18th World IMACS / MODSIM Congress, Cairns, Australia 13-17 July 2009
http://mssanz.org.au/modsim09

Experience Weighted Attraction in the First Price Auction and Becker DeGroot Marschak

Duncan James[1] **and Derrick Reagle**[1]

[1]*Economics Department, Fordham University, NY, USA*
Email: dujames@fordham.edu

Abstract: In this paper we explore the performance of Experience Weighted Attraction (EWA) in two different auction institutions: First Price Sealed Bid (FPSB), and Becker DeGroot Marschak (BDM). Our results suggest that learning has some promise as a possible explanation for previously documented cross-institutional choice anomalies usually attributed to risk aversion. Specifically, a range of parameters within EWA which provide for a learning algorithm that is "close" to reinforcement learning will then generate overbidding within the FPSB and within BDM. Such overbidding, if observed with human subjects, might be attributed by an outside observer to risk aversion (in the case of FPSB) or risk seeking (in the case of BDM); clearly such a risk-preference-based explanation is inconsistent across institutions. A learning-algorithm-based alternative explanation, such as might be offered by EWA, would at least be consistent across institutions; that is to say, a single EWA parameter set can generate the required stylized facts across the different bidding institutions in question.

Keywords: *Auctions, Risk Aversion, Learning, Experience Weighted Attraction*

1. INTRODUCTION

Do time series data from auctions tell us about risk aversion – or something else? Bidders' risk preference parameters play a key role in the Constant Relative Risk Aversion Model (CRRAM) of bidding in the first price sealed bid auction (FPSB). For that matter, inferring bidders' risk preference parameters is the whole point to the Becker-DeGroot-Marschak procedure (BDM), which can be thought of as a version of the second price sealed bid auction. But when researchers implement these institutions in laboratory experiments, the results from the FPSB and BDM contradict one another: results from the FPSB interpreted via CRRAM, imply that on average subjects are risk averse; results from BDM imply that on average subjects are risk seeking (Isaac and James (2000), Berg, Dickhaut, and McCabe (2005)). This raises the question of whether there is something instead of or in addition to risk preferences that is influencing these auction data time series.[1]

This paper seeks to examine the possibility that a learning model might explain some of the regularities – which are also anomalies – in this area. We choose to use Experience Weighted Attraction (EWA) learning as a possible explanation, and we do so because of its ability to nest a number of other learning models. This means that we can start with a relatively general model, and then investigate which specific parameter combinations within the general model allow for possible matching of experimental regularities or "stylized facts".

2. BACKGROUND

The first model of bidding behavior in the FPSB, due to Vickrey, assumes that each bidder's value is independently and privately drawn from a uniform distribution whose support is common knowledge. Vickrey then derives that a utility-maximizing-risk-neutral bidder's bid, b_i, should be related to her own (private) value, v_i, for the object at auction according to:

$$b_i = [(n - 1)/n]v_i \qquad (1)$$

where n is the total number of bidders participating in the auction. In order to explain systematic violations of the Vickrey model in experimental data, Cox, Roberson, and Smith (followed by Cox, Smith, and Walker (1982)) introduced non-risk-neutral preferences into the derivation of the FPSB bid function giving:

$$b(v_i,r_i) = v_{lower} + [(N - 1)/(N - 1 + r_i)] (v_i - v_{lower}) \qquad (2)$$

where:

v_{lower} = the lower bound on the support, $[v_{lower}, v_{upper}]$, of a uniform distribution from which object values are drawn

v_i = bidder i's value of the object being auctioned

N = the number of bidders in the auction

r_i = the exponent in bidder i's Constant Relative Risk Aversion utility function.

In contrast, BDM was developed at the outset as a way to measure risk preferences. Risk preferences are not added to the BDM as a way to better explain empirical data, but instead are assumed to be representable by expected utility, and recoverable as a consequence of optimal bidding behavior in the second price auction.

Specifically, an individual whose risk preferences have the expected utility property will value a two-state lottery as:

$$\text{Certainty Equivalent} = E[U(w)] = p_{high}[U(w_{high})] + p_{low}[U(w_{low})] \qquad (3)$$

[1] Other attempts at addressing deviations from risk neutral Nash Equilibrium (RNNE) predictions have been made by other researchers. Cox, Roberson, and Smith (1982) suggest the possibility that the utility of the act of winning an auction – separate and additional to the expected utility of the different possible end-state monetary payoffs – could induce bidding in excess of the RNNE prediction (which CRRAM otherwise interprets as "risk aversion"). Engelbrecht-Wiggins (1989) models the possibility that the opportunity cost represented by foregone payoffs associated with strategies not chosen (dubbed "regret") might induce bidding in excess of the RNNE prediction. Andreoni and Miller (1995) examine the behavior of genetic algorithms in a number of institutional settings, and find some "overbidding".

