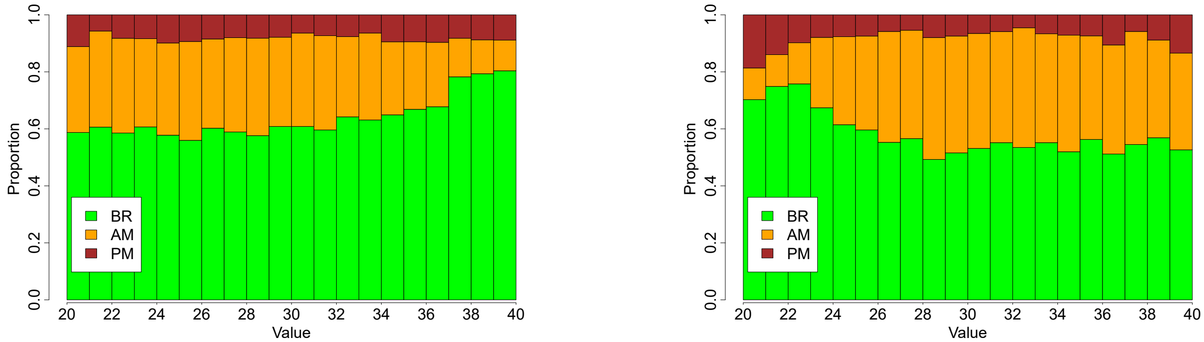

(a) Forward auctions

(b) Reverse auctions

Figure 3: Posterior distribution of bidding rules by value/cost

inspection reveals a potential increase in use of the BR rule for the values and costs for the most favorable three value segments.

4.5 A four-rule HMBM

A key conceptual point of the HMBM is that the employment of one's bidding rule is not conditional upon one's realized value/cost. To address the evidence against this point presented in Figure 3, we reformulate a new HMBM model by adding a fourth rule that, which we call the BA (BR/AM) rule, in which an adopter bids according to the BR for realized values (costs) exceeding (falling short of) a certain threshold, and according to the AM rule for those that don't. According to the trends presented in Figure 3 we set this threshold at thirty-five in the Forward frame and twenty-five in the Reverse frame.

We report in detail the estimation and evaluation of the four-rule model in Appendix D. In summary, over the course of the one hundred periods of auctions, the percentage of BA bidders grows from 10% to approximately 40%. As BA bidders follow the AM rule three-quarters and the BR rule one-quarter of the time, the net result is the proportion of bids generated by the PM, AM and BR rules roughly are 10%, 45% and 45% respectively by then end of an experimental session. Also, the addition of the BA rule does not meaningfully change the estimated parameters of the rules themselves. As for improving model fit, there is less trend in the distribution of bidding rules across values/costs. However, there is almost no improvement in the values of in-sample prediction goodness-of-fit measures.

5 Conclusion

We introduced the HMBM, which provides a dynamic generalization of mixture models of bidding in auctions. Within this flexible framework, we considered the case where auction outcomes impact the transition probabilities between strategic best response rules and simple rules of thumb. We applied the HMBM to a new experiment and found an interesting rule switching dynamic. After experiencing a loss reinforcement, the probability a subject who follows the rational BR rule switches to a rule of thumb spikes. In contrast, when a subject follows AM, the predominantly used rule of thumb, the probability of switching to the rational BR rule spikes after experiencing a win reinforcement.

We don't provide, nor are aware of, a model of optimizing behavior that generates such dynamic responses. However, we conjecture the following. Individuals appreciate there is a potential benefit to choosing a strategic rule. But ascertaining which strategic rule offers the highest reward is a difficult cognitive task, especially in our auction setting. This generates uncertainty in the value of the calculated strategic rule. On the other hand, the value of following a rule of thumb is easy to perceive and cognitively simple to execute. Consequently, a negative reinforcement undermines the bidder's confidence in his ability to successfully bid strategically. This leads to abandoning the pursuit of the less certain-higher reward strategy in favor of the easier to value rule of thumb. In the other direction, a win reinforcement when following the rule of thumb stirs a sense of lost opportunity from not trying to identify and pursue a more profitable strategic rule.

References

Albert, James H and Siddhartha Chib (1993), "Bayesian analysis of binary and polychotomous response data." *Journal of the American Statistical Association*, 88, 669–679.

Alós-Ferrer, Carlos and Michele Garagnani (2023), "Part-time bayesians: Incentives and behavioral heterogeneity in belief updating." *Management Science*. Forthcoming.

Ansari, Asim, Ricardo Montoya, and Oded Netzer (2012), "Dynamic learning in behavioral games: A hidden markov mixture of experts approach." *Quantitative Marketing and Economics*, 10, 475–503.

Bajari, Patrick and Ali Hortacsu (2005), "Are structural estimates of auction models reasonable? evidence from experimental data." *Journal of Political Economy*, 113, 703–741.

Baumol, William J. and Richard E. Quandt (1964), "Rules of thumb and optimally imperfect decisions." *The American Economic Review*, 54, 23–46.

Breaban, Adriana, Cary Deck, and Erik Johnson (2022), "Emotional differences between isomorphic auctions." *Journal of Behavioral Finance*, 23, 427–437.

Camerer, Colin F. and Teck-Hau Ho (1999), "Experience-weighted attraction in games." *Econometrica*, 67, 827–874.

Chen, Kay-Yut and Charles R Plott (1998), "Nonlinear behavior in sealed bid first price auctions." *Games and Economic Behavior*, 25, 34–78.

Cox, James C., Bruce Roberson, and Vernon L. Smith (1982), "Theory and behavior of single object auctions." In *Research in Experimental Economics* (Vernon L. Smith, ed.), 537–579, JAI Press, Greenwich.

Crawford, Vincent P. and Nagore Iriberri (2007), "Level-k auctions: can a nonequilibrium model of strategic thinking explain the winner's curse and overbidding in private-value auctions?" *Econometrica*, 75, 1721–1770.

Dittrich, Dennis A.V., Werner Gth, Martin G. Kocher, and Paul Pezanis-Christou (2012), "Loss aversion and learning to bid." *Economica*, 79, 226–257.

Engelbrecht-Wiggans, Richard and Elena Katok (2007), "Regret in auctions: Theory and evidence." *Economic Theory*, 81–101.

Engelbrecht-Wiggans, Richard and Elena Katok (2008), "Regret and feedback information in first-price sealed-bid auctions." *Management Science*, 54, 808–819.

