Optimal Bidding Strategy for Keyword Auctions and Other Continuous-time Markets

Fang Wu
HP Labs, Palo Alto, CA 94304

February 16, 2007

Abstract

This paper models continuous-time mass bidding markets, such as keyword auctions and market-based resource allocation systems, as a stochastic dynamic system that fluctuates around an average value under the influence of its users. The user's objective to maximize his long-term average utility is formulated as a stochastic control problem. The optimal bidding strategy is calculated both analytically and numerically. It is shown that market fluctuations tend to decrease expected system revenue, thus search engines like Google and Yahoo have an incentive to create a secondary stable market such as a futures market or a reservation market.

1 Introduction

Today, sponsored search keyword auctions, in which sponsors bid for advertising slots displayed alongside the first few search results of one or more keywords, have generated hundreds of billions of revenue for Internet giants like Google and Yahoo, and have consequently stimulated much interest in the research community [2, 6, 7, 18, 20]. Previous work on this topic has focused on two major issues: the design of a well-behaved bidding mechanism for the search company, and the determination of an optimal bidding strategy for an individual sponsor. The most common and natural approach is to model the problem as a one-period auction game and then study its Nash equilibrium [14, 18].

While this approach easily borrows a number of techniques and results from classical auction theory, the final equilibrium is doomed to be static due to the one-period assumption. It cannot fully explain the empirical fact that the bids oftentimes fluctuate around an average value and would not stabilize for a long period. Indeed, modern information techniques have enabled sponsors to monitor and update their bids instantly and effortlessly. Hence, a keyword auction can be viewed as a continuous-time bidding market that, unlike traditional finite-round auctions, is particularly susceptible to fluctuations and dynamic in nature [6, 3, 11, 19].

A second more practical approach formulates the bidding problem as an optimization problem whose various unknown parameters can be estimated from historical data via statistical learning methods [13, 17]. [13] presents an intelligent bidding agent which develops a future look-ahead bidding plan by solving an integer programming problem under a budget constraint. This approach is implementation friendly and truly dynamic, but it lacks a mathematical form at the level of game theory that can be thoroughly analyzed.

To better study the dynamic aspects of keyword auctions while preserving the mathematical elegance, we take a third approach somewhere in between game theory and statistical learning. We formulate the auction market as a continuous-time stochastic dynamical system. The degree of competition present in the market is described by a high-level *market variable* or *competitive index*. The sponsor's influence on the market when he modifies his bid is reflected by a change in the market variable. On one hand, the sponsor may want to increase his bid to achieve a better advertising position. On the other hand, he does not want to overbid too much to saturate the market to an extent that it is no longer worth the price (it is pointless to raise the bid after one has achieved the top position). The problem faced by the sponsor is thus a stochastic control problem in essence: Given the market situation (described by the market variable) and his utility function, how much should he bid? The solution of this problem is calculated in Section 3 using standard methodologies of stochastic optimal control theory.

Our approach not only suits for keyword auctions but is also ideal for other computing markets that permit continuous-time bidding. We closely examine one such market called the *Tycoon* system, which is an established market-based system for managing computer resources in distributed clusters. It is based on a simple proportional share allocation scheme that is convenient for quantitative analysis. In Section 2 we describe its allocation mechanism in detail. In Section 3 we use it as a paradigm of continuous-time bidding markets to demonstrate how to calculate the optimal bidding strategy. We also study how the optimal strategy can be affected by the various market parameters, especially the market volatility. We show that the users tend to bid less when the market fluctuates more around the same mean, so fluctuations hurt the Tycoon managers.

In Section 4 we apply our solution method to keyword auctions. In the end we reach the same conclusion that the sponsors tend to decrease their bids under fluctuations and the search engine has an incentive to open a secondary futures market or reservation market to increase its revenue.

2 The model

2.1 The market variable

In a keyword auction a sponsor's payoff at some time t is completely determined by his own bid and his competitors' bids, so in principle it is not hard for him to determine his best response to the others. If there exists a set of bids in which everyone's bid is the best response to his competitors, it is called a *Nash equilibrium*. A Nash equilibrium is dynamically stable in the sense that when everyone bids the equilibrium, no one has incentive to adjust his bid on his own.

Most of the previous work in the literature focused on solving and studying the Nash equilibria of an auction game. While this approach is conceptually clean and provides important insight into keyword auctions, it has its limitations too. First, these models make *static* assumptions. For example, the value of every advertising position to every sponsor is assumed to be constant. In the real world, however, sponsors enter and leave the sponsored search market frequently, and their budgets keep changing all the time. Their values therefore are not constant. Second, these models produce static solutions. Although a static Nash equilibrium is dynamically stable, it is not clear whether or how the system will reach such an equilibrium. In fact, it has been observed that the bidding prices in a first-price keyword auction may undergo complex oscillations over a very long period [3, 6, 11, 19]. Third and most importantly, in order for a sponsor to play his best response strategy specified by the Nash equilibrium, he has to know the values of the advertising positions to all his opponents, or at least the distributions of these values. This impractical requirement makes the Nash solution hard to realize, especially when the system contains a large number of keywords and sponsors.

In this paper we choose an alternative approach which is particularly efficient for complex bidding markets involving lots of players. Instead of assuming the bidding prices to be endogenous in the equilibrium, we assume that for a fixed individual sponsor under study, his price for acquiring a certain level of utility is exogenous. He needs not attend to the bid of every other player. Instead, he treats all his competitors as one big *market*, and characterizes their collective behavior with one single random variable y, which we call the *market variable* or the *competitive index*. This high level modeling greatly simplifies our problem and is ideal for real-world applications.

We further explain the notion of the market variable with two examples.

