

An Experimental Evaluation of Regret-Based Econometrics

Noam Nisan
The Hebrew University of Jerusalem
and Microsoft Research
noam@cs.huji.ac.il

Gali Noti
The Hebrew University of Jerusalem
and Microsoft Research
gali.noti@cs.huji.ac.il

ABSTRACT

Using data obtained in a controlled ad-auction experiment that we ran, we evaluate the regret-based approach to econometrics that was recently suggested by Nekipelov, Syrgkanis, and Tardos (EC 2015). We found that despite the weak regret-based assumptions, the results were (at least) as accurate as those obtained using classic equilibrium-based assumptions. En route we studied to what extent humans actually minimize regret in our ad auction, and found a significant difference between the “high types” (players with a high valuation) who indeed rationally minimized regret and the “low types” who significantly overbid. We suggest that correcting for these biases and adjusting the regret-based econometric method may improve the accuracy of estimated values.

Keywords

Econometrics; Behavioral Economics; Cognitive Biases; Sponsored Search Auctions; Regret

1. INTRODUCTION

The field of econometrics combines observational data with specific modeling assumptions in order to estimate parameters of interest. It goes beyond mere statistical analysis in that it assumes specific models of how the parameters of interest relate to the type of data observed. On the one hand, these assumptions provide power to econometric analysis by allowing it to estimate parameters that do not directly appear in the data, but, on the other hand, the correctness of the whole estimation strongly depends on the correctness of the utilized model, i.e., on the extent to which the situation at hand indeed conforms to the theoretical model.

In game-like scenarios, typical models assume that players are at an equilibrium. This is of course not a trivial assumption as we know that humans are not fully rational and certainly do not always “find” an equilibrium point (see, e.g., [8, 7] and the references therein). Beyond the usual

human lack of rationality, equilibrium assumptions are especially hard to justify in complex scenarios such as those found in electronic auctions. This may be due to a variety of reasons such as computational hardness, dependence on private information or the prior, non-intuitiveness of the situation, or repeated aspects of the game.

A recent paper [11] suggested using a much weaker assumption than the equilibrium assumption. This was demonstrated in the specific context of ad auctions (of the type that are also known as “sponsored-search auctions”), but can be applied to every repeated-game scenario. The weaker assumption that they promoted was that players *minimize regret* in the specific sense used in the regret-minimization literature [4] (sometimes known as “Hannan consistent”). This notion assumes that players manage to achieve at least as much utility as they could have gotten from playing any *fixed* action repeatedly. This makes minimal assumptions about the players’ learning and rationalizing ability. Certainly, if the players reach a Nash equilibrium, they must all be minimizing their regrets, but the regret-minimization assumption is strictly weaker, and, for example, holds even if the players reach any equilibrium from the much wider families of correlated or even coarse equilibria.

In [11], ad-auction data from Microsoft was analyzed under the regret-minimization assumption with the goal of estimating advertisers’ values that are not known to the search engine (Microsoft), by using the bids that are known to the search engine. Classic econometric methods were applied to this task in [15] and [1], and the average value estimates that were obtained in [11] seemed to be more or less in line with those from the two other methods ([15] and [1]). However, since their data set lacked the real values of the advertisers it is not clear how well each of these three econometric methods did, and, in particular, it is not clear how good are the estimates that were obtained based on the weak regret-minimization assumption.

It turned out that we had previously performed an experiment that yields exactly the type of data that can allow evaluation of the success of econometric methods in an ad-auction context [13]. In our experiment (described in Section 2), human players participated in a simulation of a stream of ad auctions. Unlike in field data, in our controlled experiment we assigned valuations to the players, and so we can directly compare the value estimates of the econometric methods with these real (assigned) values.

In the present paper we use this experimental data to evaluate the performance of the regret-based method for estimating players’ values, in comparison with the classic

©2017 International World Wide Web Conference Committee (IW3C2), published under Creative Commons CC BY 4.0 License.
WWW 2017, April 3–7, 2017, Perth, Australia.
ACM 978-1-4503-4913-0/17/04.
<http://dx.doi.org/10.1145/3038912.3052621>



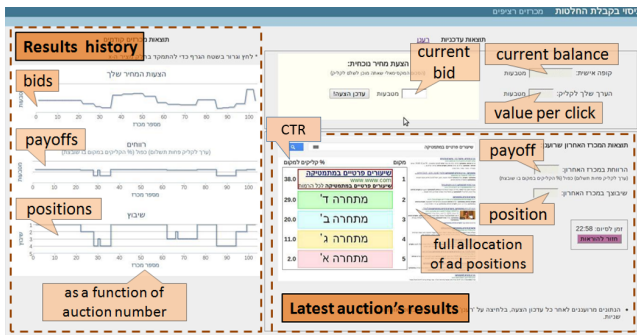



Figure 1: Game interface.

equilibrium-based estimation methods. We start by asking to what extent humans succeed in minimizing their regret in the repeated auction, and find that the answer depends on the “type” (private value) of the player (Section 3): while higher-type players indeed rationally minimize their regret, the lower-type players remain with high levels of regret. In Section 4 we specify our implementation of the regret-based estimation method, and in Section 5 we evaluate this method in comparison with the classic equilibrium-based methods. We find that the regret-based method manages to perform at least as well as the other methods, even in the truthful VCG auction where the equilibrium prediction is particularly strong. That is, our findings suggest that the weaker assumptions and general approach of the regret-based method are sufficient for this estimation task. In addition, we find that while all methods perform reasonably well for the higher-type players, they result in high errors for the lower-type players. In Section 6 we suggest several approaches to improve the accuracy of the estimates, using the regret-based econometrics and the understanding of the type-based bias that we identified.