Erev, Ido and Alvin E. Roth (1998), "Predicting how people play games: reinforcement learning in experimental games with unique, mixed strategy equilibria." *American Economic Review*, 88, 848–81.

Filardo, Andrew I. and Stephen F. Gordon (1998), "Business cycle durations." *Journal of Econometrics*, 85, 99–123.

Filiz-Ozbay, Emel and Erkut Y Ozbay (2007), "Auctions with anticipated regret: Theory and experiment." *American Economic Review*, 97, 1407–1418.

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

Geman, Stuart and Donald Geman (1987), "Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images." In *Readings in computer vision: issues, problems, principles, and paradigms* (Martin A. Fischler and Oscar Firschein, eds.), 564–584, Morgan Kaufmann Publishers Inc., San Francisco, CA, USA.

Geweke, John (1991), "Evaluating the accuracy of sampling-based approaches to the calculation of posterior moments." Staff Report 148, Federal Reserve Bank of Minneapolis.

Goeree, Jacob and Charles A. Holt (2002), "Quantal response equilibrium and overbidding in private-value auctions." *Journal of Economic Theory*, 104, 247–272.

Greiner, Ben (2004), "An online recruitment system for economic experiments." In *Forschung und wissenschaftliches Rechnen* (Kurt Kremer and Volker Macho, eds.), volume 63 of *Ges. fur Wiss. Datenverarbeitung*, 79–93, GWDG Bericht.

Grundl, Serafin and Yu Zhu (2022), "Robust inference in first-price auctions: Overbidding as an identifying restriction." *Journal of Econometrics*.

Guth, Werner, Radosveta Ivanova-Stenzel, Manfred Konigstein, and Martin Strobel (2003), "Learning to bid - an experimental study of bid function adjustments in auctions and fair division games." *Economic Journal*, 113, 477–494.

Holt, Charles A (1980), "Competitive bidding for contracts under alternative auction procedures." *The Journal of Political Economy*, 433–445.

Isaac, Mark, Svetlana Pevnitskaya, and Kurt S. Schnier (2012), "Individual behavior and bidding heterogeneity in sealed bid auctions where the number of bidders is unknown." *Economic Inquiry*, 50, 516–533.

James, Duncan and Derrick Reagle (2009), "Experience weighted attraction in the first price auction and becker degroot marschak." In *18th World IMACS Congress and MODSIM09 International Congress on Modelling and Simulation* (R.D. Braddock Anderssen, R.S. and L.T.H. Newham, eds.), 1437–1441, Modelling and Simulation Society of Australia and New Zealand and International Association for Mathematics and Computers in Simulation.

Kagel, John H (1995), "Auctions: a survey of experimental research." In *The Handbook of Experimental Economics* (Alvin E Roth and John H Kagel, eds.), first edition, chapter 7, 501 – 585, Princeton University Press, Princeton, NJ.

Kagel, John H and Dan Levin (1993), "Independent private value auctions: bidder behaviour in first-, second-, and third-price auctions with varying numbers of bidders." *The Economic Journal*, 868–879.

Kagel, John H and Dan Levin (2011), "Auctions: a survey of experimental research, 1995–2010." Manuscript, Ohio State University.

Karmeliuk, Maria, Martin G Kocher, and Georg Schmidt (2022), "Teams and individuals in standard auction formats: Decisions and emotions." *Experimental Economics*, 25, 1327–1348.

Kirchkamp, Oliver and J Philipp Reiß (2019), "Heterogeneous bids in auctions with rational and boundedly rational bidders: theory and experiment." *International Journal of Game Theory*, 48, 1001–1031.

Li, Xiaomin and Colin Camerer (2021), "Hidden markov modeling of the cognitive process in strategic thinking." *Available at SSRN 3838911.*

Maruotti, Antonello, Marco Fabbri, and Matteo Rizzolli (2022), "Multilevel hidden markov models for behavioral data: A hawk-and-dove experiment." *Multivariate Behavioral Research*, 57, 825–839.

McKelvey, Richard D and Thomas R Palfrey (1995), "Quantal response equilibria for normal form games." *Games and economic behavior*, 10, 6–38.

Neugebauer, Tibor and Reinhard Selten (2006), "Individual behavior of first-price auctions: The importance of information feedback in computerized experimental markets." *Games and Economic Behavior*, 54, 183–204.

Ratan, Anmol and Yuanji Wen (2016), "Does regret matter in first-price auctions?" *Economics Letters*, 143, 114–117.

Selten, Reinhard and Joachim Buchta (1998), "Experimental sealed bid first price auctions with directly observed bid functions." In *Games and Human Behavior: Essays in Honor of Amnon Rapoport* (David V. Budescu, Ido Erev, and Rami Zwick, eds.), chapter 5, 79–102, Lawrence Erlbaum Ass., Mahwah (NJ).

Shachat, Jason, J. Todd Swarthout, and Lijia Wei (2015), "A hidden markov model for the detection of pure and mixed strategy play in games." *Econometric Theory*, 31, 729752.

Shachat, Jason and Lijia Tan (2023), "How auctioneers set reserve prices in procurement auctions." *European Journal of Operational Research*, 304, 709–728.

Shachat, Jason and Lijia Wei (2012), "Procuring commodities: first-price sealed-bid or english auctions?" *Marketing Science*, 31, 317–333.

Shachat, Jason and Lijia Wei (2016), "Discrete rule learning and the bidding of the sexes." Technical Report SSRN 2289734, Social Science Research Network.

Vickrey, William (1961), "Counterpeculation, auctions, and competitive sealed tenders." *Journal of Finance*, 16, 8–37.

A English translations of the instructions in Mandarin

A.1 Saliva collection instructions

This is an experiment in the economics of decision making. For research purposes, we need to collect your saliva samples before and after experiment. We will use the experimental data and saliva samples marked by your ID number, so all of your personal information will be completely confidential.

In this study, all instruments are provided by School of Medicine at Xiamen University, which are safe, reliable, non-toxic and completely harmless to human body.

Please do not eat anything for the duration of the whole experiment. If you need to drink water, you can only drink the purified bottle water that we offer you after we have collected the first Saliva sample.

Before the experiment, we will follow the below steps to collect your saliva sample:

1. All participants line up and enter the laboratory one at a time.
2. Upon entering you must gargle three times using the provided cup of purified water. Do not drink any of the water. Then take the seat the monitor directs you to.
3. Read these instructions quietly at your desk. During this time do not eat and drink anything.
4. After 5 minutes, the monitor reads aloud the hormone test instructions. You receive an ID card and a collection tube (males receive a tube marked with black numbers, and females receive a tube marked with a red number.)
5. The monitor demonstrates the extraction process of saliva samples.
6. With the help of monitor, you extract your own saliva samples.
7. The monitor collects the tubes.