2.1.1 Tycoon

Tycoon is a market-based system for managing computer resources in distributed clusters like PlanetLab, the Grid, or a Utility Data Center (UDC) [9]. The basic idea is that users have a limited supply of credits. Consuming users pay providing users to use computer resource. Users who provide resources can, in turn, spend their earnings to use resources later. As of October 2006, the Tycoon grid cluster contains 76 machines in Europe and the U.S., hosting up to 760 virtual machines [1].

The Tycoon system implements a price anticipating resource allocation mechanism in which each user places a bid to each (virtual) machine, and the price of the machine is determined by the total bids placed [9]. Formally, if a user bids x and the other users bid a total of y, then the price is $x + y$, and the user receives a fraction

$$v(x, y) = \frac{x}{x + y} \tag{1}$$

of the total resource. Clearly in Tycoon a user needs not worry too much about the bid of every other user; only the total bid y matters. For this reason we may call y the *market variable*. It measures the degree of competition in the system. Assuming that the user's utility is proportional to the amount of resource he receives, $v(x, y)$ can be regarded as his received *value* when he bids x and the *market* is y. As y increases, the market becomes more and more competitive and the user needs to increase his bid to secure the same utility level.

For each fixed y, $v(x, y)$ can be regarded as a function of x, and hence can be represented as a curve on the $v - x$ plane. We call it a *value curve*. In Fig. 1(a) we plot three value curves, for $y = 1, 2, 3$ respectively. As can be seen there, the user's value curve lowers as the market becomes more competitive (larger y).

2.1.2 Keyword auctions

There are two most widely used auction mechanisms for sponsored search. One is Overture's first-price auction, in which the sponsor pays his bid when a user clicks on his advertising link. The other is Google's *generalized second-price auction*, or *GSP* [7]. In a GSP the sponsors are assigned the best to worst advertising positions in decreasing order of their bids, so the sponsor with the highest bid wins the best position, the second highest sponsor wins the second best position, and so on. When an end user clicks on a sponsor's advertising link, the sponsor pays the search engine the next highest bid, i.e., the bid of the sponsor whose position is one step below him.

Now imagine a new sponsor who has just entered the market. He sees the existing bids of other sponsors, which he sorts in decreasing order: $x_1 \geq x_2 \geq \cdots \geq x_n$. He immediately knows that at this moment he needs to bid no less than x_i to grab position i, for $i = 1, \ldots, n$, and by doing so he commits to pay $x_i + \epsilon$ or x_i every time a user clicks on his link, for first-price and second-price auction respectively. Suppose his value v_i for position i is proportional to the click-through-rate of that position. We can write $v(x_i) = v_i$, where $v(x)$ is the value he receives by bidding x.

Fig. 1(b) shows two specific keyword auctions hosted by Overture, one for the keyword "printer" and one for the keyword "cartridge", collected from `uv.bidtool.overture.com` on September 29, 2006 at 22:15pm. When plotting the data we rather arbitrarily pick the value (click-through-rate) of position i to be

$$v_i = 1 - \frac{i-1}{40} \quad \text{for } i = 1, \ldots, 40, \tag{2}$$

so that the first position is normalized to have rate 1. One position down decreases the rate by 1/40, and the positions below 40 are never clicked on.[1]

As can be seen from the figure, the value-bid data points for each keyword follow a concave pattern similar to the value curves of Tycoon. Hence we may

[1] Reference [5] reports that the click-through-rate decays in position neither linearly nor exponentially. Nevertheless, the main point of this section is to justify the market variable rather than exact modeling. To this end we still assume linear decreasing for simplicity.

(a)

(b)

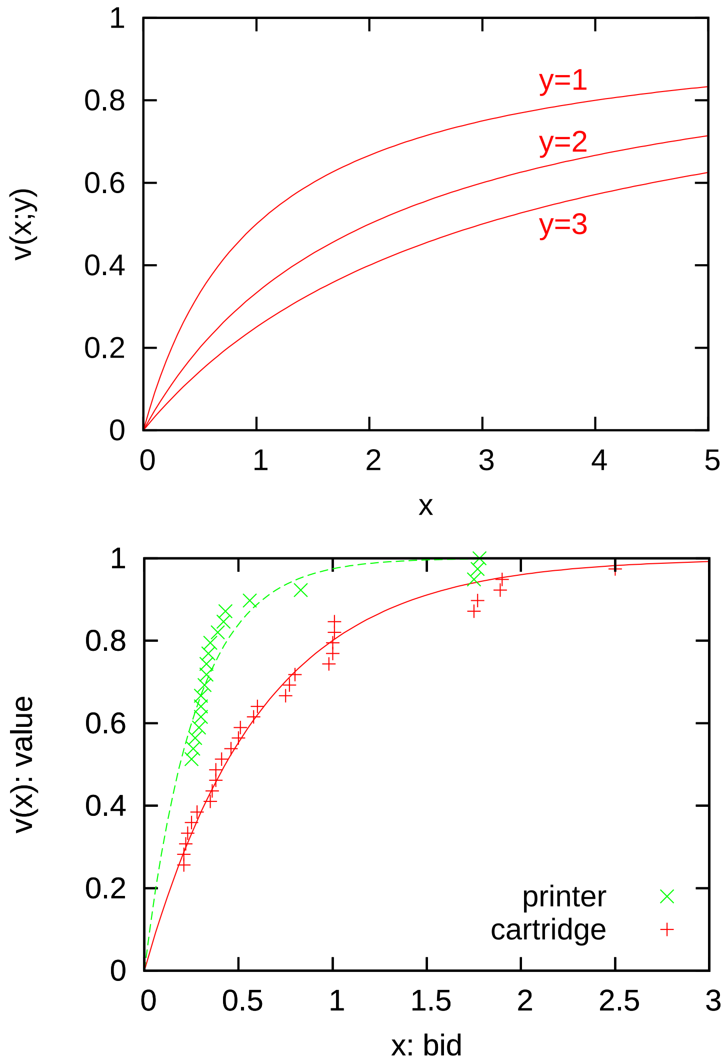

Figure 1: Value curves of (a) Tycoon and (b) Overture keyword auction. The reason why “cartridge” ($y = 0.600$) is more competitive than “printer” ($y = 0.267$) is because the printer companies extract most of their profit from ink cartridges rather than from printers. In fact, in our collected data, the No. 1 printer company HP bids \$10.00 for the keyword “cartridge” (beyond the figure border) but only \$0.56 for “printer”!