Also, an individual has a dominant strategy to bid their value in a second price auction, so their bid function is simply:

$$b_i = v_i = E[U(w)] \qquad (4)$$

As such, were an expected utility maximizer to bid for a two-state lottery in a second price auction, we should be able to observe her certainty equivalent over the probabilities and payoffs associated with that lottery. In turn, an economist could then use such certainty equivalents to draw inferences about the utility functions of such expected utility maximizers.

3. EWA AND BIDDING

In EWA, an agent makes decisions according to a difference equation system that includes variables that try to accommodate various cognitive processes. Most notable of these variables are δ, which parameterizes the ex post allure of payoffs to strategies other than the strategy chosen by the agent, and ϕ, which parameterizes the extent to which the agent remembers past play. Higher (lower) values of δ imply a greater (lesser) ex post "awareness" of the payoffs associated with other strategies than the strategy actually chosen by the agent. Higher (lower) values of ϕ are associated with longer (shorter) spans of memory of past "stimuli" (which could be either payoffs to actually chosen strategies or "imagined" foregone payoffs associated with strategies not taken, or both). The system of equations is then:

$$A^j_i(a,t) = \phi N(t-1)A^j_i(a,t-1)/N(t) + [\delta+(1-\delta)I(s^j_i,s_i(t))]\pi_i(s^j_i,s_{-i}(t))/N(t) \qquad (5)$$

$$N(t) = (1-\kappa)\phi N(t-1) + 1 \quad \text{for } t \geq 1 \qquad (6)$$

$$P^j_i(a, t+1) = e^{f(A(a,t))}/\Sigma e^{f(A(a,t))} \qquad (7)$$

where: A^j : the "attraction" of a particular strategy j

N : a measure of experience

P^j_i = the probability that a particular strategy j is chosen by player i

I = an indicator function; I = 1 if s^j_i=$s_i(t)$, 0 otherwise

$\pi_i(s^j_i,s_{-i}(t))$ = payoff associated with strategy profile $(s^j_i,s_{-i}(t))$

κ : affects the accumulation of attractions

$f(\cdot)$ = a monotonically increasing function scaling the attractions

Once the system of equations is initialized, an EWA agent makes choices in the following manner. First, the agent calculates the attraction, $A^j(t)$, for each strategy j, based on past experience with actual and or foregone payoffs associated with the various strategy profiles. Second, the agent forms a probability measure over the possible strategies by employing the scaling of attractions implied by the equation for P^j_i. Finally, the agent draws a strategy choice from the probability measure so constructed.

One can alter $f(\cdot)$ to give more (or less) weight to strategies with relatively high attractions. For instance, if $f(A(a,t)) = \lambda(A(a,t))$, then as λ goes to infinity the strategy having the maximum attraction is chosen with probability one.[2] We vary the value of λ among 1, 10, and infinity in order to document EWA's performance across varying levels of randomization over strategies.[3]

[2] Proof available upon request.

[3] One should note that the classic learning models which EWA nests as particular parameterizations have long-run convergence properties which have been characterized analytically. Two notable recent contributions in this area are those of Beggs (2002) and Hopkins (2002). Among their many respective results, the following are of direct interest for our paper. First, Beggs shows that in individual choice problems, reinforcement learning (e.g. δ=0) converges towards expected value maximization. Second, Hopkins shows that versions of reinforcement learning and fictitious play (e.g. δ=1) can both be characterized as transformations of replicator dynamics, and as such that both will converge to similar strategies, only at different speeds. (Fictitious play converges faster than reinforcement

3.1 Implementation of EWA

Having outlined the auction theoretic and learning models to be used in this study, we now wish to outline their implementation by computer. We implement FPSB by drawing a value for the EWA agent's opponent from a [$0.00, $2.00] uniform distribution, multiplying this by 1/N (which for two bidders, including the EWA agent, is ½), and comparing this number to the EWA agent's bid. If the EWA agent's bid is higher, it receives its object value minus its bid; otherwise, it receives zero. The EWA agent's object value is varied across simulations, but not within, between $0.50, $1.00, and $1.50. For the n=2 case this implies risk neutral optimal bids of $0.25, $0.50, and $0.75, respectively, which are comparable to the risk neutral optimum in each case for BDM.

We implement BDM by first making a draw of the opponent's bid from a uniform distribution on the support [$0, $2.00], and then comparing this value to the EWA agent's bid. If the EWA agent's bid is higher than that of the opponent, the EWA agent receives the results of a draw from a lottery having a 50% probability of the high state payoff (varied in cross-section, but not in time series, between $0.50, $1.00, and $1.50) and a 50% probability of the low state payoff, $0.00 (constant). If the EWA agent's bid is lower than that of the opponent, the EWA agent receives an amount equal to the opponent's bid. Thus for a single round of BDM, there are two random draws to be made.[4]

4. EWA BIDDING IN AUCTIONS AND ITS COMPARISON WITH HUMAN BIDDING

We document the "short run" behavior of EWA by comparing cross-sectional mean EWA bids in each auction institution at particular points in time with canonical experimental results, or stylized facts, in each auction institution.[5] We concentrate on whether there is a particular set of parameters for which EWA can reproduce the following broad patterns in human bidding data: sustained overbidding in the FPSB, and decreasing overbidding in BDM.