The steps of sampling saliva are as follows:

1. Open the package of the disposable syringe.
2. Carefully remove the protective cover together with the needle inside.
3. Put your tongue against the gums of your lower teeth, so the saliva will be pooled on top of your tongue. If there is too little saliva, please gently suck before the saliva collection.
4. Put the syringe into your month, and slowly pull the plunger to absorb saliva.
5. Inject your saliva into the tube and closed the tube. If the amount of saliva does not reach 0.75ml, marked with a line, please repeat above steps.

At the end of the experiment, you need to collect your saliva sample again.

Please do not talk with other subjects. If you have any questions, please raise your hand, and the monitor will answer your questions.

Thank you for your cooperation.

A.2 English translation of experiment instructions: Auction task

Overview

This is an experiment in the economics of decision making. The instructions are simple, and if you follow them carefully and make good decisions, you may earn a considerable amount of money that will be paid to privately you in cash at the end of this experiment.

It is important that you read these instructions carefully, so that you can understand the task in the experiment. We will also ask you a series of review questions after you finish reading these instructions. This is to test your understanding of the tasks in the experiment. You will be allowed to refer to the instructions as you answer the questions and as you participate in the experiment.

You are prohibited from talking to other participants. Turn off your mobile phones, and only use the computer to complete the experiment. If you have any questions, you should raise your hand and someone will come to answer them.

How you earn money

In todays experiment, you will participate in two kinds of auctions: forward and reverse auctions. In each period, you will be randomly matched with other participants into a group. Each group has three or six bidders. In each auction, you must submit a bid which will determine whether you earn profits in the period.

In a forward auction, the experimenter makes a single object of a fictitious good available for purchase within each group. In each period of forward auction, you will receive a random number between 20 and 40. It is equally likely that you receive any number in this range. Your random number is your private value of the object in this period. Only you know your value, all bidders values are selected from the range of 20 to 40 with the same probabilities.

After receiving your value, you will submit a bid. Your bid must be between 0 and 60 with the minimum interval of 0.01. If your bid is the highest in your group, you purchase the object and pay the amount of your bid. In the event of a tie, we will randomly select one of the highest bidders to purchase the object. If you purchase the object, then you will trade it for experimental currency equal to your private value. In this case:

$$\text{Your profit} = \text{Your value} \;\; \text{Your bid.}$$

(Note, your profit could be positive, zero, or negative.)

If you dont purchase the object successfully then your profit in the period is zero.

In the reverse auction, the experimenter purchases an object of the fictitious good from one participant in each group. In each period of reverse auction, you will receive a random number between 20 and 40. It is equally likely that you receive any number in this range. Your random number is your private cost to produce the object. Only you know your cost, all bidders costs are selected from the range of 20 to 40 with the same probabilities.

After receiving your cost, you will submit a bid. Your bid must be between 0 and 60 with the minimum interval of 0.01. If your bid is the lowest in your group, you sell the object and receive the amount of your bid. In the event of a tie, we will randomly select one of the lowest bidders to sell the object. If you sell the object, then you will trade it for experimental currency equal to your value. In this case:

Your profit= Your bid Your cost.

(Your profit could be positive, zero or negative)

If you dont sell the object successfully, you dont pay your cost and your profit is zero.

The profit of each period will be added up to determine your total profit. Your total earning are in experimental dollars, and we will convert them RMB at the exchange rate:

1 RMB = 10 $ Experimental.

We will pay you this money plus the 10 RMB show-up fee privately at the conclusion of the experiment.

How to use the computer interface

The example for Forward Auction:

The follow figure is the bidding window for a forward auction. You can see there are 3 subjects in your group in this period. And your value in this period is 36.12 points. You enter your bid in the blank box. Then click the Submit your bid button.

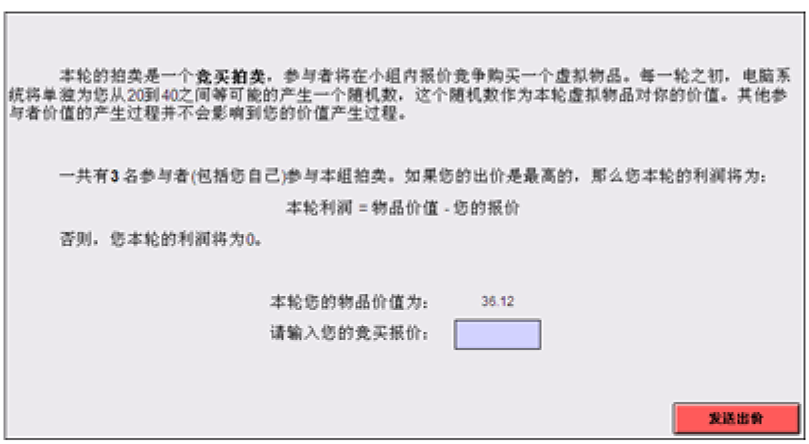

If you submit the bid of 30 in this period and you win the auction, you will earn the profit of 36.12 30.00 = 6.12.

The example for Reverse Auction:

The follow figure is the bidding window for the reverse auction. You can see there are 3 subjects in your group this period. Your cost in this period is 30.85 points. You enter your bid in the blank box. Then click the Submit your bid button.

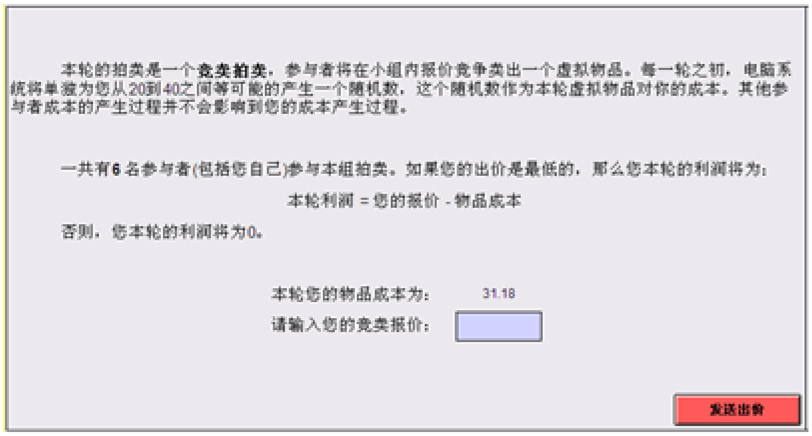

If you submit the bid of 31 in this period and you win the auction, you will earn the profit of 31.00 - 30.85 = 0.15.