try to fit the data with a family of concave value curves:

$$v(x; y) = 1 - e^{-\frac{x}{y}}, \tag{3}$$

where x is the variable and y is the parameter. Like in Tycoon, every y uniquely determines a value curve, and the larger y, the lower is the curve. The least-square fit results are plotted in Fig. 1(b) as two curves: $y = 0.267$ for "printer" and $y = 0.600$ for "cartridge". They serve as good approximations to the real data. In other words, although the sponsor's received value really depends on $n + 1$ variables, namely his own bid x and other players' bids $x_1, \ldots, x_n$, we can approximate it with a two variable function $v(x; y)$. Thus, like in Tycoon, we may call y the *market variable* or *competitive index* for keyword auctions, because it describes the collective behavior of all players and measures the competitiveness of the market.

In summary, this paper assumes the existence of a market variable y, which, together with the user's bid x, uniquely determines his value.

2.2 The value function

We could have used a two-parameter function, e.g. $1-e^{-\frac{x-y_1}{y_2}}$, to fit the keyword auctions data in the previous section. Then we would have two market variables y_1 and y_2 instead of one, which would give an even better approximation to the market. However, the main innovation of this paper lies in the stochastic control approach other than modeling, so we want to keep our model as simple as possible and stay with one market variable, although our later analysis could easily extend to more than one variable.

The choice $v(x, y) = 1-e^{-\frac{x}{y}}$ for keyword auctions is of course arbitrary. We could have chosen another functional form, say the Tycoon one, as long as it satisfies the following properties:

1. $v(x, y)$ is increasing in x and decreasing in y.
2. For any fixed y, $v(x, y) \to 0$ as $x \to 0$, and $v(x, y) \to 1$ as $x \to \infty$.
3. For any fixed x, $v(x, y) \to 1$ as $y \to 0$, and $v(x, y) \to 0$ as $y \to \infty$.

Property 1 says the bidder gets more when he bids more, and gets less when the market becomes more competitive. Property 2 and 3 reflect the finiteness nature of the bidding object, so the bidder's value eventually saturates. (In Tycoon he cannot get more than the whole machine, and in keyword auctions he cannot get higher than the first position.)

In Section 4 we will use the function $v(x, y) = e^{-\frac{y}{x}}$.

2.3 The utility function

A bidder will not participate in an auction unless by placing a bid he can gain some (positive) utility, which in general can be written as a function of his value and cost. In Tycoon a user's cost is just his bid x. In first-price auctions, a

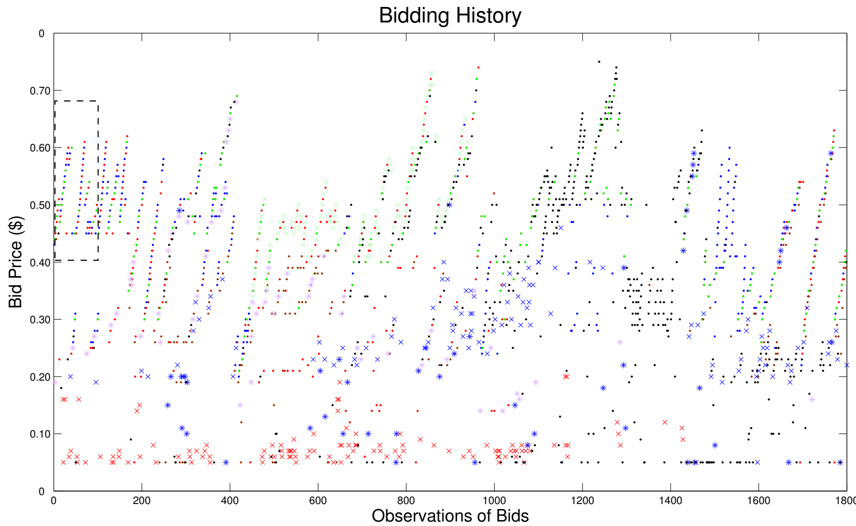

Figure 2: This figure is adapted from [19]. Nine most active bidders are represented as colored symbols, while the rest bids are represented as black dots.

bidder pays x for every click, so his cost in unit time is vx. In GSP however, a sponsor pays the next highest bid, so his cost can be lower than vx. But he can always choose his bid just slightly higher than the next highest bid so that the two bids are essentially equal, in which case his cost is again vx. In all cases, the bidder's utility is determined by v and x and can be written as $u(v, x)$.

When the market is static, i.e. when the market variable y is fixed, he can only operate on a fixed value curve $v(x; y)$, which implies that his utility $u = u(v(x, y), x)$ is uniquely determined by his bid. Therefore his optimal bidding strategy is given by any $x^* \in \arg\min_x u(v(x, y), x)$.

The presupposition that the market is static, however, is very much open to question. [3, 6, 11, 19] have reported cycles in both first and second-priced auctions. Bidders change their bids all the time to outbid their competitors and to save cost. Also, once a new bidder enters the market, he affects the market immediately by making it more competitive. In the next section we discuss how to characterize a dynamic market.

2.4 The market dynamics

Fig. 2 displays the complete bidding history of one keyword in the year 2001 for Yahoo keyword auctions [19]. As can be observed, the bidding price fluctuated around some average value for the whole year and never settled down. This dynamical behavior is typical for first-price auctions [3]. Unfortunately, Google does not publicize its bidding data so we cannot visualize a second-price auction like Fig. 2. But there is some empirical evidence that second-price auctions also fluctuate [6].