2. OUR EXPERIMENT

We performed a controlled experiment where human subjects were asked to participate in a simulation of ad auctions, similar to those held by search engines like Google or Microsoft. This experiment was described in [13], which contains all the details as well as the results. In the experiment, we recruited participants in groups of five. In each instance of the experiment, the five participants simulated the roles of advertisers and had to compete in a stream of ad auctions that lasted 25 minutes. We used a flexible auction experimentation software that we developed that enabled us to control the auction details as well as the players’ knowledge and values. The auctions were conducted continuously, one auction per second, to a total of 1500 auctions within the 25-minute game. The participants could modify their bids at any time, and each auction was performed with the current settings of the bids. Each player was assigned a “type” at random, which was his private “valuation,” i.e., the monetary value that he obtained from each user who clicked on his ad (we used 21, 27, 33, 39, 45 “coins”). Each ad auction sold five ad positions with varying (commonly known) Click Through Rates (CTR) (we used 2%, 11%, 20%, 29%, 38%), which were displayed in a decreasing order of CTRs, so that the position on the top of the page received the highest CTR. Every time an advertiser with a valuation v won a position



(a) Bid modifications frequency: The average number of bid changes per minute made by a player, in each of the experimental conditions over time.

(b) Social welfare: The average social welfare achieved in each of the experimental conditions over time, on a linear scale from worst to best outcomes.

Figure 2: Previous results.¹

with CTR α , he got an income of $\alpha \cdot v$ from that auction. This income was added to his balance and the appropriate payment according to the auction rule was deducted from his balance. The players were given a graphical user interface in which they could modify their bids as often as they wished, and follow the results of the auctions so far. Figure 1 shows a screen shot of the user interface.


The experiment had a two-way (2x2) between-participant design; thus there were four experimental conditions. The two factors were:

1. **Payment Rule (the Auction Mechanism):** We compared the (theoretically appealing) VCG payment rule with the (commonly used) GSP payment rule. Both VCG and GSP auctions make the same allocation of positions – by decreasing order of bids – but their payment rule is different. Unlike GSP, the VCG is truthful; i.e., in every VCG auction it is a dominant strategy for every player to bid his true value (see [5, 15]).
2. **Valuation Knowledge:** While the starting point of analyzing behavior in auctions is the “valuation” of the bidder, it is questionable to what extent users are explicitly aware of this valuation. We compared the case where bidders were directly given their valuation (given value, GV), and were explained its significance, and the case where bidders were not directly given the valuation, but rather only see their payoffs – information from which the valuation may be deduced, but could alternatively be directly used to guide the bidding (deduced value, DV).

There were a total of 24 experimental sessions, 6 sessions for each of the 4 experimental conditions (thus there were 12 sessions for each factor). The groups (of five players each) were randomly assigned to the four experimental conditions, giving a total of $n = 120$ participants. For further details regarding the experimental setup see [13].

Out of the various results of this experiment, let us mention two that have particular relevance to the question at hand regarding the assumptions of the econometric model. The first result is that players in no way seemed to converge to an equilibrium. In fact, players kept modifying their bids

¹Figures 2a and 2b are taken from [13].



(a) By experimental conditions (b) By player types

Figure 3: Relative momentary regret over time. The momentary regret is the gap, computed separately for each minute, between the utility achieved by the players and the utility that would have been achieved had they played the fixed strategy that was optimal for that minute.

throughout the 25 minutes, and the frequency of bid modification increased over time. Figure 2a shows the average activity level (i.e., frequency of bid modification) as a function of time in the four experimental conditions. This was true in all auction formats, even when using the VCG auction rule with explicitly given valuations for which truthful bidding is a dominant strategy and so we would have expected truthful bidding as a strong and stable equilibrium prediction.

On the other hand, despite any lack of convergence to equilibrium, it seems that the auction was able to quickly achieve close to the “correct” (social-welfare maximizing) allocation of the slots, attaining over 80-90% of the optimal level of welfare² toward the end of the session (see Figure 2b), as well as extracting revenue to the auctioneer that matches and even exceeds the theoretically expected revenue. This suggests that the different auction formats were able to extract the value information from the users and utilize it, despite not having reached an equilibrium. This would thus seem to indicate that it should be possible to deduce the (hidden) value information from the (visible) auction bids, even though an equilibrium assumption does not seem to hold.