Phases in the experiment

There are four different 25 period phases in todays experiment. They are forward auction with 6 bidders, forward auction with 3 bidders, reverse auction with 3 bidders and reverse auction with 6 bidders.

At the beginning of each phase, you will see a note about which phase is starting. The note will be shown as follow:

History of Past Periods

In each auction period, you can see your history information in the top of the window. Here you will find information about your past values (costs), you past bids, the past winning prices, and your past purchases and sales, etc.. The history of past periods will be shown as follow:

Review Questions

Please raise your hand if you have any questions. After 5 minutes we will go through the answers together.

1. If you are in a forward auction,

Suppose your value is 26 and you bid 28. If you buy the unit, your payoff =   .

If you dont buy the unit, your payoff =   .

Suppose your value is 30 and you bid 30. If you buy the unit, your payoff =   .

If you dont buy the unit, your payoff =   .

Suppose your value is 37 and you bid 35. If you buy the unit, your payoff =   .

If you dont buy the unit, your payoff =   .

2. If you are in a reverse auction,

Suppose your cost is 26 and you bid 28. If you sell the unit, your payoff =   .

If you dont sell the unit, your payoff =   .

Suppose your cost is 30 and you bid 30. If you sell the unit, your payoff =   .

If you dont sell the unit, your payoff =   .

Suppose your cost is 37 and you bid 35. If you sell the unit, your payoff =   .

If you dont sell the unit, your payoff =   .

A.3 Post-experiment survey (implemented on the computer)

We are interested in whether there is a correlation between participants bidding behavior and some socio-psychological factors. Your information will be very helpful for our research.

Female and Male will have different questions. All the information will be used only for research purpose. This information will be strictly confidential.

1. What is your ethnic origin?

(a) White (b) African (c)Hispanic (d) Non-Chinese Asian (e) Chinese Han (f)Chinese Minority

2. What is your age?

3. How many siblings do you have?

4. How many elder brothers, younger brothers, elder sisters and younger sisters do you have?

5. Would you describe your personality as (please choose one)?

(a) optimistic (b)pessimistic (c)neither

6. Which of the following emotions did you experience during the experiment (you may choose more than one)?

(a) anger (b) anxiety (c) confusion (d) contentment (e) fatigue (f) happiness (g) irritation (h) mood swings (i) withdrawal

7. What is your major field of study?

(a) Arts/Humanities (b) Economics (c) Other Social Science (d) Mathematics (e) Chemistry/Biology/Physics (f) Other Natural Science (g)Engineering (h)Sports

For female participants only:

8. How many days ago was the start of your last menstruation?

9. On average, how many days are there between your menstrual cycles?

10. How many days does your menstruation last on average?

11. Do you currently use a hormone-based contraceptive?

12. Do you currently experience any symptoms of PMS (Premenstrual Syndrome)?

B Deriving the strategic BR bidding rule

In the appendix, we present derivations for the BR rule in the case of a single kink defined by the condition $H + \kappa > b_{BR}(H) > (1 + \rho)H$ in the forward auction frame, and $L + \kappa < b_{BR}(L) < (1 + \rho)L$ in the reverse auction frame.

We first consider the forward auction frame. Suppose $b_{BR}(v_i)$ is a strictly increasing, continuous and bounded function, and thus has a continuous inverse. Now consider bidder j who follows the BR rule; he chooses the bid p_j to maximize his expected utility conditional upon his value v_j.[13]

Step one: Calculate the kink point v_R^*

By assumption the kink occurs at the inverse of b_{BR} evaluated at the maximum possible bid of the PM rule. The kink v_F^* satisfies,

$$b_{BR}(v_R^*) = H(1 + \rho).$$

Solving for v_F^*,

$$H(1 + \rho) = \frac{(L + \pi_{AM}\kappa)(1 + \rho)}{(1 + \rho(1 - \pi_{PM}))M} + \frac{M - 1}{M} v_F^*, \text{ or}$$

$$v_F^* = \frac{H(1 + \rho)M}{(M - 1)} - \frac{(L + \pi_{AM}\kappa)(1 + \rho)}{(1 + \rho(1 - \pi_{PM}))(M - 1)}$$

Step two: Calculate the bidding strategy

(1) The bidding strategy when $v_j \leq v_F^*$ is:

$$b_{BR}(v_j | v_j \leq v_F^*) = \frac{(L + \pi_{AM}\kappa)(1 + \rho)}{(1 + \rho(1 - \pi_{PM}))M} + \frac{M - 1}{M} v_j,$$

where $M = \eta(n - 1) + 1$.

(2) When $v_j > v_F^*$, the expected utility is:

[13] Note, we are considering the one-shot best response and accordingly suppress the index of auction iterations.

$$E[V(p_j|v_j)] = \eta(v_j - p_j)^{\frac{1}{\eta}} \cdot (p_j > p_j) = \eta(v_j - p_j)^{1/\eta} U(v_j)^{n-1}$$

The first order condition to p_j:

$$U(v_j)b'(v_j) + \frac{\eta(n-1)(1-\pi_{PM})}{H-L} b(v_j) = \frac{\eta(n-1)(1-\pi_{PM})}{H-L} v_j$$

$$U(v_j)^{M-1}b'(v_j) + \frac{(M-1)(1-\pi_{PM})}{H-L} U(v_j)^{M-2} b(v_j) = \frac{(M-1)(1-\pi_{PM})}{H-L} U(v_j)^{M-2} v_j$$

Since the bidding function is continuous when $v_j = v_F^*$

$$\frac{\mathrm{d}}{\mathrm{d}v_j}(U(v_j)^{M-1} b(v_j)) = \frac{(M-1)(1-\pi_{PM})}{H-L} U(v_j)^{M-2} v_j$$

$$U(v_j)^{M-1} b(v_j) = \int_L^{v_F^*} x \mathrm{d}U(x)^{M-1} + \int_{v_F^*}^{v_j} x \mathrm{d}U(x)^{M-1} + C$$

Since $b(v_F^*) = H(1+\rho)$, we have

$$C = U(v_F^*)^{M-1} H(1+\rho) - \int_L^{v_F^*} x \mathrm{d}U(x)^{M-1}$$

.