Based on the empirical data, we assume that our market variable follows a *mean reverting process* [4], in the sense that it fluctuates randomly around some average value. Our particular choice is the standard geometric Ornstein-Uhlenbeck (Dixit-Pindyck) process [8, 16]:

$$\frac{dy}{y} = \kappa(c - y)\, dt + \sigma\, dw. \tag{4}$$

Here time t is continuous, a setting most natural for markets like Tycoon and sponsored search auctions in which bidders can update their bids at any time. w is the standard Wiener process (Brownian motion). κ is the mean reversion rate and c is the mean reversion level. When y falls below c the first mean reversion term pushes y upward, and when y exceeds c the mean reversion term drags y downward. σ is the volatility that measures the degree of fluctuation. It is assumed that $\kappa, c, \sigma > 0$ and $\kappa c > \sigma^2/2 \equiv s$.

Now we consider what happens after a new bidder enters the market. As soon as he places a bid the market becomes more competitive, so the competitive index y will tend to increase. We make the simplest assumption that the increase is proportional to the new bid, so the market dynamics with this entering effect appears like

$$dy = \kappa y(c + ax - y)\, dt + \sigma y\, dw. \tag{5}$$

Eq. (5) is our complete market dynamics.[2] The model parameters can be estimated from historical data in principle. Simple as it looks, it captures mean reversion, random fluctuation and market competition all in one equation.

2.5 The bidder's problem

We assume throughout this paper that the bidder has a quasi-linear utility density in the form $u = \text{value} - \beta \cdot \text{cost}$, where $\beta > 0$ is his utility density for spending one dollar. For Tycoon this amounts to

$$u(x, y) = \frac{x}{x + y} - \beta x. \tag{6}$$

For keyword auctions we have

$$u(x, y) = v - \beta v x = v(x, y)(1 - \beta x), \tag{7}$$

where $v(x, y) \geq 0$ is the click-through-rate of his position, as before. We note that in both cases $u < 0$ when $x > \beta^{-1} \equiv b$.

The bidder's total expected utility over a period of T is therefore

$$U(T) = \mathbb{E}\left[\int_0^T u(x, y) dt\right]. \tag{8}$$

[2] It is shown in the Appendix that $y > 0$ with probability one, so y can never be negative.

A bidder usually stays in the bidding market for a long time. Suppose he seeks to maximize his long-term average utility. His objective can then be formulated as follows:

$$\max_{x(y)\in X} \lim_{T\to\infty} \frac{1}{T}\,\mathbb{E}\left[\int_0^T u(x,y)dt\right], \tag{9}$$

where X is his strategy space containing all continuous functions $x(y) : \mathbb{R}_+ \to \mathbb{R}_+$. Note that the bidder's strategy x should only depend on the market situation y and not on t, because of our infinite horizon formulation.

Alternatively, we could have formulated the bidder's problem as follows. Suppose he wants to maximize his long-term total *value* defined by

$$\max_{x(y)\in X} \mathbb{E}\left[\int_0^T v(x,y)dt\right], \tag{10}$$

subject to a budget constraint

$$\mathbb{E}\left[\int_0^T x(y)dt\right] = B \quad \text{for Tycoon,} \tag{11}$$

or

$$\mathbb{E}\left[\int_0^T v(x,y)x(y)dt\right] = B \quad \text{for keyword auctions.} \tag{12}$$

We can transform this constrained problem into an unconstrained problem by introducing a Lagrange multiplier (shadow price) β. This is actually equivalent to the previous quasi-linear utility formulation. For this reason we may rightly call $b \equiv \beta^{-1}$ the *budget factor*. The larger budget the bidder has, the less he cares about money and the bigger is b.

2.6 Our model is not a game

Our high-level modeling of the market with one descriptive variable y suits only for mass markets. In small markets the deviation of any single user's strategy may have a big impact on the market itself and be sufficient to invalidate Eq. (5). In such cases it makes more sense to formulate the problem as a multi-player game and solve for its equilibria.

Contrary to the traditional game theory approach, our approach focuses on just one user and does not attempt to solve any equilibrium. Nor do we attempt to answer questions like, if all user adopt the same optimal strategy (to be given in Section 3), whether they can bootstrap the collective behavior described by Eq. (5). Any convincing answer to such questions requires constructing and solving a continuous-time stochastic game, a topic which we leave out due to its technical difficulty.

3 The solution for Tycoon

We recapitulate the Tycoon user's objective:

$$\max_{x(y)\in X} \lim_{T\to\infty} \frac{1}{T}\,\mathbb{E}\left[\int_0^T u(x,y)dt\right], \tag{13}$$

where

$$u(x,y) = \frac{x}{x+y} - \beta x, \tag{14}$$

$$dy = \kappa y(c + ax - y)\,dt + \sigma y\,dw. \tag{15}$$

3.1 Deterministic markets

We first solve the bidder's problem (13) for deterministic markets ($\sigma = 0$), whose dynamics is governed by

$$\dot{y} = \kappa y(c + ax - y). \tag{16}$$

Suppose the bidder has found an optimal strategy $x^*(y)$. Because $u(x,y) < 0$ for $x > b$, he should keep his bid under b. Consider the function $f(y) = c + ax^*(y) - y$. We see that $f(0) \geq 0$ and $f(c+ab) \leq 0$. Because f is continuous, there must exist some $y^* \geq 0$ such that $f(y^*) = 0$, not necessarily unique. The one-dimensional dynamical system (16) will eventually converge toward such a y^*, and $x^*(y)$ will converge to the number $x^*(y^*)$.