3. TO WHAT EXTENT DO HUMANS MINIMIZE REGRET?

We start by asking to what extent does the assumption that players minimize regret in the repeated game hold in our experiment with human players? I.e. we compare the actual utility achieved by each one of our participants to the optimal utility he could have achieved had he played an optimal fixed bid in a sequence of auctions. More formally, let \mathbf{b}^t denote the vector of actual bids played by the five players at time t , and denote by $U_i(b, \mathbf{b}_{-i}^t | v_i)$ the utility of player i whose value is v_i by bidding b at time t . Thus, player i ’s *actual utility* in a sequence of auctions $(\mathbf{b}^t)_t$ is $Actual_i((\mathbf{b}^t)_t | v_i) = \sum_t U_i(b_i^t, \mathbf{b}_{-i}^t | v_i)$, where b_i^t is the actual bid played by bidder i at time t . The *optimal utility* that player i could have obtained by playing a fixed bid repeatedly in all auctions $(\mathbf{b}^t)_t$ is $Opt_i((\mathbf{b}^t)_t | v_i) =$

²Our count here is within the possible range of allocations, such that a random allocation would get 50% of the welfare. In absolute terms, around 90-95% of the welfare was achieved. See [13] for details.

$max_b \sum_t U_i(b, \mathbf{b}_{-i}^t | v_i)$. The *regret* of player i whose value is v_i in a sequence of auctions $(\mathbf{b}^t)_t$ is defined to be the difference between his optimal and actual utilities in these auctions, i.e.,

$$Regret_i((\mathbf{b}^t)_t | v_i) = Opt_i((\mathbf{b}^t)_t | v_i) - Actual_i((\mathbf{b}^t)_t | v_i)$$

For the VCG auction formats, the optimal fixed bid is certainly the dominant strategy of bidding the true value, but for the GSP auction formats the optimal bid depends on the bids of the other players and is only evident in hindsight (and thus could not even theoretically be known to our bidders in real time.)

We define the *regret of a set S of players* just additively over the players in S . That is, $Actual_S = \sum_{i \in S} Actual_i$, and $Opt_S = \sum_{i \in S} Opt_i$. Then, the regret of S at a sequence of auctions $(\mathbf{b}^t)_t$ is $Regret_S((\mathbf{b}^t)_t) = Opt_S((\mathbf{b}^t)_t) - Actual_S((\mathbf{b}^t)_t) = \sum_{i \in S} Regret_i((\mathbf{b}^t)_t | v_i)$. Since the level of regrets depends on the magnitude of the utilities, we present the regret levels as percentages of the corresponding optimal outcomes. Specifically, we define the *relative regret* of S at a sequence of auctions $(\mathbf{b}^t)_t$ by


$$RelativeRegret_S((\mathbf{b}^t)_t) = \frac{Regret_S((\mathbf{b}^t)_t)}{Opt_S((\mathbf{b}^t)_t)}$$

First, in order for any regret minimization to be possible, we would need to see some learning by the bidders as time progresses, and indeed that is what we find. Figures 3a and 3b show the relative “momentary regret” of the players over time by experimental conditions and by player types, respectively, where the momentary regret is the regret computed separately in each minute of the experiment. Consistent with the regret-minimization assumption, we find that the momentary regret decreased over time in all conditions and types. Specifically, for each of the four experimental conditions, the average momentary regret in the first third of the game is significantly higher than in the last third³ (N=6 sessions, Wilcoxon paired two-sided signed rank test, $p < 0.05$ except for VCG-GV for which $p = 0.06$) and the same for each of the five types of players (N=24, $p < 0.001$ for each type except for 33 for which $p < 0.03$). The decrease was faster at the beginning when players acquired experience in the game. This seems to be consistent with the suggestion raised by [11], that higher levels of regret may indicate that bidders are in their initial learning phase. The regret toward the end of the game reached low levels of around 15% regret for the different conditions, but, as can be seen in Figure 3b, remained substantial for the lower-type players.

At this point we can look at the overall regret achieved by our players, and answer the question to what extent is it true that bidders in ad auctions minimize regret, at least after the initial learning phase. We compute the “total regret” of a bidder based on the second half of the game (750 auctions), where regret levels stabilized, indicating that players had completed the initial learning phase and were experienced enough in the game.⁴ Looking according to the experimental conditions, the total regret seems reasonably low: a 10-20% loss (Figure 4a). The total regret in GSP-DV

³When considering momentary regret, the first and last thirds of the game refer to the first and last 8 minutes, respectively. The results are robust to other choices of partitioning.

⁴Qualitatively our findings are robust to modifications of the definition of the initial learning phase.



(a) By experimental conditions (b) By player types

Figure 4: Relative total regret of the bidders according to experimental conditions and according to types of players. The total regret is computed over the second half of the auctions game (750 auctions), and is presented as a percentage of the corresponding optimal outcome.

was higher than in the other three conditions, as could be expected, since this condition combines the two difficult settings – GSP and DV – where bidders needed to learn both their opponents’ behavior and their own values. Yet, the difference was significant at the 5% level only in comparison with the GSP-GV condition.

However, a different story is revealed when examining bidders’ regret by their types: we found that the type of a player, i.e., the value assigned to him at random and used to determine his payoffs in all auctions, also had a significant effect on his regret-minimization performance. Figure 4b shows the relative total regret for the five types of players. It can be seen that the lower the type is, the higher the player’s total regret: the highest-type players achieved low levels of regret – a less than 6% loss – while the lowest-type players remained with very high levels of regret – with a more than 50% loss of their optimal outcome.⁵ Specifically, there is a significant negative correlation between the player’s type and his total regret ($N = 120$, $r = -0.68$, $p < 0.001$). These differences between types were consistent throughout the game, as can be seen in Figure 3b. Therefore, high levels of regret may indicate not only that a player is in his initial learning phase, as suggested by [11], but also that he might be of a lower type relative to his opponents.