Thus, we solve the best response bidding strategy in forward auctions:

$$b(v_j) = \frac{G_F(v_F^*)[H(1+\rho) - v_F^*]}{G_F(v_j)} + v_j - \frac{\int_{v_F^*}^{v_j} G_F(x)\mathrm{d}x}{G_F(v_j)},$$

where

$$G_F(x) = \left(\pi_{PM} + (1-\pi_{PM})\frac{x-L}{H-L}\right)^{M-1}.$$

Next we solve for the BR in the reverse auction frame. By assumption the kink occurs at the inverse of b_{BR} evaluated at the minimum possible bid of the PM rule. In the reverse auction, assume there is only have one kink point since $L+\kappa < b_{BR}(L) < (1+\rho)L$

Step one: Calculate the kink point v_R^*

The kink v_R^* satisfies,

$$b_{BR}(v_R^*) = L(1+\rho),$$

$$L(1+\rho) = \frac{(H+\pi_{AM}\kappa)(1+\rho)}{(1+\rho(1-\pi_{PM}))M} + \frac{M-1}{M}v_R^*,$$

and

$$v_R^* = \frac{L(1+\rho)M}{(M-1)} - \frac{(H+\pi_{AM}\kappa)(1+\rho)}{(1+\rho(1-\pi_{PM}))(M-1)}.$$

Step two: Calculate the bidding strategy

(1) The bidding strategy when $v_j \geq v_R^*$ is:

$$b_{BR}(v_j|v_j \geq v_R^*) = \frac{(H+\pi_{AM}\kappa)(1+\rho)}{(1+\rho(1-\pi_{PM}))M} + \frac{M-1}{M}v_j,$$

where $M = \eta(n-1)+1$.

(2) When $v_j < v_R^*$, the expected utility is:

$$E[V(p_j|v_j)] = \eta(p_j - v_j)^{1/\eta} \cdot Prob(p_j < p_i) = \eta(p_j - v_j)^{1/\eta} U(v_j)^{n-1}$$

The first order condition with respect to p_j:

$$U(v_j)b'(v_j) - \frac{\eta(n-1)(1-\pi_{PM})}{H-L}b(v_j) = -\frac{\eta(n-1)(1-\pi_{PM})}{H-L}v_j$$

$$U(v_j)^{M-1}b'(v_j) - \frac{(M-1)(1-\pi_{PM})}{H-L}U(v_j)^{M-2}b(v_j) = -\frac{(M-1)(1-\pi_{PM})}{H-L}U(v_j)^{M-2}v_i$$

Since the bidding function is continuous when $v_j = v_R^*$

$$\frac{\mathrm{d}}{\mathrm{d}v_j}(U(v_j)^{M-1}b(v_j)) = -\frac{(M-1)(1-\pi_{PM})}{H-L}U(v_i)^{M-2}v_j$$

$$U(v_j)^{M-1}b(v_j) = \int_H^{v_R^*} x\mathrm{d}U(x)^{M-1} + \int_{v_*}^{v_j} x\mathrm{d}U(x)^{M-1} + C$$

Since $b(v_R^*) = L(1+\rho)$, we have

$$C = U(v_R^*)^{M-1}L(1+\rho) - \int_H^{v_R^*} x\mathrm{d}U(x)^{M-1}$$

.

Thus, we solve the best response bidding strategy in reverse auctions:

$$b(v_j) = \frac{G_R(v_R^*)[L(1+\rho) - v_R^*]}{G_R(v_j)} + v_j - \frac{\int_{v_*}^{v_j} G_R(x)\mathrm{d}x}{G_R(v_j)}$$

where

$$G_R(x) = \left(\pi_{PM} + (1-\pi_{PM})\frac{H-x}{H-L}\right)^{M-1}$$

C The MCMC estimation of the HMBM

This paper adopts a Bayesian methodology to make inferences regarding the parameters and unobserved components of the HMBM. The HMBM is a statistical process with three components. First, we can write the state space of bid functions as,

$$\begin{pmatrix} b_{AM}(v_{it}, f_{it}) \\ b_{PM}(v_{it}, f_{it}) \\ b_{BR}(v_{it}, f_{it}, n_{it}) \end{pmatrix} = \begin{pmatrix} \kappa_{f_{it}} & 1 & 0 \\ 0 & (1+\rho_{f_{it}}) & 0 \\ \alpha(f_{it}, n_{it}) & \beta_1(f_{it}, n_{it}, \rho_{f_{it}}) & \beta_2(f_{it}, n_{it}, \rho_{f_{it}}) \end{pmatrix} \begin{pmatrix} 1 \\ v_{it} \\ v_{it}^2 \cdot I(v_{it} > v_F^*) \end{pmatrix} + \begin{pmatrix} \epsilon_{AM_{it}} \\ \epsilon_{PM_{it}} \\ \epsilon_{BR_{it}} \end{pmatrix}.$$

For the best response rule we adopt a spline function approximation. This approximation is linear for valuations less than or equal to the kink at $v_F^* = (1+\rho)H$, and is quadratic for valuations exceeding the kink. We specify the approximation of the best response rule for the reverse auction frame in a similar fashion.

Let Φ denote the matrix of bidding rule parameters and $\Sigma = (\sigma_r^2)$ denote the vector of heteroscedastic variances of the bidding rule perturbations. Second, Γ indicates the matrix of the parameters of the transition matrix P, where each row of Γ consists of the multinomial probit model parameters of the corresponding bidding rule. Third, Π_1 is the multinomial distribution governing the initial assignment of bidding rules and has the parameters π_{AM} and π_{PM}. The output of the HMBM consists of the observable the sequences of bids and auctions outcomes, B and O respectively, for each bidder, and the unobservable sequence of bidding rules, S, adopted by each bidder. Also, for notational convenience, let V be the collection of all V_i's.

