Bidding with Limited Statistical Knowledge in Online Auctions

Chong Jiang
Coordinated Science Laboratory and Department of Electrical and Computer Engineering
jiang17@illinois.edu

Carolyn L. Beck
Coordinated Science Laboratory and Department of Industrial and Enterprise Systems Engineering
beck3@illinois.edu

R. Srikant
Coordinated Science Laboratory and Department of Electrical and Computer Engineering
rsrikant@illinois.edu

University of Illinois at Urbana-Champaign
Urbana, IL 61801, USA

ABSTRACT

We consider online auctions from the point of view of a single bidder who has an average budget constraint. By modeling the rest of the bidders through a probability distribution (often referred to as the mean-field approximation), we develop a simple bidding strategy which can be implemented without any statistical knowledge of bids, valuations, and query arrival processes. The key idea is to use stochastic approximation techniques to automatically track long-term averages.

Keywords

Online auctions, Budget constraints, Stochastic approximation

1. INTRODUCTION

Sponsored search is a type of online advertisement where paid links are shown next to the search results of relevant queries, yielding substantial revenue for search engine providers like Microsoft and Google. It is also a problem of theoretical interest, since bidders are competing for overlapping keywords and targeted demographics, and where the number of bidders, ads, and ad slots on webpages are all large. Even though the ad displayed for each query can be modeled as a second price auction, the aforementioned other constraints complicate the problem.

We consider a model from a single bidder's point of view, and in our model the bidder has an average budget constraint, expressed as a constraint on the expected payment, but which can be interpreted as a constraint on the time average of the payment. As in common in this literature, if we model all other bidders' strategies by a probability distribution over their largest bid, then we get a simple static optimization problem for each bidder. However, the solution to this static optimization problem requires knowledge of the probabilities of the maximum of others' bids, as well as the distribution of one's own valuation, statistics of the arrival process of relevant queries, and in more complicated settings with auctions involving multiple ad slots, we also require click-through rates, the distributions of the second highest bid of the opponents, etc. The main contribution of this work is to show that the average budget model considered here allows us to compute the optimal bid using stochastic approximation without the aforementioned statistical details. In this short abstract, we present simulation results which suggest that, despite not knowing the statistics fully, the expected regret in the payoff, and any budget overdraft or underdraft are both very small.

1.1 Relationship to Prior Work

The distributional assumption that we make about the opponents' bids is often called the mean-field approximation, which has been studied in a Markov decision problem context in [4, 3]. In [4], the focus is on learning the distribution of the valuation, while the focus of [3] is on budget constraints. Our model is closer to that of [3, 1], but our use of an average budget constraint rather than the strict budget constraint they use allows us to obtain a solution that does not require statistical knowledge of the system parameters. We show that under an average budget constraint, an underbidding factor also appears in the solution, which we then estimate through stochastic approximation (SA) [7, 5, 2]. In [3], a Markov Decision Problem (MDP) must be solved to obtain the factor by which one underbids an item's true valuation. However, this MDP is nearly intractable and therefore, a fluid approximation is used to calculate this factor in a more tractable manner. However, even this approximation still relies on the knowledge of the probability distribution of the opponents' maximum bid. In [1], under an additional assumption of homogeneous bidders, this underbidding strategy is shown to be the unique fluid mean-field equilibrium.

Compared to an MDP formulation, the use of an SA approach (when used with a small but fixed stepsize) allows one to track non-stationary behavior on the bids, valuations, relevant query arrivals, etc. Additionally, our formulation allows us to change our bid at whatever time scale we choose. For example, suppose that bids are placed at multiple geographical locations, then computational considerations involved in keeping track of the massive amount of data generated from multiple bids may prevent a query-by-query update of the underbidding factor. Then it is natural to fix a time period over which the underbidding factor is held constant, and changed only at the beginning of each period based on the expenditure, aggregated over all locations, during the previous period. Our formulation allows us to incorporate such practical constraints.

We note that the average budget constraint has been used previously in [8] in a different context, from the point of view of a search engine provider. In that model, there is no bidding involved and the goal is to maximize either click-through rates or impressions subject to an average budget constraint.

2. MODEL AND ALGORITHM

We define the sequence of queries to be indexed by $i \in \mathbb{N}$, with a valuation v_i, drawn i.i.d. from some distribution f_v. Suppose the highest bid from the other bidders is b'_i, drawn i.i.d. from $f_{b'}$. We also assume an average budget constraint of B for each time period, indexed by $t \in \mathbb{N}$. We allow the number of queries per time period to be random, and let N_t denote the total number of queries during the time periods $\{1, 2, \ldots, t\}$, where N is assumed to have independent and stationary increments. Thus the number of queries in time period t is $N_t - N_{t-1} \sim f_{N_1}$, where $N_0 = 0$. Two such examples of N are $N_t = r \cdot t$ and a Poisson process of rate r.

