MANAGEMENT SCIENCE
Vol. 57, No. 9, September 2011, pp. 1703–1719
ISSN 0025-1909 | EISSN 1526-5501 | 11 | 5709 | 1703

http://dx.doi.org/10.1287/mnsc.1110.1408

Cyclical Bid Adjustments in Search-Engine Advertising

Xiaoquan (Michael) Zhang
Department of Information Systems, Business Statistics and Operations Management, Hong Kong University of Science and Technology, Clear Water Bay, Kowloon, Hong Kong, zhang@ust.hk

Juan Feng
Department of Information Systems, College of Business, City University of Hong Kong, Kowloon, Hong Kong, juafeng@cityu.edu.hk

Keyword advertising, or sponsored search, is one of the most successful advertising models on the Internet. One distinctive feature of keyword auctions is that they enable advertisers to adjust their bids and rankings dynamically, and the payoffs are realized in real time. We capture this unique feature with a dynamic model and identify an equilibrium bidding strategy. We find that under certain conditions, advertisers may engage in cyclical bid adjustments, and equilibrium bidding prices may follow a cyclical pattern: price-escalating phases interrupted by price-collapsing phases, similar to an "Edgeworth cycle" in the context of dynamic price competitions. Such cyclical bidding patterns can take place in both first- and second-price auctions. We obtain two data sets containing detailed bidding records of all advertisers for a sample of keywords in two leading search engines. Our empirical framework, based on a Markov switching regression model, suggests the existence of such cyclical bidding strategies. The cyclical bid-updating behavior we find cannot be easily explained with static models. This paper emphasizes the importance of adopting a dynamic perspective in studying equilibrium outcomes of keyword auctions.

Key words: bid adjustment; Edgeworth cycle; keyword auction
History: Received February 26, 2010; accepted June 1, 2011 by Pradeep Chintagunta and Preyas Desai, special issue editors. Published online in *Articles in Advance* August 4, 2011.

1. Introduction

Keyword advertising, or sponsored search, is one of the most successful advertising models on the Internet. One unique feature of keyword auctions is that they enable advertisers to adjust their bids and rankings dynamically, and the payoffs are realized in real time. We study bid adjustments in keyword advertising and demonstrate dynamic interactions among advertisers.

We argue that one of the most distinctive features of search advertising is precisely the dynamic nature of interactions among advertisers. Compared to traditional auctions, keyword auctions possess unique characteristics. First, bidders are not required to be present at the same place at the same time. This nonsynchronicity in time naturally leads to bidders' learning and dynamic interactions. Second, unlike bids in a typical open auction, which can move in only one direction, bids in keyword auctions can be either higher or lower than previous bids. As a result, the strategy space does not shrink with the submission of new bids. Third, keyword auctions do not close. The payoffs are realized in real time while the auctions are still in progress. These properties enable advertisers to evaluate positions, examine competitors, and adjust their bids in real time. The bidders' capability to learn and react makes these keyword auctions markedly different from traditional auctions and creates a need to model them differently. For example, in traditional second-price auctions, it is well known that a bidder does not want to deviate from her true value by raising her bid; however, in second-price keyword auctions, a bidder can adopt a strategy called "bid jamming" such that her bid is just one cent below her competitor's.[1] Bid jamming allows a bidder to increase her competitor's cost at almost no cost for herself. Without dynamic learning, it is impossible to implement this strategy. This study suggests that bid adjustments are an important feature of keyword auctions and highlights the importance of using dynamic models to examine this market.

Figure 1 shows some bid-adjustment patterns in two popular search engines (Yahoo.com and Baidu.com) that adopt first- and second-price auction mechanisms, respectively.[2] These patterns foreshadow our core findings and suggest that (1) cyclical

[1] This strategy is commonly suggested by search-engine marketing experts and is widely used in practice (Ganchev et al. 2007, Stokes 2010).

[2] A detailed description of the data is provided in §3. The two figures are based on data extracted from 15 days from Yahoo! and 30 days from Baidu.