Consider the joint posterior density function of the HMBM parameters and the unobserved realized state sequences, $h(S, \Phi, \Sigma, \Gamma, \Pi_1|B, V, O)$. We first assume that parameters of the bidding rules are independent from the auction outcomes. Then we express this joint density as the product of the marginal density of HMBM parameters conditional on the observed bids, values, and outcomes; and the unobserved states with the marginal density of the states conditional upon action choices and outcomes.

$$h(S, \Phi, \Sigma, \Gamma, \Pi_1|B, V, O) = h(\Phi, \Sigma, \Gamma, \Pi_1|S, B, V, O)h(S|B, V, O).$$

Next, we assume the distribution of the bidding rule parameters are independent of the distribution of the rule transition probability parameters when both are condition upon the observable elements (B, V, O). This allows us to state

$$h(S, \Phi, \Sigma, \Gamma, \Pi_1|B, V, O) = h(\Phi, \Sigma|S, B, V)h(\Gamma, \Pi_1|S, O)h(S|B, V, O). \tag{2}$$

This product of three conditional posteriors permits a simple Markov Chain procedure of sequentially sampling from these distributions. The MCMC approach relies on augmenting the parameter space with filtered values of the unobserved states S, and using Gibbs sampling procedures generate sequential draws from the marginal distributions of equation 2. After a large number of iterations, indexed by l, the empirical density of these draws converges in probability to the true posterior density functions (Geman and Geman, 1987). The first iteration of MCMC algorithm, $l = 0$, starts with the construction of S^0 by random draws from the uniform multinomial distribution, Π_1^0 from the uniform Dirichlet distribution, and draws for all values parameter values Φ^0, Σ^0, and Γ^0 from the standard uniform distribution.

Iterations $l > 0$, consists of the following three steps.

Step 1: Sample $\Phi^{(l)}$ and $\Sigma^{(l)}$ by using $S^{(l-1)}$, B, V;
Step 2: Sample $S^{(1)}$ by using $\Gamma^{(l-1)}$, O, and $S^{(l-1)}$; and
Step 3: Sample $\Gamma^{(l)}$ by using $S^{(l)}$, $\Psi^{(l)}$ and O.

The Gibbs sampler is run for a large number of iterations until the empirical distribution of all the parameters has converged according to the convergence test of Geweke (1991). Then the sampling procedure is allowed to continue to run for another number of iterations to build up an empirical distribution that corresponds to the posterior distribution of the HMBM parameters. It is from this empirical distribution that we conduct statistical inferences. We now describe the three steps of an iteration of the Gibbs sampler in detail.

Step 1: Sampling the rule parameters $\{\mathbf{\Phi},\ \mathbf{\Sigma}\}^{(l)}$

Given the values of S, We can summarize the rules summarized as linear models:

$$b_s(v_{it}) = \phi_{s0} + \phi_{s1}v_{it} + \phi_{s2}v_{it}^2 + \epsilon_{sit}.$$

Define B_s to be the vector of bids when subjects adopt rule s, and V_s to be the matrix of right-hand side variables when subjects adopt rule s. We start by assuming the prior joint distribution of $(\phi_s^{(l)}, \sigma_s^{(l)})$ follows the Normal-Inverse Gamma distribution, N-IG$(\tilde{\phi}_s, \tilde{A}_s, \tilde{\sigma^2}_s)$. Note the prior parameters are generally set to zero, except for the following point prior assignments that enforce parameter restrictions: $\phi_{PM0} = 0$, $\phi_{AM1} = 1$, and $\phi_{AM2} = \phi_{PM2} = 0$.

The posterior distribution of $(\phi_s,\ \sigma_s^2)$ has the Normal-Inverse Gamma form:

$$\phi_s \sim \mathbf{N}\left(\bar{\phi}_s, \bar{A}_s \cdot \sigma_s^2\right)$$
$$\sigma_s^2 \sim \mathbf{IG}\left(\frac{\bar{\nu}_s}{2}, \frac{\bar{\nu}_s \cdot \bar{\sigma}_s^2}{2}\right)$$

where $\bar{\nu}_s$ is the degrees of freedom of the linear model.

We draw values for $(\sigma_s^2)^{(l)}$ from the posterior inverse-gamma distribution with $\bar{\nu}_s$ and $\bar{\sigma}_s^2$, which can be calculate by following formulas.

$$
\begin{aligned}
\bar{A}_s &= \left(\tilde{A}_s + V_s' \cdot V_s\right) \\
\bar{\phi}_s &= \bar{A}_s^{-1} \cdot \left(\tilde{A}_s \cdot \tilde{\phi}_s + V_s' \cdot B_s\right) \\
\bar{\sigma}_s^2 &= \bar{\nu}_s^{-1} \left(\bar{\nu}_s \tilde{\sigma}_s^2 + \left(B_s - \bar{\phi}_s V_s\right)' \left(B_s - \bar{\phi}_s V_s\right) + \left(\bar{\phi}_s - \tilde{\phi}_s\right)' \tilde{A}_s \left(\bar{\phi}_s - \tilde{\phi}_s\right)\right)
\end{aligned}
$$

Where $\tilde{\phi}_s$, $\tilde{A}_s$ and $\tilde{\sigma}_s^2$ are the prior parameters. Value of $\phi_s^{(l)}$ can be drawn from the normal distribution with the variance $(\sigma_s^2)^{(l)}$. The prior parameters $\tilde{\phi}_s$, $\tilde{A}_s$ and $\tilde{\sigma^2}_s$ are chosen to be non informative. To avoid the switching of estimate rules, we restrict $\sigma_{PM}^2 > \sigma_{BR}^2 > \sigma_{AM}^2$. The restriction is consistent with the E-MLE results of Shachat and Wei (2012).

Step 2: Sampling the state sequences $\mathbf{S}^{(l)}$

To generate $s_{it}^{(l)}$, we start at $t = 100$ and recursively calculate state probabilities. Then we determine state $s_{it}^{(l)}$ by taking a realization from the standard uniform distribution and comparing it to the calculated state probabilities. The formula for the state probabilities are

$$
\Pr(s_{it}^{(l)} = r) = \frac{\Theta(r)}{\sum_{j \in \{AM,PM,BR\}} \Theta(j)},
$$

where

$$
\Theta(r) = \begin{cases}
\Pr\left(b_{it}|s_{it}^{(l)} = r, v_{it}, \phi_r^{(l)}, \sigma_r^{(l)}\right) \Pr\left(s_{it}^{(l)} = r|s_{it-1}^{(l-1)}, o_{it-1}, \Gamma^{(l-1)}\right) & \text{if } t = 100 \\
\\
\Pr\left(b_{it}|s_{it}^{(l)} = r, v_{it}, \phi_r^{(l)}, \sigma_r^{(l)}\right) \Pr\left(s_{it}^{(l)} = r|s_{it-1}^{(l-1)}, o_{it-1}, \Gamma^{(l-1)}\right) & \\
\quad \times \Pr\left(s_{it+1}^{(l)}|s_{it}^{(l)} = r, o_{it}, \Gamma^{(l-1)}\right) & \text{if } 2 < t < 100 \\
\\
\Pr\left(b_{it}|s_{it}^{(l)} = r, v_{it}, \phi_r^{(l)}, \sigma_r^{(l)}\right) \Pr\left(s_{it}^{(l)} = r|\Pi_1^{(l-1)}\right) & \\
\quad \times \Pr\left(s_{it+1}^{(l)}|s_{it}^{(l)} = r, o_{it}, \Gamma^{(l-1)}\right) & \text{if } t = 1
\end{cases}.
$$