Since we observe only v_i, we seek to find a bidding strategy $b(v)$ that will maximize the expected reward while satisfying the average budget constraint, i.e.

$$\max_b E_{v,b',N_1}\left[\sum_{i=1}^{N_1}\left(v_i - b'_i\right) I_{\left\{b_i > b'_i\right\}}\right]$$

subject to

$$E_{v,b',N_1}\left[\sum_{i=1}^{N_1} b'_i I_{\left\{b_i > b'_i\right\}}\right] \le B.$$

Using a Lagrange multiplier λ^*, we now find

$$\max_b E_{v,b',N_1}\left[\sum_{i=1}^{N_1}\left(v_i - b'_i - \lambda^* b'_i\right) I_{\left\{b_i > b'_i\right\}}\right].$$

Note that the bid is a function of v and N_1, but we can these values as given when optimizing for b, so the above maximization can be equivalently written as

$$\max_b E_{b'}\left[\sum_{i=1}^{N_1}\left(v_i - b'_i\left(1+\lambda^*\right)\right) I_{\left\{b_i > b'_i\right\}}\right]$$
$$= \max_b \sum_{i=1}^{N_1}\int_0^{b_i}\left(v_i - b'_i\left(1+\lambda\right)\right) f_{b'}\left(b'\right) db'.$$

Differentiating w.r.t. b and equating to 0 yields

$$\left(v_i - b'_i\left(1+\lambda^*\right)\right) = 0$$
$$\Longrightarrow b_i = \frac{v_i}{1+\lambda^*}. \tag{1}$$

By complementary slackness, λ^* is either the positive solution to

$$E_{v,b',N_1}\left[\sum_{i=1}^{N_1} b' I_{\left\{\frac{v}{1+\lambda^*} > b'\right\}}\right] - B = 0 \tag{2}$$

if it exists, or 0 otherwise. Note that this optimal bidding is similar in form to that found in [3], where they refer to $\frac{1}{1+\lambda^*}$ as the bid shading (a.k.a. underbidding) factor.

Now we focus on computing λ^* when the distributions of b', v and N_1 are unknown. Later we will comment on how to use any available partial knowledge of these distributions. The form of (1) and (2) suggests the following stochastic approximation update rule, where ϵ_t is the stepsize:

$$\lambda_{t+1} = \left[\lambda_t + \epsilon_t\left(\sum_{i=N_t+1}^{N_{t+1}} b'_i I_{\left\{b_i > b'_i\right\}} - B\right)\right]^+ \tag{3}$$

and a bidding rule $b_i = \frac{v_i}{1+\lambda_{t(i)}}$, where $t(i)$ is the time period that includes query i, i.e.

$$N_{t(i-1)} < i \le N_{t(i)}.$$

Applying convergence results from [2, 5], we can show that under mild conditions and a sufficiently slowly decreasing sequence ϵ_t, we have that $\lambda_t \to \lambda^*$ a.s. regardless of the initial condition λ_0. One can also use a fixed stepsize (i.e., ϵ_t is a constant) to track any possible non-stationarities in the random processes involved. In this case, instead of almost sure convergence, one can provide probabilistic guarantees on how close λ_t is to λ^* in steady-state.

While the convergence (or closeness) of λ_t to λ^* is important, it is also important to keep track of the total regret (the difference between the expected optimal payoff and the realized payoff under our algorithm) and the amount by which we overshoot or undershoot the budget (called overdraft and underdraft, respectively) as functions of time. More precisely, if the amount that the bidder has spent up to time t is $\tilde{B}_t$, then we define the underdraft to be $B \cdot t - \tilde{B}_t$. If the underdraft is negative, then one can also call it an overdraft. In the longer version of the paper, we will study these quantities (regret and underdraft) analytically; in this version, we provide simulation results (in the following section).

Note that while convergence may hold for any λ_0, the time required to approach λ^* and the magnitude of the regret and budget underdrafts are dependent on the particular value of λ_0. This suggests that if we are given partial or approximate information about the system distributions, we should take into account $\hat{\lambda}^*$, the solution to (2) using the available distributions in place of $f_{b'}$, f_v, and f_{N_1}, when deciding λ_0.