From the above argument we see that, because the bidder optimizes his long-term average utility, the function values of $x^*(y)$ at points other than y^* do not really matter. He could always place the same *constant* bid $x^*(y) = x^*$ neglecting the market, and the market itself would be led to $y^* = c + ax^*$. Of course, he should choose his bid x^* in such a way that $u(x^*, y^*)$ is maximized.

Specifically, for Tycoon we have

$$u(x^*, y^*) = \frac{x^*}{x^* + y^*} - \beta x^* = \frac{x^*}{(a+1)x^* + c} - \beta x^*. \tag{17}$$

The maximum utility is achieved at

$$x^* = \frac{1}{a+1}\left(\sqrt{\frac{c}{\beta}} - c\right)^+, \tag{18}$$

where "+" denotes positive part. The bidder will only bid when his budget is large enough ($b > c$), and his optimal long-term average utility is

$$u^* = \frac{(1 - \sqrt{\beta c})^2}{a+1} \quad (\beta c < 1). \tag{19}$$

3.2 The steady state

Now let us inspect stochastic markets ($\sigma > 0$). Suppose the bidder uses a bounded continuous strategy $x(y)$. Instead of converging to a fixed number, it can be shown that when the market follows the stochastic dynamics specified by Eq. (5) the market variable y will converge to a *steady-state distribution*, whose density function is given by

$$\pi(y) = \frac{m}{\sigma^2 y^2} \exp\left[\int_{y_0}^{y} \frac{\kappa(c + ax(z) - z)}{sz} dz\right], \tag{20}$$

where m is some positive normalization factor. The bidder's optimization problem can then be written equivalently as

$$\max_{x(y)\in X} \int_0^\infty \pi(y)u(x(y), y)dy. \tag{21}$$

3.3 Mass markets

There is a special limit case of stochastic markets which bears a closed-form solution. When the market consists of a large population of bidders one individual's impact can be neglected. This corresponds to the $a = 0$ limit, in which the bidder only needs to bid the price that maximizes his instant utility $u(x, y)$. From Eq. (6) we find

$$x^*(y) = \left(\sqrt{\frac{y}{\beta}} - y\right)^+, \tag{22}$$

and

$$u(x^*, y) = (1 - \sqrt{\beta y})^2 \quad \text{for } y < b. \tag{23}$$

Note how Eq. (22) and (23) are connected with Eq. (18) and (19) in the $a = 0$ limit. Furthermore, the steady-state distribution, which can be derived from Eq. (20), turns out to be a Gamma distribution

$$\pi(y) = \frac{A^B}{\Gamma(B)} e^{-Ay} y^{B-1}, \tag{24}$$

where

$$A = \frac{\kappa}{s}, \quad B = \frac{\kappa c}{s} - 1, \tag{25}$$

with mean and variance

$$\mu = c - \frac{s}{\kappa}, \quad \Sigma^2 = \left(c - \frac{s}{\kappa}\right)\frac{s}{\kappa}. \tag{26}$$

The optimal long-term average utility and spending under this distribution can then be calculated.

3.4 The general solution

We now solve the optimal strategy for the general case. We start by relaxing the normalization factor m in Eq. (20) to rewrite the bidder's problem as a constrained one:

$$\max_{x(y)\in X} \int_0^\infty \pi(y)\left[\frac{x(y)}{x(y)+y} - \beta x(y)\right] dy \quad \text{s.t.} \quad \int_0^\infty \pi(y)dy = 1, \tag{27}$$

where the constraint is just the normalization condition. Observe that when $y > b = \beta^{-1}$, it holds that

$$u(x,y) = \left(\frac{1}{x+y} - \beta\right) x \le \left(\frac{1}{y} - \beta\right) x \le 0. \tag{28}$$

Hence the bidder should never bid when $y > b$, which will only cause him to lose utility and waste time in the more competitive region (placing a bid increases the drift). Pick y_0 in Eq. (20) to be $y_0 = \inf\{y : x(z) = 0 \text{ for all } z \ge y\}$. Then $y_0 \le b < \infty$. By means of Lagrange multiplier, we now need to solve

$$\max_{x(y)} \left\{ \int_0^{y_0} \pi(y)\left[\frac{x(y)}{x(y)+y} - \beta x(y)\right] dy + \lambda\left[1 - \int_0^\infty \pi(y)dy\right]\right\}, \tag{29}$$

where we have replaced the upper endpoint of the first integral with y_0, since any optimal strategy $x(y)$ vanishes on $[y_0, \infty)$.

We can solve Eq. (29) using calculus of variations. The variation of $\pi(y)$ is given by

$$\delta\pi(y) = \pi(y) \int_{y_0}^{y} \frac{\kappa a\, \delta x(z)}{sz} dz. \tag{30}$$

The total variation of Eq. (29) can be calculated as follows:

$$\begin{aligned}
\delta U &= \int_0^{y_0} \left\{ \delta\pi\left[\frac{x}{x+y} - \beta x\right] + \pi\left[\frac{y}{(x+y)^2} - \beta\right]\delta x \right\} dy - \lambda \int_0^\infty \delta\pi dy \\
&= \int_0^{y_0} \left\{ \delta\pi\left[\frac{x}{x+y} - \beta x - \lambda\right] + \pi\left[\frac{y}{(x+y)^2} - \beta\right]\delta x \right\} dy \\
&= \int_0^{y_0} \pi\left[\frac{x}{x+y} - \beta x - \lambda\right] \int_{y_0}^{y} \frac{\kappa a\, \delta x(z)}{sz} dz\, dy + \int_0^{y_0} \pi\left[\frac{y}{(x+y)^2} - \beta\right]\delta x\, dy \\
&= \int_0^{y_0} \left\{ -\frac{\kappa a}{sy} \int_0^{y} \pi(z)\left[\frac{x(z)}{x(z)+z} - \beta x(z) - \lambda\right] dz \right. \\
&\qquad \left. + \pi(y)\left[\frac{y}{(x(y)+y)^2} - \beta\right]\right\} \delta x(y)\, dy,
\end{aligned} \tag{31}$$

where in the second "=" we have replaced $\int_0^\infty \delta\pi dy$ with $\int_0^{y_0} \delta\pi dy$. This is valid because the fact that $x(y) = 0$ on $[y_0, \infty)$ and Eq. (30) imply $\delta x(y) = \delta\pi(y) = 0$ on $[y_0, \infty)$. The first-order condition for the optimal $x^*(y)$ reads

$$-\kappa a \int_0^y \pi(z)\left[\frac{x^*(z)}{x^*(z)+z} - \beta x^*(z) - \lambda\right] dz + s\pi(y)\left[\frac{y^2}{(x^*(y)+y)^2} - \beta y\right] = 0. \tag{32}$$