Figure 1 Close-ups of the Bidding History

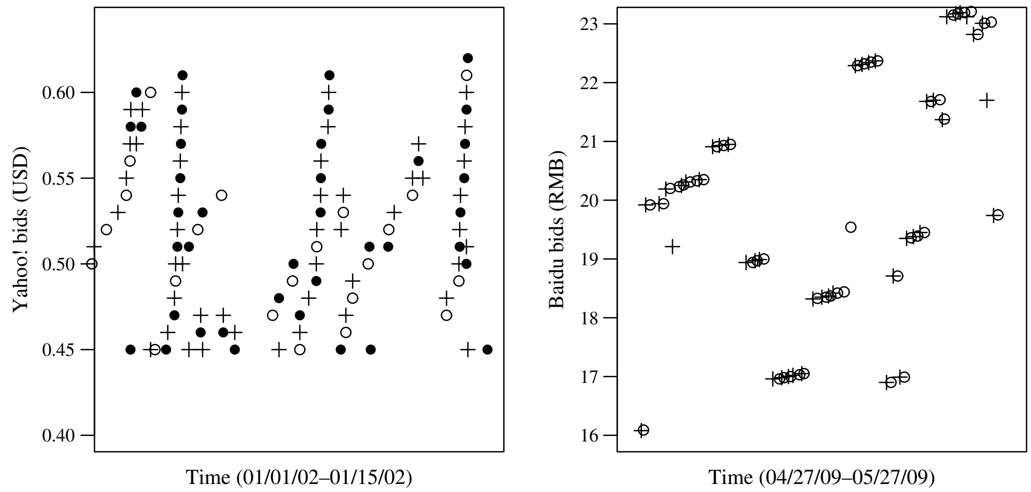

bid adjustments take place frequently, and (2) static models cannot describe these dynamics properly.

In this paper, we first establish a theoretical framework of search advertising and focus on dynamic reactions among bidders. Our model suggests that under certain equilibrium conditions, for both first- and second-price auctions, bids can display cyclical patterns and may not converge to a static equilibrium. Using data directly obtained from search engines, we demonstrate the existence of cyclical patterns in both first- and second-price auctions.

Keyword auction-based search-engine advertising attracts much academic interest across the disciplines of marketing (Goldfarb and Tucker 2011; Rutz and Bucklin 2007, 2008; Wilbur and Zhu 2009; Chen et al. 2009b; Feng 2008; Katona and Sarvary 2010; Yang and Ghose 2009), information systems (Ghose and Yang 2009; Feng et al. 2007; Chen et al. 2009a; Animesh et al. 2010, 2011; Liu and Chen 2006; Liu et al. 2010), economics (Varian 2007, Edelman et al. 2007, Arnold et al. 2011), and computer science (e.g., Ganchev et al. 2007). Edelman et al. (2007) and Varian (2007) establish theoretical models of the market and find that prices and allocations of positions are equivalent to those obtained from a static dominant-strategy Nash equilibrium of a multi-item Vickrey auction. Borgers et al. (2007) suggest that the models of Edelman et al. (2007) and Varian (2007) focus on only a small subset of multiple possible equilibria. Most existing studies have treated keyword auctions as if they were in a static environment; these static models can leave important equilibrium outcomes unnoticed. Although Cary et al. (2008) find that dynamic responses in their model may lead to a static equilibrium under specific conditions, they also suggest that researchers need to rigorously examine how results in static equilibria can be generalized to dynamic games. Different from prior models of static games, we examine equilibrium bidding in a dynamic setting. Whereas Zhang and Feng (2005) and Feng and Zhang (2007) consider price cycles only in first-price keyword auctions, this paper presents theoretical and empirical results for both first- and second-price auctions.

This paper is also related to price wars, an everlasting theme in the literature of dynamic price competition. Collusion, demand and cost shocks, and information asymmetry resulting from noisy signals or detection lags (Stigler 1964; Green and Porter 1984; Rotemberg and Saloner 1986; Porter 1983; Lee and Porter 1984; Ellison 1994; Athey et al. 2004; Athey and Bagwell 2001, 2004, 2008) can all induce price-war patterns. In reality, firms react to each other's actions continually, especially in industries with only a small number of firms.

One possible outcome predicted by our model, the cyclical pattern of bid adjustments, is similar to Edgeworth cycles in the literature of duopoly price competition. In these studies, two firms with identical marginal costs undercut each other's price in an alternating manner. Edgeworth (1925) proposes this theory in criticizing the standard "Bertrand competition" result that both competitors set their prices equal to their marginal costs. Maskin and Tirole (1988a, b) formally characterize Edgeworth cycles with the concept of Markov perfect equilibrium (MPE) in a class of sequential-move duopoly models. Building on Maskin and Tirole's model, Noel (2007a, b) identifies Edgeworth cycles in retail gasoline prices. Our paper extends these studies to model heterogeneous players competing for multiple, heterogeneous objects. In our model, bidders have different private valuations for clicks, and different ad positions also offer different values. The structure of the optimal strategy in this study is thus different from that of the symmetric-player framework that other researchers examine.