One extension to the basic model is to replace the second price auction with either a generalized second price (GSP) or Vickrey-Clarke-Groves (VCG) auction with m ad slots, and which takes into account click-through rates. Another extension is to consider a fixed total budget over a finite time horizon. One could treat this using the average budget formulation proposed here, and stopping when the budget is exhausted or the time horizon is reached. Based on the simulation results that follow, any budget underdraft or regret at the end of the time horizon will be small. More work is needed to make these extensions precise and also to design strategies to minimize the regret and budget underdraft.

3. SIMULATION RESULTS

We verify that this algorithm works well using numerical simulations on synthetic data generated from fixed distributions (unknown to the algorithm). Following [6, 9], we use log-normal distributions to model bids and valuations. Namely, we set our budget to be $B = 0.3$ with time periods of 3 queries each, and generate $v \sim [\ln \mathcal{N}(0,1)]_0^{10}$ and $b' \sim [\ln \mathcal{N}(1,1)]_0^{10}$ independently, where $\ln \mathcal{N}\left(\mu, \sigma^2\right)$ is the log-normal distribution and $[X]_a^b$ indicates the truncated distribution of X on $[a, b]$. These assumptions model a budget-constrained problem where our valuations are less than the maximum bid from the rest of the bidders, in expectation.

If these parameters were known, solving (2) would yield the optimal Lagrange multiplier $\lambda^* \approx 1.18$. One goal of the simulation is then to verify how well our algorithm tracks this Lagrange multiplier without knowledge of the system parameters. We numerically simulate our algorithm starting with $\lambda_0 = 0$, $\lambda_0 = 1$, and $\lambda_0 = 1.5$, with a time-varying stepsize $\epsilon_t = t^{-0.5}$ and a time horizon of $3 \cdot 10^3$, and plot the empirical means based on 10^6 sample-paths in Figure 1. The same initial conditions and parameters, but with $\epsilon_t = t^{-0.8}$ and a time horizon of $3 \cdot 10^4$, is shown in Figure 2. Simulation results for constant ϵ_t and for $\epsilon_t = t^{-1}$ are omitted due to space considerations. In the figures, we plot the empirical behavior of λ_t, the regret $\sum_{i=1}^{N_{t+1}} (v_i - b'_i) \cdot \left(I_{\left\{\frac{v_k}{1+\lambda^*} > b'_k\right\}} - I_{\left\{\frac{v_k}{1+\lambda_k} > b'_k\right\}} \right)$, and the budget underdraft $\sum_{s=1}^{t} \left[B - \sum_{i=N_s+1}^{N_{s+1}} b'_i I_{\{b_i > b'_i\}} \right]$.

For $\epsilon_t = t^{-0.5}$, we see that the choice of λ_0 affects the initial transients of λ_t, and whether $\lambda_0 \gtrless \lambda^*$ determines the initial sign of the underdraft. Furthermore, after λ_t has become close to λ^*, the regret increases very slowly in t. For $\epsilon_t = t^{-0.8}$, we additionally see that the choice of λ_0 can affect the convergence rate of λ_t, with underestimates overshooting and then slowly descending towards λ^* due to the asymmetry in the update equation for λ_t.

In the simulations presented, the regret and budget underdraft are within a few percent of the total profit and budget, respectively, at the end of the time horizon. As would be expected, there is a dependence of the regret and budget underdraft on the initial condition λ_0, with smaller regret when λ_0 is close to λ^*. This suggests that whenever possible, the algorithm should be warm-started with λ_0 set to an approximation of λ^*. As mentioned before, if the exact distributions of v, b', and N_1 are unknown, but one has empirical estimates of these distributions (or perhaps full information for some), one could solve (2) for $\hat{\lambda}^*$ using those surrogate distributions. Additionally, one could also intentionally bias the initial choice λ_0 away from λ^*, possibly to avoid budget overdrafts or to speed up convergence.

4. CONCLUSION

We have considered the problem of auctions with a large number of bidders and an average budget constraint. Through the use of a mean-field approximation, we can formulate the problem as a static optimization problem for each player.