Step 3: Sampling $\Pi_1^{(l)}$ and $\Gamma^{(l)}$

For Π_1 we assume a uniform Dirichlet prior, $D(\Pi_{1,s}; d_{AM}, d_{PM}, d_{BR})$ by setting the shape parameters, d_s to one. The Dirichlet distribution is the conjugate prior for the multinomial distribution, and the shape parameters of the conditional posterior are simply the number of occurrences of each state in the first element of the sequences $S^{(l)}$. Let n_j^0 be the number incidences of $s_{i,1}^{(l)} = j$. Thus, the conditional posterior is

$$h(\Pi_{1,s}|S^{(l)}) = h(\Pi_{1,s}; d_{AM} + n_{AM}^0, d_{PM} + n_{PM}^0, d_{BR} + n_{BR}^0).$$

We generate $\Pi_1^{(l)}$ by sampling from this distribution.

To characterize the marginal conditional posterior distribution for the rule transition probability matrix parameters Γ and to generate a sample $\Gamma^{(l)}$, we use the methods introduced by Albert and Chib (1993) and Filardo and Gordon (1998). First, reduce the three indices of equation 1 to two by normalizing the BR rule as follows[14];

$$\Psi_{it}^{j1} = \Psi_{it}^{jAM} - \Psi_{it}^{jBR} \text{ and } \Psi_{it}^{j2} = \Psi_{it}^{jPM} - \Psi_{it}^{jBR}.$$

Now we can restate the transition probabilities; Given S, Γ and the following inequality constraint:

$$\begin{aligned}
\Pr(s_{it} = BR|s_{it-1}, o_{it-1}) &= \Pr(\Psi_{it-1}^{s_{it-1}1} \leq 0, \Psi_{it-1}^{s_{it-1}2} \leq 0) \\
\Pr(s_{it} = PM|s_{it-1}, o_{it-1}) &= \Pr(\Psi_{it-1}^{s_{it-1}1} \geq \Psi_{it-1}^{s_{it-1}2}, \Psi_{it-1}^{s_{it-1}1} \geq 0). \\
\Pr(s_{it} = AM|s_{it-1}, o_{it-1}) &= \Pr(\Psi_{it-1}^{s_{it-1}2} \geq \Psi_{it-1}^{s_{it-1}1}, \Psi_{it-1}^{s_{it-1}2} \geq 0)
\end{aligned}$$

We construct realized values of the Ψ_{it}^{j1} and Ψ_{it}^{j2} by using $\Gamma^{(l-1)}$, $s_{it}^{(l)}$, and o_{it} with perturbations randomly generated from the appropriate truncated bivariate normal distributions. These realized normalized indices are now simple linear regression models with unit variance. With the known variance and assumed normal distributed prior with zero mean, the conditional posterior distribution is also normal and we can generate draws similarly to how we do in Step 1.

D The estimation results of four bidding rules

In response to our three bidding rule model not generating a distribution of rules strongly independent of value, we consider a four bidding rule model by adding the rule BA. The BA rule follows the AM rule for low quality values and the BR rule for high quality values. For forward auctions the switching value is thirty-five and twenty-five for reverse auctions.

[14]Note that this normalization highlights that γ_{jk} is not identified, only the differences between rules indices under neutral reinforcement are.

Formally stated for the Forward and Reverse auction frames,

$$b_{BA}(v_{it}|f_{it}=F)=\begin{cases}\kappa_F+v_{it}+\epsilon_{AM_{it}} & \text{if } v_{it}\leq 35,\\ \dfrac{(L+\pi_{AM}\kappa)(1+\rho)}{\big[1+\rho(1-\pi_{PM})\big]M}+\dfrac{M-1}{M}v_{it}+\epsilon_{BR_{it}} & \text{if } \ 35\leq v_{it}\leq v^*_{Fn}\\ \dfrac{G_F(v^*)[H(1+\rho)-v^*]}{G_F(v_i)}+v_{it}-\dfrac{\int_{v^*}^{v_i}G_F(x)\mathrm{d}x}{G_F(v_i)}+\epsilon_{BR_{it}} & \text{if } \ v_{it}>v^*_{fn}\end{cases},$$

and for the Reverse auction frame,

$$b_{BA}(v_{it}|f_{it}=R)=\begin{cases}\kappa_R+v_{it}+\epsilon_{AM_{it}} & \text{if } v_{it}\geq 25,\\ \dfrac{(H+\pi_{AM}\cdot\kappa)(1+\rho)}{(\big[1+\rho(1-\pi_{PM})\big]M}+\dfrac{M-1}{M}v_{it}+\epsilon_{BR_{it}} & \text{if } v^*_{Rn}\leq v_{it}\leq 25\\ \dfrac{G_R(v^*)[L(1+\rho)-v^*]}{G_R(v_i)}+v_{it}-\dfrac{\int_{v_*}^{v_i}G_R(x)\mathrm{d}x}{G_R(v_i)}+\epsilon_{BR_{it}} & \text{if } v_{it}<v^*_{Rn}\end{cases}.$$

We next report on the parameter estimates of the alternative bidding rules. Table **??** reports on the new estimated coefficients for the AM and PM rules, and inspection reveals slightly more aggressive mark-up for both rules in the two frames compared to the three rule HMBM. **??** reports the revised estimates for the parameters of the BR rule. All of the parameter estimates are inline with those of the three rule HMBM estimates, with the notable exception of an increased estimated value of the error variance term σ^2_{BR}.