Thus, we see that players who were assigned (by chance) with “poor” types turn out to play significantly less rationally than players who were lucky to get “rich” types. While we were surprised at first to see such a significant gap, on second thought this behavior seems quite intuitive: when the low-value players play rationally they tend to “lose” in the auction, i.e., win the low CTR slots. This is quite frustrating and so they keep trying to “win,” but to no avail since they “should” be losing the auction according to their true value. This gap in rational play is consistent with other irrational behaviors reported in [13] – overbidding and high frequency of bid changes – that were also correlated with the player types. It may be interesting to relate this to other settings where it was found that “the poor act irrationally” (see, e.g., [2, 3, 14]). However, our controlled setting proves that this “irrational” behavior cannot be explained by any characteristics of the poor themselves (e.g., lower education), but rather only being poor (in an inferiority) relative to the

⁵Notice that the high percentage of regret of low-type players is not so high (though still higher) in absolute additive terms, since they tend to win the low CTR slots anyway.

others affected players’ rationality. These findings may be explained in terms of “auction fever” [9, 6], as a form of an “illusion of control” bias [10], or in terms of a stronger bias toward “winning.”

Finally, we should also note here that the regret is positive: in principle it is possible to achieve negative regret (in the GSP auction) using time-varying bids. Yet not a single one of the 60 players who participated in our GSP auctions managed to achieve negative total regret. I.e., none of our players managed to utilize the *dynamics* in the repeated game to their advantage and to play better than the fixed-bid benchmark.

4. REGRET-BASED ECONOMETRICS

Before we proceed to the evaluation results, let us formally specify our implementation of the regret-based method, suggested by [11], for estimating bidders’ valuations from their bids. For this estimation task, we now assume that we observe only the bids played by the bidders and the fixed CTRs.

The regret-based estimation method is based on the assumption that players use learning strategies that minimize their regret in the repeated game. That is, they use bidding strategies by which, over time, their utilities are not much worse than the optimal utilities they could have achieved by playing the best fixed bid in hindsight. Thus, the method sets the value estimate for a bidder, in a sequence of auctions, to be the value that minimizes his regret in these auctions. Again, this assumption is weaker than the standard approach in econometrics, which relies on the assumption that the observed game reaches a static equilibrium, and it is general for different games (and, in particular, for different auction rules).

Our implementation of the regret-based method is given in the “Regret-Minimization” estimation procedure (Algorithm 1). The procedure is given as input a sequence of bid profiles $(\mathbf{b}^t)_t$ (in our implementation we used $(\mathbf{b}^t)_{t=751}^{1500}$, excluding the first half of the game as an initial learning phase, as mentioned above), as well as the players’ utility functions $U_i(b, \mathbf{b}_{-i}^t | v)$. Regret-Minimization estimates value \hat{v}_i of each bidder i as follows: it begins by fixing the sets of possible valuations V_i and of possible bids B_i to consider (we set $V_i = B_i = \{1, 2, \dots, 60\}$ for all bidders). Then, for every possible value $v \in V_i$, it computes $Regret_i((\mathbf{b}^t)_t | v)$ – the regret of player i had his value been v – which in turn requires the computation of $Actual_i((\mathbf{b}^t)_t | v)$ and $Opt_i((\mathbf{b}^t)_t | v)$. Notice that in VCG, finding $Opt_i((\mathbf{b}^t)_t | v)$ (in line 8) does not require iterating over all $b \in B_i$, as this optimal outcome is always obtained at $b = v$. The estimate \hat{v}_i is then the value that minimizes this regret.^{6 7}

We applied the Regret-Minimization procedure to each of the bidders in the experiment, and we evaluate the accuracy of the estimations in the next section. Figure 5 presents the regret curves as a function of value, according to the types of players and the experimental conditions. The minimum point for each player serves as the estimate of his value, and is clearly visible in the graphs. As can be seen, this minimum

⁶The specific implementation in [11] selected the value that minimizes the *relative* regret.

⁷Only for a single bidder of the 120 bidders in our data, the optimal v was not unique but could only be located in the range of values ($v \in [49, 52]$). We used as our estimate the middle point in this range.

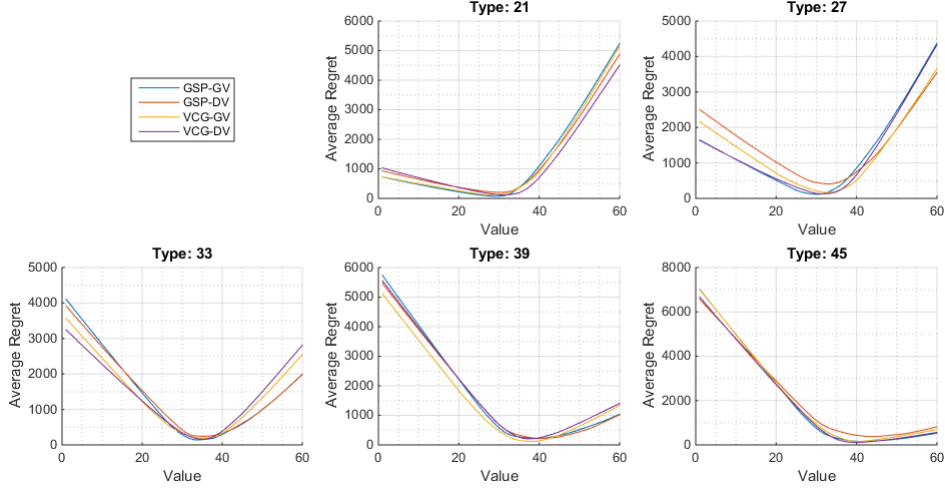


Figure 5: Average total regret as a function of value according to player types and experimental conditions. The value that gives the minimal total regret for a player is given as the estimate for the player’s value.