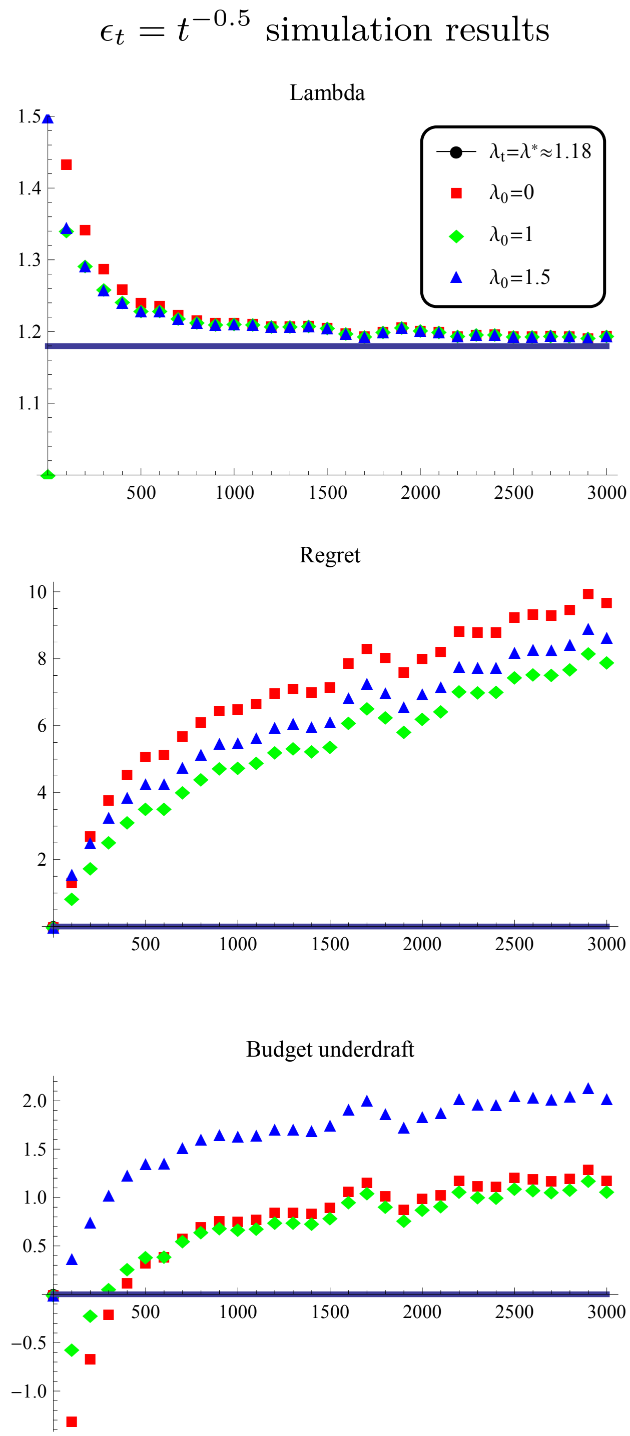

Figure 1: Empirical means of λ_t, regret, and budget underdraft vs. t in the $\epsilon_t = t^{-0.5}$ simulation. For comparison, at the end of the time-horizon, the optimal profit is 871 and the total budget is 300; the regret and budget underdraft are both approximately 1% of these values.