We contribute to the literature in a few ways. First, whereas most models in the search-advertising literature are static, we theoretically demonstrate the existence of cyclical bidding patterns in a dynamic framework. Our results therefore emphasize the importance of developing dynamic models to better understand bidding strategies in keyword auctions. Second, few previous studies use field data directly obtained from search engines. Our empirical analysis offers a quantitative assessment of how bidders behave in the real world. Finally, this paper is the first to document Edgeworth-like cycles in a setting unrelated to gasoline prices.

This paper proceeds as follows. Section 2 develops the general theoretical framework and examines the existence of cyclical bid adjustments. In §3, we describe two data sets and present the empirical analysis. Section 4 discusses the implications of our findings, and §5 concludes.

2. The Model

2.1. Setup

Our notation and model setup closely follow Maskin and Tirole's (1988b) framework. Noel (2007a, b) adopts a similar theoretical framework in his model of retail gasoline prices. Consider two risk-neutral advertisers competing for two advertising positions for a keyword.[3] Denote θ_i as advertiser i's exogenous valuation for a click directed from the ad, which measures the profit that an advertiser expects to obtain from a click-through.

Assume that each advertising position has a positive inherent expected click-through rate (CTR) $\tau_j \in (0, 1)$, which is solely determined by the position, j, of the advertisement. More specifically, the higher an advertisement's position, the more click-throughs it attracts; thus $\tau_1 \geq \tau_2$.[4] Accordingly, an advertiser i's expected revenue from the winning position j can be represented by $\theta_i \tau_j$.

In the auction of advertising positions, advertisers submit their bids for the amount that they are willing to pay for one click. The ad with the higher bid is displayed at the top position, and the other ad is displayed at the second position. Major search engines often rank ads by a quality score, in addition to the bids. We extend our model to accommodate this practice in §4.1. The payment scheme for each position can be either first price (that is, advertisers pay their own bids), or second price (that is, the advertiser at the top position pays the bid of the second position, and the other advertiser pays the reserve price). The competition for advertising positions can last arbitrarily long. To capture the dynamic feature of bid adjustments, we focus on a game that takes place in discrete time in infinitely many periods with a discount factor $\delta \in (0, 1)$. Similar to Maskin and Tirole (1988b), we assume that advertisers adjust their bids in an alternating manner; that is, in each period t, only one advertiser is allowed to update her bid, and in the next period, only the other bidder is allowed to update.[5]

Previous studies used the "locally envy-free" assumption to rule out the possibility of dynamic interactions in bids.[6] The assumption helps reduce the model to a one-shot game. Because the objective of this paper is to study dynamic interactions, we can no longer adopt this assumption. Instead, we focus on a dynamic-equilibrium approach to study keyword auctions.

2.2. Main Model

Following (Maskin and Tirole 1988a, p. 553), we assume "recent actions have a stronger bearing on current and future payoffs than those of the more distant past." In each period, a bidder's strategy depends only on the variables that directly enter her payoff function (that is, the bid set by the other bidder in the last period and the resulting allocation of ad positions). We focus on the perfect equilibrium of this game, which means, starting from any period, the advertiser who is about to act selects the bid that maximizes her intertemporal profit, given the subsequent optimal strategies of the other advertiser and of her own. We are interested in stationary properties; therefore, initial conditions are irrelevant.

We make the following assumptions:

1. The inherent CTR (τ_j) for each position j is exogenously given and remains constant across all periods.

[3] In the theoretical model, we do not discuss the case of more than two bidders or the case when the number of bidders is different from the number of slots. As Maskin and Tirole (1988a, p. 551) argue, "at the cost of simplicity, considering more players yields no additional insights." In our context, this means that bidding wars among more than two bidders can often be broken into segments of bidding wars between just two bidders. We can see in Figure 1 that advertisers are often paired together by their bid adjustments.

[4] This is a common assumption in the literature that many industry reports confirm (Breese et al. 1998). This reflects the fact that ads in higher positions attract more attention.

[5] This assumption is necessary to establish our dynamic, sequential-move equilibrium. As Maskin and Tirole (1988a, p. 549) argue, "The fact that, once it has moved, a firm cannot move again for two periods implies a degree of commitment."

[6] Edelman et al. (2007, p. 243) wrote, "By the folk theorem, however, such a game will have an extremely large set of equilibria, and so we focus instead on the *one-shot, simultaneous-move, complete information stage* [italics added by authors] game, introducing restrictions on advertisers' behavior suggested by the market's dynamic structure. We call the equilibria satisfying these restrictions 'locally envy-free.' "

2. Both bidders are willing to pay at least r, the minimum reserve price, to participate in the auction; that is, $\theta_i - r > 0, \forall i$.