Algorithm 1 Regret-Minimization: regret-based estimation of bidders’ valuations.

- 1: **Input:** actual bid sequence $(\mathbf{b}^t)_t$ and functions $U_i(b, \mathbf{b}_{-i}^t|v)$ modeling players’ utilities
 - 2: **Output:** players’ value estimates $\hat{v}_1, \dots, \hat{v}_n$
 - 3: **procedure** REGRET-MINIMIZATION
 - 4: **for** bidder $i \in N$ **do**
 - 5: denote the set of possible valuations by V_i and the set of possible bids by B_i
 - 6: **for** $v \in V_i$ **do**
 - 7: $Actual_i \leftarrow \sum_t U_i(b_i^t, \mathbf{b}_{-i}^t|v)$
 - 8: $Opt_i \leftarrow \max_{b \in B_i} \sum_t U_i(b, \mathbf{b}_{-i}^t|v)$
 - 9: $Regret_i((\mathbf{b}^t)_t|v) \leftarrow Opt_i - Actual_i$
 - 10: $\hat{v}_i = \arg \min_{v \in V_i} Regret_i((\mathbf{b}^t)_t|v)$
 - 11: **return** the estimates $\hat{v}_1, \dots, \hat{v}_n$
 - 12: **end procedure**
-

is achieved at a higher value as we move from lower to higher types, consistent with the regret-minimization model. Also notice the asymmetry in the slopes of the regret, as well as how it changes between low and high types, showing that low types would have suffered low regret had their values been lower (as they tend to get the low CTR positions anyway) and higher regret had their values been higher, while the opposite is observed for the higher types.

5. EVALUATING REGRET-BASED ECONOMETRICS

We are now ready to evaluate the success of the Regret-Minimization method in estimating players’ private valuations, in comparison with standard econometric methods that rely on the equilibrium assumption. The same regret-based estimation procedure (Algorithm 1) works for both the VCG and GSP auctions, with the payment rule taken into account in the calculation of the utilities, $U_i(b, \mathbf{b}_{-i}^t|v)$. However, since these auction mechanisms have different equilibria, standard methods use different procedures in these two

cases, and thus we perform the evaluation separately for each mechanism.

We assess the quality of the estimation methods in percentages based on the mean squares of the relative errors. Specifically, for every player i whose true value is v_i , we compute the relative estimation error: $error_i = \frac{1}{v_i} |v_i - \hat{v}_i|$, where \hat{v}_i is the value estimate for player i . The estimation error on a set of players S is: $error(S) = \sqrt{\frac{1}{|S|} (\sum_{i \in S} error_i^2)}$. We base the evaluation on the second half of the auctions game: a stream of 750 auctions, excluding the first half as an initial learning phase (as mentioned above).

5.1 Evaluating Estimations in VCG Auctions

We start by considering the VCG auction. The usual econometric treatment will note that players have dominant strategies, so it should be a strong prediction that they all bid these dominant strategies in equilibrium. The classic econometric method will thus take as the model that each player bids his true value in each round, plus an error term: $b_i^t = v_i + \epsilon_i^t$, where v_i is the true value and b_i^t is the bid in round t . The estimate of the hidden value v_i from the visible data of the b_i^t s that were played in a sequence of T auctions will be $\hat{v}_i = (\sum_{t=1}^T b_i^t) / T$, as this average minimizes the sum of the squares of ϵ_i^t : $\sum_t (v - b_i^t)^2$. We applied the Regret-Minimization method and the standard method of taking the average bid (“Average-Bid”) to estimate the valuations of the bidders in the VCG sessions. Overall, considering all 60 VCG bidders, Regret-Minimization and Average-Bid performed very similarly, with total estimation errors of 23.37% and 23.38%, respectively. The errors distributed similarly, both with high variance between players: $\mu = 17.69\%$ and $\sigma = 15.39\%$ using Regret-Minimization, and $\mu = 17.37\%$ and $\sigma = 15.79\%$ using Average-Bid. The similarity persists when considering errors by information settings or by types of players, as can be seen in Figures 6a and 6b. Thus, the estimation obtained by the Regret-Minimization method, which is general for all auction formats, does not fall from the estimation of the standard approach, even though the




Figure 6: Estimation results in VCG sessions: Regret-Minimization (RM) vs. Average-Bid (AB) methods. (a) and (b) present the estimation errors by value-information conditions and by types of players, respectively. (c) presents the average estimates of the two methods according to types of players. The “differences table” below presents the averages of the differences between the value estimates and the true values of the bidders, and the differences between the value estimates using the two methods, by types, as indicated in the first column. Significant differences are marked by *, **, and ***, for significance at levels of 10%, 5%, and 1%, respectively.

specific equilibrium prediction for VCG is strong and is of the simple strategy of bidding the true value.