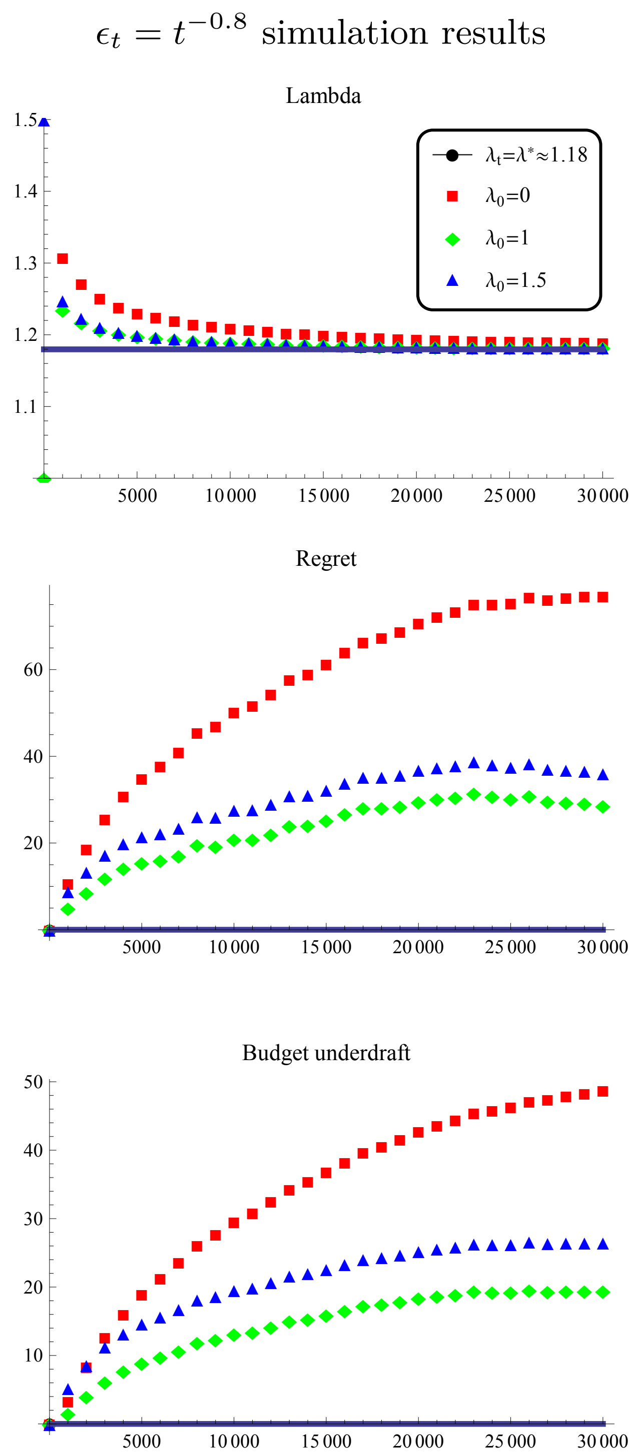

Figure 2: Empirical means of λ_t, regret, and budget underdraft vs. t in the $\epsilon_t = t^{-0.8}$ simulation. For comparison, at the end of the time-horizon, the optimal profit is $8.71 \cdot 10^3$ and the total budget is $3 \cdot 10^3$; the maximum absolute regret and budget underdraft are approximately 1% and 2% of these values, respectively.

One could solve this explicitly given knowledge of the opponents' bid distributions and the valuation distribution, but by applying stochastic approximation, we provide an algorithm that can converge to the optimal bid even without such information. Furthermore, this formulation is significantly simpler to analyze than the dynamic auctions considered previously under a strict budget constraint, and is capable of multiple extensions, including to VCG auctions that incorporate click-through rates. The computations involved in implementing this algorithm are rather minimal, especially if λ_t is updated per time period instead of per query. Our simulation results on synthetic data offer empirical support for bounds on the regret and any budget overdrafts or underdrafts, but of course depending on the assumed distributions, λ_0, choice of ϵ_t, and time horizon T. Warm-starting with an estimate for λ^* based on limited statistical knowledge is also considered. In the longer version of this paper, we will give analytical results on the regret and budget underdraft/overdraft, as well as detail the extensions to multiple slot auctions and the total budget case.

5. ACKNOWLEDGMENTS

Research supported in part by AFOSR MURI FA 9550--10-1-0573.

6. REFERENCES

[1] S. Balseiro, O. Besbes, and G. Weintraub. Auctions for online display advertising exchanges: Approximations and design. *Columbia Business School Research Paper*, 2012.

[2] V.S. Borkar. *Stochastic Approximation: A Dynamical Systems Viewpoint*. Cambridge University Press, 2008.

[3] R. Gummadi, P. B. Key, and A. Proutiere. Optimal bidding strategies in dynamic auctions with budget constraints. In *Proceedings of the 49th Annual Allerton Conference on Communication, Control, and Computing*, pages 588–588, 2011. Longer version available at http://ssrn.com/abstract=2066175.

[4] K. Iyer, R. Johari, and M. Sundararajan. Mean field equilibria of dynamic auctions with learning. In *Proceedings of the 12th ACM conference on Electronic commerce*, EC '11, pages 339–340, 2011. Longer version available at http://ssrn.com/abstract=1799085.

[5] H. Kushner. Stochastic approximation: a survey. *Wiley Interdisciplinary Reviews: Computational Statistics*, 2(1):87–96, 2010.

[6] S. Lahaie and D. M. Pennock. Revenue analysis of a family of ranking rules for keyword auctions. In *Proceedings of the 8th ACM conference on Electronic commerce*, pages 50–56. ACM, 2007.

[7] H. Robbins and S. Monro. A stochastic approximation method. *The Annals of Mathematical Statistics*, pages 400–407, 1951.

[8] B Tan and R Srikant. Online advertisement, optimization and stochastic networks. *IEEE Transactions on Automatic Control*, 57(11):2854–2868, 2012.

[9] D. R. M. Thompson and K. Leyton-Brown. Revenue optimization in the generalized second-price auction. In *Proceedings of the 14th ACM conference on Electronic commerce*, June 2013. To appear.