By introducing

$$r(y) = \frac{y}{x^*(y) + y}, \tag{33}$$

it can also be written as

$$-\kappa a \int_0^y \pi(z)[1 - r(z) - \beta x^*(z) - \lambda]\, dz + s\pi(y)[r(y)^2 - \beta y] = 0. \tag{34}$$

Taking derivatives with respect to y, we find

$$-\kappa a(1 - r - \beta x^* - \lambda) + s\frac{\dot{\pi}}{\pi}(r^2 - \beta y) + s(2r\dot{r} - \beta) = 0. \tag{35}$$

Taking log-derivatives with respect to y in Eq. (20), we have

$$s\frac{\dot{\pi}}{\pi} = \frac{\kappa c - 2s + \kappa a x^*}{y} - \kappa. \tag{36}$$

Plugging Eq. (36) into Eq. (35), we obtain

$$2s\kappa^{-1}r\dot{r} + \left(\frac{c - 2s\kappa^{-1}}{y} - 1\right) r^2 + \frac{ax^*}{y} r^2 + \beta y + ar - a + a\lambda - \beta(c - s\kappa^{-1}) = 0. \tag{37}$$

Using $rx^* = (1-r)y$ we finally obtain a first-order nonlinear differential equation for $r(y)$:

$$2s\kappa^{-1}r\dot{r} + \left(\frac{c - 2s\kappa^{-1}}{y} - 1\right) r^2 - a(1-r)^2 + \beta y + a\lambda - \beta(c - s\kappa^{-1}) = 0. \tag{38}$$

Letting $y \to 0$ we find the asymptotics $r(y) \sim \sqrt{y}$ near $y = 0$, so the boundary condition is $r(0) = 0$.

We still need an equation to determine the value of λ though. The first-order condition of Eq. (29) with respect to m yields

$$\lambda \int_0^\infty \pi(y) dy = \int_0^{y_0} \pi(y) \left[\frac{x^*(y)}{x^*(y) + y} - \beta x^*(y)\right] dy. \tag{39}$$

In probabilistic language this is

$$\lambda = \mathbb{E}\left[\frac{x^*(y)}{x^*(y) + y} - \beta x^*(y)\right] = \mathbb{E}\, U(x^*(y), y), \tag{40}$$

where the expectation of y is taken over the steady-state distribution. Thus we see that λ can be interpreted as the long-term average utility under the optimal bidding strategy.

Eq. (38) and (39), together with the boundary condition $r(0) = 0$ constitute a complete set of equations for $r(y)$ and therefore $x^*(y)$.

Remark. The optimal strategy can also be solved using a dynamic programming approach [15]. Let $J(y,t,T)$ be the maximum expected utility the bidder can gain during the time interval (t,T) when the market variable starts from y at time t. It must satisfy the Hamilton-Jacobi-Bellman equation

$$0 = \max_x \left\{ u(x,y) + \frac{\partial J}{\partial t} + \frac{\partial J}{\partial y}\kappa y(c+ax-by) + \frac{\partial^2 J}{\partial y^2} sy^2 \right\}. \tag{41}$$

In the infinite horizon limit we must have [15]

$$\lim_{T\to\infty} \frac{\partial J}{\partial t} = \text{constant}, \tag{42}$$

and the optimal strategy can be solved by representing $\partial J/\partial y$ and $\partial^2 J/\partial y^2$ in terms of u, x and y. More details can be found in [15].

3.5 The numerical solution

Eq. (38) in general cannot be solved analytically, so in this section we resort to numerical solutions. Provided with the input of parameters $(\kappa, a, c, \sigma, \beta)$, Algorithm 1 describes a simple heuristic that solves the optimal $x^*(y)$ and the corresponding λ. As an example, we plot the numerical solution for three sets of parameters in Fig. 3.

Algorithm 1 Finding the optimal strategy for Tycoon

Step 1. Pick any $\lambda \in (0,1)$. Numerically solve Eq. (38) with the boundary condition $r(0) = 0$ until $r(y) = 0$ for some $y = y_0$.

Step 2. Calculate $\pi(y)$ from Eq. (20) and set

$$\lambda' = \frac{\displaystyle\int_0^{y_0} \pi(y) \left[\frac{x(y)}{x(y)+y} - \beta x(y) \right] dy}{\displaystyle\int_0^{\infty} \pi(y) dy}. \tag{43}$$

Step 3. Set $\lambda \leftarrow \lambda'$ and go to Step 1, unless $|\lambda - \lambda'| < \epsilon$ for some pre-defined $\epsilon > 0$.

As can be seen from the figure, when y is small there is very little competition so the bidder can bid a low price to receive a large fraction of resource. When y exceeds a certain threshold the market eventually becomes so competitive that the user cannot afford to offer any bid given his budget.