Figure 6a shows that the estimation errors of both methods are somewhat larger in the deduced-value setting (DV) than in the given-value setting (GV), as could be expected, and the regret-based method seems somewhat better in the former setting and somewhat worse in the latter. However, these effects were not statistically significant.

A significant effect is revealed for both methods when considering errors by types of players: as expected from the finding in Section 3 that the low types deviate significantly from playing rationally, the estimation errors of both methods for the low types are far larger than those for the high types (see Figure 6b). Specifically, for each of the two methods, there is a significant negative correlation between the estimation error and the player type ($N = 60$, $\rho = -0.65$ and $\rho = -0.52$ for Regret-Minimization and Average-Bid, respectively, $p < 0.001$). The total estimation error of both methods for the lowest-type players is reaching around 40%, while for the highest-type players the error is much lower – around 13%. Despite the overall similarity between the two methods, Figure 6b shows that while Regret-Minimization is more accurate for the middle-valued types, Average-Bid is more accurate for the lower-type players.⁸ In addition, as opposed to the significant and gradual differences between estimation errors for different types using Regret-Minimization (consistent with the gradual pattern of regret levels shown in Figure 4b), using Average-Bid the differences are only due to the error in estimating the lowest type’s valuation, which was significantly higher than that of each of the other types. But still, the errors of the two methods were not significantly different at the 5% level also when tested for each type separately.

Finally, we found that Regret-Minimization and Average-Bid are different in terms of the average estimates that they give. Figure 6c compares the average estimates obtained by the two methods and the bidders’ true values (by types).

⁸This pattern is robust for different choices that we tried of initial learning phases (i.e., other than the first half of the game), and for some choices the interaction between method and player types reaches statistical significance at 5% level.

As can be seen, Regret-Minimization significantly overestimates the valuations of the three lower-type bidders, and significantly underestimates the valuations of the highest-type bidders. The direction of mistakes by Regret-Minimization gradually changes from low to high types. These findings suggest that the low-type bidders tend to play as if their valuation were higher and the opposite holds for the bidders with the highest value, who tend to play as if their valuation were lower than it really was. Interestingly, these tendencies do not clearly arise from the average estimates of the Average-Bid method, and thus must be hidden in the dynamic strategies played by the players relative to their opponents. In Section 6 we demonstrate how we may improve the accuracy of estimations by taking advantage of the differences between the two methods.

5.2 Evaluating Estimations in GSP Auctions

For the GSP auction the situation is much more complicated for equilibrium-based econometric methods since there are no dominant strategies and there exist multiple equilibria. There are two basic approaches in the literature for deducing bidders’ valuations in GSP auctions. In [15] it is suggested that the players should reach the equilibrium of the full-information one-shot GSP game that gives the VCG-prices (hence the “VCG-like” equilibrium). Assuming that this is indeed the case, then at each time step t in a sequence of auctions, one may deduce values \hat{v}_i^t for all players i from the actual bids b_i^t , such that the bids are this VCG-like equilibrium of these deduced values. The final estimate is then the average of these \hat{v}_i^t . Some complications arise when this is attempted on real data since it is often the case that the bids do not correspond to an equilibrium of any tuple of values. In these cases we follow [15] and perturb the bid observations in the minimal possible way so as to satisfy the equilibrium constraints, and set the final estimates to the perturbed values.⁹ These and other complications of the

⁹In fact, only 13.3% of the auctions were consistent with the equilibrium inequalities without perturbing their data. However, similar to [15], we observed that the required perturbations were relatively small.




Figure 7: Estimation results in GSP sessions: Regret-Minimization (RM) vs. Best-Response (BR) vs. VCG-like-NE (VNE) methods. (a) and (b) present the estimation errors by value-information conditions and by types of players, respectively. (c) presents the average estimates of the three methods according to types of players. The “differences table” below presents the averages of the differences between the value estimates and the true values of the bidders, by types, as indicated in the first column. Significant differences are marked by *, **, and ***, for significance at levels of 10%, 5%, and 1%, respectively.

“VCG-like-NE” method are discussed in [15] and in the full version of this paper.

A second method was suggested by [1] where bidders participate in a large number of auctions, and receive feedback that can vary from auction to auction. The basic assumption is that each bidder is best-responding to the *distribution* that he faces (by placing a single bid). Specifically, given a sequence of auctions, define functions $Q_i(b_i)$ and $TE_i(b_i)$ as the expected CTR and the expected total expenditure, respectively, of bidder i by bidding b_i . Thus, his expected utility with valuation v is $Q_i(b_i) \cdot v - TE_i(b_i)$. Against *smooth* distributions the best bid would be a strictly increasing function of the value. In these cases, the valuation of bidder i who maximizes his expected utility by bidding b_i can be recovered using the first-order condition by $\hat{v}_i = \frac{\partial TE_i(b_i)/\partial b_i}{\partial Q_i(b_i)/\partial b_i}$. When applying this method to actual data complications arise, and there are many possible implementations. We have tested several implementation variants, and in the implementation we chose (referred to as the “Best-Response” method), we used the average bid that a bidder played as his best-response to the distribution of the bids of the others (since bids were not constant), and found the value by optimizing directly using grid search (since the empirical derivatives had their own complications). Details of implementation and complications of this method are discussed in [1] and in the full version of this paper.