In the FPSB, we find that low δ is critical to overbidding. Overall, in terms of we find that as λ increases from 1 to 10, the region supporting overbidding *shrinks* to (ϕ=0,δ=0), (ϕ=0.5,δ=0), (ϕ=0,δ=0.5) (notably *not* including (ϕ=0.5,δ=0.5)). Then, as λ increases from 10 to ∞, the region *expands* to reclaim (ϕ=0.5,δ=0.5) and add (ϕ=1.0,δ=0). We also note that as λ approaches infinity, a dependence on initial values develops, as the learning model seizes on early positive payoffs due to overbidding and makes this the basis of its subsequent bidding. One can see this in the difference in mean bids across different starting values, for given values of ϕ and δ, that is apparent for $\lambda=\infty$, but otherwise generally absent.

In BDM we note that as λ increases, the importance of the initial bid in determining the path of subsequent bidding increases in what appears to be a monotonic manner. The effect of high λ is such that either extreme underbidding or extreme overbidding could be observed for the same parameterization of ϕ and δ, depending upon whether the initial bid is low or high, respectively. Furthermore, for high values of λ, bidding behavior does not converge towards the risk neutral optimum, but rather sticks near the initial bid. This suggests that λ=1 as a defensible parameterization. Similar to the case of FPSB, we observe overbidding between Periods 2 and 50 for parameterizations featuring either $\phi<1.0$ or $\delta<0.5$, or both. However, given that an additional "stylized fact" for BDM is a movement towards the risk neutral optimum between Period 10 and Period 50 (James (2007)), the parameterization (ϕ=1,δ=0) stands

---

learning.) Hence we would expect that in the individual choice problems posed in this paper that EWA would eventually tend toward expected value maximization. However, EWA's choices in the short run (say, over the length of a typical human subjects experiment) and in particular choice settings (auctions) may exhibit transitory asymmetries which might potentially reproduce broad patterns in human bidding. It is this possibility that is of primary interest in this paper.

[4] Finally, we note that the processes just outlined need to be initialized. We vary the method by which we do this, to check for robustness of results. The key items in terms of initialization are initial attractions to strategies (that is, bids), and the initial bids. The initial attractions over strategies are set as either a) uniform (akin to uninformed priors) or b) proportional to expected values. The initial bid was varied on either side of the risk neutral optimum bid.

[5] Tables of computational results available upon request.

out as the best match across institutions. (ϕ=1,δ=0) allows initial overbidding in BDM that reduces through time, and also allows sustained overbidding in the FPSB through the 50 period window that approximates the length of a human subjects experiment.

5. CONCLUSIONS

It is possible to recreate, at least in broad outline, some key stylized facts about human bidding behavior by means of learning algorithms. This suggests that the role of learning in bidding tasks is likely something which experimental economists and bidding theorists should not ignore. At a minimum, failing to account for the possibility that asymmetry in bidding might be caused by the adjustment dynamics of a learning algorithm might lead researchers (who are seeking to explain bidding behavior) to attempt to explain too much of the asymmetry in bidding by recourse to too limited a set of parameters (for example, curvature of the utility of wealth function).

REFERENCES

Andreoni, James and John H. Miller (1995), Auctions with Artificial Adaptive Agents. *Games and Economic Behavior*, 10, 39-64.

Becker, Gordon M., Morris H. DeGroot, and Jacob Marschak (1964), Measuring Utility by a Single-Response Sequential Method. *Behavioral Science*, 9, 226-232.

Beggs, Alan (2002), On the Convergence of Reinforcement Learning. *Oxford University Discussion Paper*, no. 96.

Berg, Joyce, John Dickhaut, and Kevin McCabe. (2005), Risk Preference Instability Across Institutions: A Dilemma."*Proceedings of the National Academy of Sciences*, 102, 4209-4214.

Camerer, Colin and T. H. Ho (1999), Experience-Weighted Attraction Learning in Normal Form Games. *Econometrica*, 67, 827-874.

Cox, James C., Bruce Roberson, and Vernon L. Smith (1982), Theory and Behavior of Single Object Auctions. In: Vernon L. Smith, ed., *Research in Experimental Economics*, Vol. 2, Greenwich, CT: JAI Press.

Engelbrecht-Wiggins, John (1989), The Effect of Regret on Optimal Bidding in Auctions. *Management Science*, 685-692.

Hopkins, Ed (2002), Two Competing Models of How People Learn in Games. *Econometrica*, 70, 2141-2166.

Isaac, R. Mark and Duncan James (2000), Just Who Are You Calling Risk Averse? *Journal of Risk and Uncertainty*, 20, 177-187.

James, Duncan (2007), Stability of Risk Preference Parameter Estimates within the Becker-DeGroot-Marschak Procedure, *Experimental Economics*, 10(2), pp.123-141.

Vickrey, William (1961), Counterspeculation, Auctions, and Competitive Sealed Tenders. *Journal of Finance*, 16, 8-37.