3.6 Comparative statics

We study the effect of the parameters a and s through comparative statics analysis. In Table 1(a) we fix the parameters κ, c, β, s and increase a. The

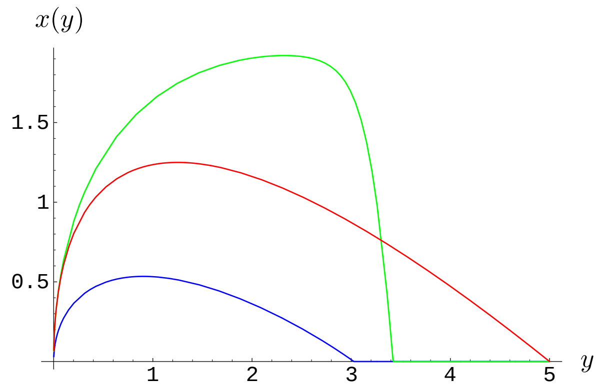

Figure 3: Optimal bidding strategy for Tycoon as a function of the competitive index. The chosen parameters (κ, a, c, s, β) are $(1, 0, 1, .2, .2)$ for the red line, $(1, .2, 1, .1, .1)$ for the green line, and $(1, .5, .8, .5, .25)$ for the blue line. The long-term average utility λ for the three lines are 0.306, 0.417 and 0.478, respectively.

resulting optimal long-term average utility λ and the optimal expected spending $\mathbb{E}\, x(y)$ are listed. We see that increasing a decreases the bidder's utility and spending. This is because markets with bigger a counteract on the bidder more strongly, so he should *decrease* his bid to alleviate competition.

a	λ	$\mathbb{E}\, x$
0	.363	.391
.1	.342	.356
.2	.324	.327
.3	.308	.301
.4	.293	.278
.5	.280	.259

(a) $\kappa = 1$, $c = 1$, $\beta = .5$, $s = .5$.

s	λ	$\mathbb{E}\, x$
0	.208	.417
.1	.233	.396
.2	.257	.376
.3	.280	.357
.4	.303	.341
.5	.324	.327

(b) $\kappa = 1$, $a = .2$, $\beta = .5$, $\mu = .5$.

Table 1: Comparative statics of a and s.

Next we study the effect of σ. Instead of fixing c, we fix the mean market competitive level in the absence of the bidder, i.e. we fix $\mu = c - s/\kappa$ (see Eq. (26)).[3] We also fix the values of κ, a, β. The effect of an increasing s is listed in Table 1(b). As can be seen, when the market becomes more volatile, the bidder *pays less to get more*!

This phenomenon can be intuitively explained as follows. Say the market variable fluctuates around a mean value 1 with standard deviation 0.8. Without the fluctuations the bidder has to compete with $y = 1$ for sure. In the presence of fluctuations, however, there is some chance that y will get to as low as 0.2,

[3] The same trends in λ and $\mathbb{E}\, x$ are observed when c is fixed in lieu of μ.

an event which the bidder can exploit in full. Of course the market could also rise to as high as 1.8, but then the bidder can choose to stay away. On average, *he rides the market low to gain utility.*

4 The solution for keyword auctions

The optimal bidding strategy for keyword auctions can be derived similarly to Tycoon. Instead of embroiling ourselves in intricate differential equations, in this section we will assume a special form of value function and solve for a special class of suboptimal strategies, which would be sufficient for us to carry out some basic comparative statics analysis to achieve a better understanding of the sponsored search market.

Specifically, we assume that the sponsor has a value function in the form $v(x,y) = e^{-\frac{y}{x}}$. It satisfies the three properties listed in Section 2.2, and is mathematically tractable.

We now ask the question: what *constant* bid x^* is optimal for the bidder? Of course the class of constant strategies is a very limited subset of the complete strategic space, but the constant x^* would still give us a rough idea about the bidder's optimal average spending, plus in the real world those bidders who do not bother to monitor the market continuously do use such lazy strategies.

When the bidder bids x^* constantly, the market variable converges to the steady-state distribution

$$\pi(y) = \frac{A^B}{\Gamma(B)} e^{-Ay} y^{B-1}, \tag{44}$$

where

$$A = \frac{\kappa}{s}, \quad B = \frac{\kappa(c + ax^*)}{s} - 1. \tag{45}$$

His long-term average utility is

$$\mathbb{E}\,U = \int_0^\infty \frac{A^B}{\Gamma(B)} e^{-Ay} y^{B-1} e^{-\frac{y}{x^*}} (1 - \beta x^*) dy = (1 - \beta x^*) \left(\frac{Ax^*}{Ax^* + 1} \right)^B. \tag{46}$$

x^* must satisfy the first-order condition. Taking log-derivatives with respect to x^*, we have

$$-\frac{1}{b - x^*} + \frac{B}{x^*} - \frac{B}{x^* + s'} + \frac{a}{s'} \log \left(\frac{x^*}{x^* + s'} \right) = 0, \tag{47}$$

where $s' = s/\kappa$. In the large market limit ($a = 0$) this simplifies to

$$-\frac{1}{b - x^*} + \frac{c - s'}{x^*(x^* + s')} = 0, \tag{48}$$

and its solution is given by

$$x^* = \frac{\sqrt{c^2 + 4(c - s')b} - c}{2} = \frac{\sqrt{(\mu + s')^2 + 4\mu b} - (\mu + s')}{2}. \tag{49}$$

We see that x^* decreases when we hold μ and increase s'. In words, the bidder tends to spend less when the market fluctuates more around the same mean. This is consistent with the comparative statics analysis in Section 3.6.

To put it in another way, *search engines like Google and Yahoo loses revenue due to fluctuations in the spot market.* Thus, they have an incentive to create a secondary stable market for sponsors to bid in, such as a futures market or a reservation market, in which the sponsors bid a price and stick to it for a certain period of time [10]. There can be many ways for Google to maintain two markets side by side. For example, its website can be designed in such a way that when a user searches for a set of keywords, with probability p the sponsor links are displayed according to the spot market bidding result, and with probability $1 - p$ they are displayed according to the reservation market bidding result. This way Google would increase its revenue.