Figures 7a and 7b compare the estimation error of the regret-based method with those of the two “classic” equilibrium-based methods. Overall, the regret-based method tends to be better than the other two methods, but the difference is statistically significant at the 5% level only in a comparison with the VCG-like-NE method.¹⁰ Specifically, over all 60 GSP bidders, the total estimation errors obtained using Regret-Minimization, Best-Response, and VCG-like-NE were 25.51%, 28.94%, and 30.60%, respectively. The average of the estimation errors using Regret-Minimization, Best-Response and VCG-like-NE, were 18.07%, 20.04%, and 21.45%, respectively, and the standard deviations were 18.00%, 20.87% and 21.82%, respectively.

¹⁰The error using the Best-Response method was lower than the error using VCG-like-NE at a significance level of 10%.

Considering estimation errors by information settings (Figure 7a), the errors of all three methods are somewhat larger in the deduced-value setting (DV), in which case the regret-based method has an advantage over the other two methods. While this appears to be quite consistent (for different selections of the initial learning phase), it is only partially statistically significant.

As in the case of VCG, the significant effect is revealed when considering errors by player types (see Figure 7b): all three methods succeed better in estimating valuations of the higher-type players (who succeeded better in minimizing their regret) than of the lower-type players (who remained with high levels of regret, see Section 3), for whom all methods perform very poorly. Specifically, for each of the three methods, the estimation errors are negatively correlated with the player type ($N = 60$, $\rho = -0.59$, $\rho = -0.47$, and $\rho = -0.57$ for Regret-Minimization, Best-Response, and VCG-like-NE, respectively, $p < 0.001$). When comparing estimation errors for each type separately, the differences are statistically significant only for types 27 and 33: the error using Regret-Minimization is lower at 5% level than the error using VCG-like-NE for types 27 and 33, and Best-Response is also better than VCG-like-NE at 5% level for type 33.

Finally, we found that the three methods are similar in terms of the average estimates that they give. As can be seen in Figure 7c, all three methods tend to overestimate the valuation of the three lowest types. Thus, also in GSP, where bidders are expected to shade their bids, lower types play as if their value were higher, in consistency with their high levels of regret described in Section 3. The estimations of the two highest types are not significantly in a particular direction relative to the values. The direction of mistakes changes gradually across types, as in the case of VCG.

6. POSSIBLE IMPROVEMENTS

We have seen, in the previous section, that the regret-based estimation method gives value estimates that are competitive with standard econometric methods both in GSP and in VCG. Thus, the weaker assumptions that the regret-based method makes seem sufficient for this estimation task. However, we have also seen how *all* methods have quite high errors, particularly on the lower-type players, who, as we


found, tend not to follow the methods’ underlying assumption of rationality. Therefore we ask: how can we improve the accuracy of the estimations using regret-based econometrics and our understanding of the biases of human players?¹¹

First, it may be useful to combine estimates of the different estimation methods, taking advantage of the differences between them, as the different methods catch different aspects of behavior and may have different sources of mistakes. For example, we averaged the estimates of the Regret-Minimization and Average-Bid methods for each of our VCG bidders, and the estimation error of the combined method was lower than the error of each of the two methods separately (23.37%, 23.38%, and 21.73%, for the Regret-Minimization, Average-Bid, and the combined method, respectively; however the differences in our results were statistically significant at the 5% level only with the Regret-Minimization method).

A second approach is to “un-bias” the estimates according to player types, taking advantage of our finding that the bias of all methods was correlated with the types of players. Since this seems to result from systematic differences in behavior between player types, where the lower-ranked players tend (irrationally) to play as if their value were higher, it is useful to know to what extent we could have improved the estimates had we known how to quantify and correct this bias. We demonstrate the potential gains of un-biasing the estimates on our GSP bidders: for every bidder, we calculated his most frequent position in the auction, and used this information to “un-bias” our estimate according to the average bias of the estimates of his (deduced) rank. For example, let \hat{v} be the estimate using Regret-Minimization of a GSP bidder who was most frequently observed in the last position. Then, we use the finding that on average Regret-Minimization overestimates the lowest type as 137% of his true value, and we fix the estimate accordingly, so that $\hat{v}_{fixed} = \hat{v}/1.37$. This correction of the bias significantly improved the results of all three methods for the GSP bidders, as can be seen in Figure 8a. Of course, since this bias-correction factor (137% in the example) came from within the sample using the real values, one can only interpret these reduced estimation errors as an upper bound on the gains from correcting by types. Our demonstration here suggests that these gains may be substantial, indicating that it is worthwhile to further understand the behavioral tendencies of the different types.

