Optimal Bidding Strategies in Dynamic Auctions with Budget Constraints

R. Gummadi†, P. B. Key‡, A. Proutiere*

***Abstract*—We consider the problem of a bidder with limited budget competing in a series of second-price auctions. A motivating example is that of sponsored search auctions, where advertisers bid in a sequence of repeated generalized second price auctions. To characterize the optimal bidding strategy, we formulate the problem as a discounted Markov Decision Process, and provide explicit solutions when the bidder is involved in a large number of auctions.**

I. Introduction

In on-line advertisement systems, such as sponsored search auction systems, advertisers are repeatedly involved in auctions to acquire advertisement spaces. We analyze the problem faced by a single advertiser with limited budget competing in a series of second-price auctions, and assume that the highest bids of the opponents are independent and identically distributed over different auctions. Under this assumption, we characterize the optimal bidding strategy. The i.i.d. assumption is motivated by recent analysis of real-world data traces [1], and can be theoretically justified when the number of competing bidders grows large [2].

II. A Discrete time Continuous State Markov Decision Model

Time. Time is discrete and indexed by $i = 0, 1, 2, \dots$

Random environment. The pay-offs and payments are determined by the random realizations of the competing bid and the user's own valuation in each time slot. The valuation is observable prior to bidding, but the competing bid is unobservable. This environment is assumed to be i.i.d. over time. Two *generic* random variables (v, w) represent the valuation and competing bid, i.e. they have the same distribution as $v(i), w(i)$ for any i.

Actions. At each time, the bidder selects a bid u from a set U.

Utility. The instantaneous utility from bidding u when the competing bid is w is defined to be $\mathbb{1}_{u>w}(v-w)$. The bidder wishes to maximize the infinite horizon discounted utility, $\sum_{i=0}^{\infty} e^{-\beta i}\mathbb{E}[\mathbb{1}_{u(i)>w(i)}(v(i)-w(i))]$.

Budget Constraint. Initially the bidder's budget is $b(0) = b$. At each time slot, the remaining budget is decreased by the payment and incremented by a fixed amount a. The balance $b(i)$ available at the beginning of slot i evolves as follows:

$$b(i+1) = b(i) + a - \mathbb{1}_{u(i)>w(i)}w(i), \quad b(0) = b. \tag{1}$$

The bidder is forbidden from taking any sequence of actions that lead to a negative balance at any point.

The MDP. Consider a discrete time Markov Decision Process (MDP) on the continuous state space, $\mathbb{R}_+$ representing the budget balance. Let $\mathcal{U}$ represent the collection of all admissible Markov Policies (i.e. $\mathcal{U}$ represents all sequences of Markov bids with the additional restriction that any sequence of bids $\mathbf{u}(i)$ which lead to a strictly positive probability on the event $\mathbf{b}(i+1) < 0$ is forbidden). Then, the value function for an initial balance $\mathbf{b}(0) = b$ is given by:

$$v_\beta(b) = \sup_{\mathcal{U}} \left(\mathbb{E}[\sum_{i=0}^{\infty} e^{-\beta i}\mathbb{1}_{u(i)>w(i)}w(i)] \right) \tag{2}$$

III. Asymptotically optimal bidding

We are interested in analyzing scenarios where the bidder is optimizing over a large number of auctions, and where in each auction, potential payments only consume a very small fraction of the budget. To do so, we let $\beta \to 0$ and rescale value function and budget as follows: $V_\beta(B) \triangleq \beta v_\beta(b) = \beta v_\beta(B/\beta)$, where $v_\beta(b)$ is defined in (2). The following theorem characterizes the limiting value function when $\beta \to 0$.

Theorem 1. *Consider the MDP defined in Section II. Let $\phi : \mathbb{R}_+ \mapsto \mathbb{R}$ be defined as:*

$$\phi(x) = ax + \mathbb{E}\left[(v - b(1+x))^+\right] \tag{3}$$

then ϕ is a convex Lipschitz function with a minimum denoted $\eta_* \triangleq \min_{x\geq 0}\phi(x)$. Let $f : \mathbb{R} \mapsto \mathbb{R}_+$ be an inverse function to ϕ defined as: $f(y) = \min\{x \geq 0 : \phi(x) = y\}$. Then, for all $B \geq 0$, $V(B) = \lim_{\beta\to 0}\beta V_\beta(B)$ is well defined, and satisfies the ODE:

$$\frac{dV}{dB} = f(V), \quad V(0) = \eta_*. \tag{4}$$

When $\beta \approx 0$, the approximately optimal bidding strategy can be written using the theorem as:

$$u^*(v, B) = \frac{v}{1 + V'(B)}. \tag{5}$$

Therefore, when the scaled balance is B, the optimal bid is obtained by shading down the true valuation by a factor equal to $\frac{1}{1+V'(B)}$.

References

[1] F. Pin and P. Key. Stochastic Variability in Sponsored Search Auctions: Observations and Models. *In proc. of ACM Conference on Electronic Commerce*, 2011.

[2] K. Iyer, R. Johari, and M. Sundararajan. Mean field equilibria of dynamic auctions with learning. *In proc. of ACM Conference on Electronic Commerce*, 2011.

†: Stanford Univ., ‡: Microsoft Research, Cambridge, UK *: KTH, Sweden.