Table 6: Mark-up bidding rule parameters: posterior means with 95% confidence intervals

Rule	Variable	Coef.	95% Confidence interval
	$D_F\cdot\kappa$	-0.59	[-0.60,-0.57]
AM	$D_R\cdot\kappa$	0.49	[0.48,0.51]
	σ^2_{AM}	0.16	[0.16,0.17]
	$D_F\cdot\rho$	-0.22	[-0.24,-0.20]
PM	$D_R\cdot\rho$	0.23	[0.21,0.25]
	σ^2_{PM}	55.22	[50.66,60.37]

Next we report on the evolution of the rule assignment distribution over the course of the 100 periods of the experiment in Figure 4. In this plot of time series we can see that adoption of

Table 7: Strategic Best Response (BR) bidding rule parameters: posterior means with 95% confidence intervals

	Forward Auction		Reverse Auction	
Variable	n=3	n=6	n=3	n=6
Intercept	1.83	0.25	7.15	4.78
	[1.44,2.18]	[-0.20,0.70]	[6.93,7.37]	[4.54,5.03]
v_i	0.88	0.94	0.83	0.89
	[0.87,0.89]	[0.92,0.96]	[0.82,0.83]	[0.89,0.90]
$(v_i - v^*) \cdot I(v > v^*)$	-0.06	0.16	-	-
	[-0.16,0.04]	[0.06,0.26]	-	-
$(v_i - v^*)^2 \cdot I(v > v^*)$	-0.02	-0.05	-	-
	[-0.04, -0.01]	[-0.07,-0.03]	-	-
v^*	33.32	32.95	21.11	22.17
	[32.32,34.33]	[32.01,33.91]	[20.60,21.61]	[21.69,22.66]
σ^2_{BR}	0.96			
	[0.93,1.01]			

the BA rule starts below ten percent and then rises close to forty percent. When we consider the independent distribution of values/costs, BA bidders follow the AM rule seventy-five percent of the time and the BR rule twenty-five percent of the time; the net result is that on average in any one auction there are equal numbers of AM and BR generated bids. This suggest more mark-up bidding than in this and other previous studies.

We now provide some evaluation of whether adding the BA rule improves the issue of independence of rule and type in the estimated posterior rule assignments. Figure **??** presents bar plots of the rule assignments by 1 ECU buckets of values and cost. These charts appear to have reduced some of the strong trends that one finds in Figure 1.

Finally, we assess whether the four rule version of the HMBM yields more in-sample predictive power of bids than the three rule version. We assess this by two measures: Root Mean

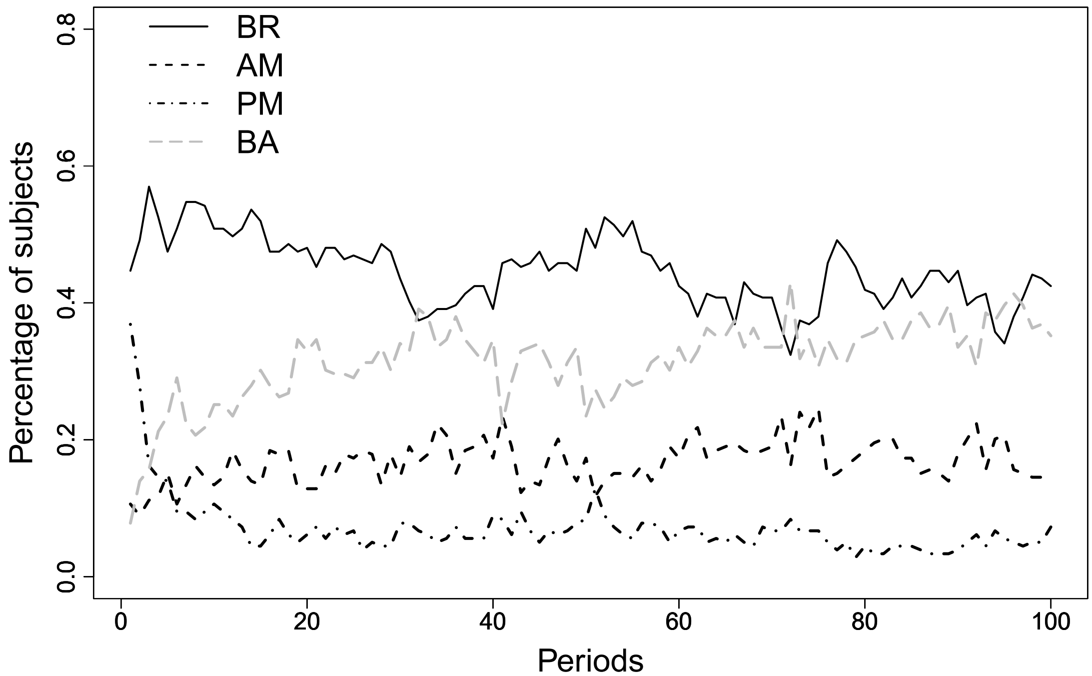

Figure 4: Estimated sequences of subject proportions using the AM, PM, BA and BR rules

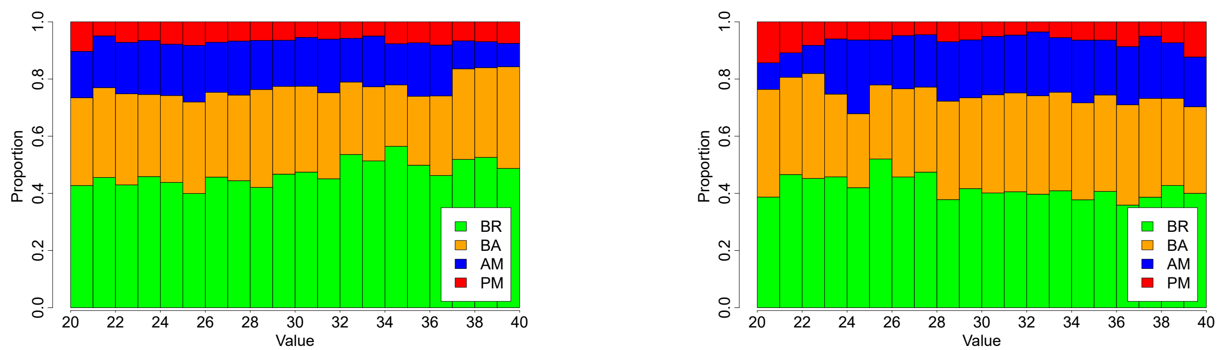

(a) Forward auctions

(b) Reverse auctions

Figure 5: Posterior distribution of bidding rules by value/cost in the four rule HMBM

Square Error (RMSE) and Mean Absolute Error (MAE). These are calculate as follows,

$$\text{RMSE} = \sqrt{\frac{\Sigma_m \Sigma_I (b_{it} - \hat{b_{it}})^2}{m \cdot I}}$$

$$\text{MAE} = \frac{\Sigma_m \Sigma_I \left| b_{it} - \hat{b_{it}} \right|}{m \cdot I},$$

where m is number of observations and I is the number of sampling iterations. The results of the three bidding rule HMBM are RMSE = 2.12 and MAE=0.91. In the four bidding rule model we so infinitesimal improvement; RMSE=2.09 and MAE=0.90.