Finally, we suggest improving the estimates of the Regret-Minimization method by taking advantage of the shape of the regret curves, which we found to be different for the different types of players (see Figure 5). Since we observed that human bidders do tend to minimize their regret, but do not minimize perfectly, we suggest setting the estimate based on the entire regret curve rather than just on the minimal regret point. For example, we set the value estimate of a bidder i to the weighted average of all possible valuations, with weights that are proportionally decreasing in

¹¹One general approach one may try is to “clean” the data, i.e., remove outlier observations that might bias the estimates. However, our attempts in this direction did not improve the estimation accuracy and even slightly increased the error, indicating that this error is not a result of the ends of the bid distributions but is inherent in the bidders’ behavior.



(a) Improving by “un-biasing” the estimates according to player types. Original and fixed (un-biased) estimation errors for the GSP bidders. (b) Improving by the “regret-based average” method, which estimates using the entire regret curve rather than just the minimal regret point.

Figure 8: Possible improvements of the value estimates.

the regret: $\hat{v}_i = \frac{\sum_{v \in V_i} v \cdot \text{Regret}_i^{-1}(\cdot|v)}{\sum_{v \in V_i} \text{Regret}_i^{-1}(\cdot|v)}$. This “regret-based average” method is completely applicable based on observed data only, and, as can be seen in Figure 8b, significantly improves the accuracy of the estimates for our bidders. In a separate paper we explore this approach further [12].

7. CONCLUSIONS AND FURTHER DIRECTIONS

We have demonstrated that, at least in our repeated ad-auction experiment, players do act to reduce their regret, and that the regret-based method suggested by [11] for estimating players’ values from their bids is (at least) competitive with “classic” equilibrium-based econometric methods. We find this to be especially significant due to the generality and simplicity of the regret-based method: while “classic” econometric methods require choosing between different interpretations of equilibria as well as among many significant implementation details, and to tailor the method to the specific equilibrium assumptions, the regret-based econometric method does not require any specific tailoring and hence is much easier to specify and implement. We furthermore speculate that this simplicity may reasonably allow this method to be further improved, and we suggested possible improvement approaches of this form.

En route we also observed irrational bidding in our experiment and identified a distinct human bias: players with “low” types tend to overbid leading to high regret, and resulting in increased estimation errors in all the econometric methods that we tried. It is an interesting challenge to use our understanding of this behavioral bias to improve the estimates, and we demonstrated that the improvements based on this bias may be substantial. Regret-minimization is a simple and powerful method that can be easily modified to capture different behavioral phenomena and to take advantage of insights from behavioral disciplines. Obviously, further evaluation of regret-based econometrics in more scenarios is called for.

8. ACKNOWLEDGMENTS

This research is supported by ISF grant 1435/14 administered by the Israeli Academy of Sciences, by Israel-USA Binational Science Foundation (BSF) grant number 2014389, and by the Adams Fellowship Program of the Israel Academy of Sciences and Humanities.

9. REFERENCES

- [1] S. Athey and D. Nekipelov. A structural model of sponsored search advertising auctions, 2010.
- [2] M. S. Barr. *No Slack: The Financial Lives of Low-Income Americans*. Brookings Institution Press, 2012.
- [3] M. Bertrand, S. Mullainathan, and E. Shafir. A behavioral-economics view of poverty. *American Economic Review*, 94(2):419–423, 2004.
- [4] A. Blum and Y. Mansour. Learning, regret minimization, and equilibria. In *Algorithmic Game Theory*. Cambridge: Cambridge University Press, 2007.
- [5] B. Edelman, M. Ostrovsky, and M. Schwarz. Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords. *American Economic Review*, 97(1):242–259, 2007.
- [6] J. E. Heyman, Y. Orhun, and D. Ariely. Auction fever: The effect of opponents and quasi-endowment on product valuations. *Journal of Interactive Marketing*, 18(4):7–21, 2004.
- [7] J. H. Kagel and D. Levin. *The Handbook of Experimental Economics, Volume II*, chapter Auctions: A Survey of Experimental Research, 1995 - 2008. Princeton: Princeton University Press, 2010.
- [8] J. H. Kagel and A. E. Roth. *The Handbook of Experimental Economics*. Princeton University Press, Princeton, N.J., 1995.
- [9] G. Ku, D. Malhotra, and J. K. Murnighan. Towards a competitive arousal model of decision-making: A study of auction fever in live and internet auctions. *Organizational Behavior and Human Decision Processes*, 96(2):89–103, 2005.
- [10] E. J. Langer. The illusion of control. *Journal of Personality and Social Psychology*, 32(2):311–328, 1975.
- [11] D. Nekipelov, V. Syrgkanis, and E. Tardos. Econometrics for learning agents. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation*, EC '15, pages 1–18, New York, NY, USA, 2015. ACM.
- [12] N. Nisan and G. Noti. A "quantal regret" method for structural econometrics in repeated games. *arXiv:1702.04254*, 2017.
- [13] G. Noti, N. Nisan, and I. Yaniv. An experimental evaluation of bidders' behavior in ad auctions. In *Proceedings of the 23rd International Conference on World Wide Web*, WWW '14, pages 619–630, New York, NY, USA, 2014. ACM.
- [14] T. P. H. Steven J. Katz. Socioeconomic disparities in preventive care persist despite universal coverage breast and cervical cancer screening in ontario and the united states. *JAMA*, 272(7):530–534, 1994.
- [15] H. R. Varian. Position auctions. *International Journal of Industrial Organization*, 25(6):1163–1178, 2007.