5 Conclusion

This paper formulates a continuous-time mass bidding market as a stochastic dynamic system that fluctuates around an average value. Such markets include sponsored search keyword auctions and market-based resource allocation systems. We have calculated for both cases the optimal bidding strategy which maximizes the user's long-term average utility. By playing this optimal strategy the market converges to a steady-state distribution. We have also calculated the expected system revenue under the limit distribution and showed that market fluctuations tend to decrease the system's expected revenue.

The few parameters of our model can be measured from historical data. On the user side, our optimal solution serves as a practical strategic guide. On the system manager side, this paper offers a theoretical ground for search engines like Google and Yahoo to open up secondary stable markets such as a futures market or a reservation market.

I thank Jason Hartline, Bernardo A. Huberman, Li Zhang and Yunhong Zhou for helpful comments and suggestions.

References

[1] `http://tycoon.hpl.hp.com/cp/monitor/`.

[2] G. Aggarwal, A. Goel, and R. Motwani. Truthful auctions for pricing search keywords. In *ACM EC'06*, 2006.

[3] K. Asdemir. Bidding patterns in search engine auctions. 2006.

[4] C. Blanco and D. Soronow. Mean reverting processes — energy price processes used for derivatives pricing & risk management. *Commodities Now*, June 2001.

[5] N. Brooks. The Atlas rank report: How search engine rank impacts traffic, July 2004.

[6] B. Edelman and M. Ostrovsky. Strategic bidder behavior in sponsored search auctions. In *ACM EC'05*, 2005.

[7] B. Edelman, M. Ostrovsky, and M. Schwartz. Internet advertising and the generalized second prize auction: Selling billions of dollars worth of keywords. Forthcoming in *American Economic Review*, 2006.

[8] D. Epstein, N. Mayor, P. Schonbucher, A. E. Whalley, and P. Wilmott. The valuation of a firm advertising optimally. *Quarterly Review of Economics and Finance*, 38(2):149–166, 1998.

[9] M. Feldman, K. Lai, and L. Zhang. A price anticipating resource allocation mechanism for distributed shared clusters. In *ACM EC'05*, 2005.

[10] J. Feng, Z.-J. M. Shen, and R. L. Zhan. Ranked items auctions and online advertisement. Forthcoming in *Production and Operations Management*.

[11] J. Feng and X. Zhang. Price cycles in online advertising auctions. In *Proceedings of the 26th International Conference on Information Systems (ICIS)*, December 2005.

[12] I. Karatzas and S. E. Shreve. *Brownian motion and stochastic calculus*. Springer, 2nd edition, 2000.

[13] B. Kitts and B. Leblanc. Optimal bidding on keyword auctions. *Electronic Markets*, 14(3):186–201, 2004.

[14] S. Lahaie. An analysis of alternative slot auction designs for sponsored search. In *ACM EC'06*, 2006.

[15] R. C. Merton. An asymptotic theory of growth under uncertainty. *Review of Economic Studies*, 42:375–93, 1975.

[16] G. E. Metcalf and A. Hasset. Investment under alternative return assumptions comparing random walks and mean reversion. *Journal of Economic Dynamics and Control*, 19:1471–1488, 1995.

[17] T. Payne, E. David, N. R. Jennings, and M. Sharifi. Auction mechanisms for efficient advertisement selection on public displays. In *Proceedings of 17th European Conference on Artificial Intelligence (ECAI-06)*, 2006.

[18] H. R. Varian. Position auctions. Forthcoming in *International Journal of Industrial Organization*.

[19] X. Zhang. Finding Edgeworth cycles in online advertising auctions. 2006.

[20] Y. Zhou and R. Lukose. Vindictive bidding in keyword auctions. In *Second Workshop on Sponsored Search Auctions*, 2006.

Appendix

In this appendix we show that if y follows the dynamics

$$dy = \kappa y(c + ax(y) - y)dt + \sigma y dw, \tag{50}$$

then it has a steady state distribution over $(0, \infty)$. Here it is assumed that $0 \leq x(y) < b$.

Define the *scale function*

$$p(y) = \int_{y_0}^{y} \exp\left\{-\int_{y_0}^{z} \frac{\kappa z'(c + ax(z') - z')}{sz'^2} dz'\right\} dz \tag{51}$$

and the *speed measure*

$$m(dy) = \frac{1}{sy^2} \exp\left\{\int_{y_0}^{y} \frac{\kappa z(c + ax(z) - z)}{sz^2} dz\right\} dy. \tag{52}$$

It suffices to show that

$$p(0) = -\infty, \quad p(\infty) = \infty, \quad m((0, \infty)) < \infty. \tag{53}$$

(See e.g. [12] Section 5.5, Exercise 5.40, pp. 352.) These conditions can be checked as follows:

$$p(0) \leq -\int_0^{y_0} \left(\frac{z}{y_0}\right)^{-\frac{\kappa c}{s}} e^{-\frac{\kappa y_0}{s}} dz = -\infty, \tag{54}$$

$$p(\infty) \geq \int_{y_0}^{\infty} \left(\frac{z}{y_0}\right)^{-\frac{\kappa(c+b)}{s}} e^{\frac{\kappa(z-y_0)}{s}} dz = +\infty, \tag{55}$$

$$m((0, y_0)) \leq \int_0^{y_0} \frac{y^{\frac{\kappa c}{s}-2}}{sy_0^{\frac{\kappa c}{s}}} e^{\frac{\kappa y_0}{s}} dy < \infty, \tag{56}$$

$$m((y_0, \infty)) \leq \int_{y_0}^{\infty} \frac{1}{sy^2} \left(\frac{y}{y_0}\right)^{\frac{\kappa(c+b)}{s}} e^{-\frac{\kappa(y-y_0)}{s}} dy < \infty. \tag{57}$$