3. The bidding space is discrete; that is, firms cannot set prices in units smaller than, for example, one cent. Let ϵ denote this smallest unit of valuation. This implies that all values and bids discussed in this paper are multiples of ϵ.

4. Both advertisers discount future profits with the same discount factor δ.

Denote b_i as a focal bidder's bid, and let b_{-i} represent the other bidder's previous bid. Define $R = (R_1, R_2)$ as a dynamic equilibrium strategy profile, which represents the dynamic reaction functions forming the perfect equilibrium. Strategy profile R is an equilibrium if and only if, for each bidder i, $b_i = R_i(b_{-i})$ maximizes the bidder's intertemporal profit at any time, given both bidders bid according to (R_i, R_{-i}) thereafter. Bidder i's single-period payoff $\pi_i(b_i; b_{-i})$, when the current competing bid is b_{-i}, is determined by her own bid b_i, as well as the position she wins. Following industry practice, we assume that if the bids are the same, the bidders will split the chance to share the top position. We can then write $\pi_i(b_i; b_{-i})$ in a first-price auction as

$$\pi_i(b_i; b_{-i}) = \begin{cases} (\theta_i - b_i)\tau_1 & \text{if } b_i > b_{-i}, \\ \frac{1}{2}(\theta_i - b_i)\tau_1 + \frac{1}{2}(\theta_i - b_i)\tau_2 & \text{if } b_i = b_{-i}, \\ (\theta_i - b_i)\tau_2 & \text{if } b_i < b_{-i}. \end{cases} \quad (1)$$

Similarly, $\pi_i(b_i; b_{-i})$ in a second-price auction can be written as

$$\pi_i(b_i; b_{-i}) = \begin{cases} (\theta_i - b_{-i})\tau_1 & \text{if } b_i > b_{-i}, \\ \frac{1}{2}(\theta_i - b_{-i})\tau_1 + \frac{1}{2}(\theta_i - r)\tau_2 & \text{if } b_i = b_{-i}, \\ (\theta_i - r)\tau_2 & \text{if } b_i < b_{-i}. \end{cases} \quad (2)$$

Define a pair of value functions, $V_1()$ and $W_1()$, for advertiser 1 as below (the other advertiser's valuation functions can be defined in the same way). Let

$$V_1(b_2) = \max_{b_1}[\pi_1(b_1; b_2) + \delta W_1(b_1)] \quad (3)$$

represent bidder 1's expected payoff if (a) she is about to move, (b) the other bidder just bid b_2 in the previous period, and (c) both bidders play according to $R = (R_1, R_2)$ thereafter. $W_1(b_1)$ is defined as

$$W_1(b_1) = E_{b_2}[\pi_1(b_1; b_2) + \delta V_1(b_2)], \quad (4)$$

which represents bidder 1's valuation if (a) she played b_1 in the previous period and (b) both bidders play optimally according to $R = (R_2, R_2)$ thereafter. Thus, $R = (R_1, R_2)$ is an equilibrium if $R_1(b_2) = b_1$ is the solution to Equation (3), the expectation in Equation (4) is taken with respect to the distribution of b_2, and the symmetric conditions hold for advertiser 2. Without loss of generality, assume $\theta_1 > \theta_2$. We can propose the following equilibrium strategy for first-price auctions:

Proposition 1 (Equilibrium Bidding Strategy (First Price)). *In a first-price auction, for heterogeneous advertisers ($\theta_1 - \theta_2 > \epsilon$) and a sufficiently fine grid (i.e., ϵ is sufficiently small compared to θ_i, $i = 1, 2$), there exist threshold values $\underline{b}_1$ and $\bar{b}_2$, such that each bidder's equilibrium strategy can be specified as*

$$\begin{aligned} b_1 = R_1(b_2) &= \begin{cases} b_2 + \epsilon & \text{if } b_2 < \underline{b}_1, \\ \bar{b}_2 & \text{otherwise}; \end{cases} \\ b_2 = R_2(b_1) &= \begin{cases} b_1 + \epsilon & \text{if } b_1 < \bar{b}_2, \\ r & \text{otherwise}, \end{cases} \end{aligned} \quad (5)$$

where r is the reserve price, and ϵ is the smallest increment.

Proof. Proofs for Propositions 1 and 2 are in the appendix. Other proofs are in the online appendix (http://blog.mikezhang.com/files/bidadjustments_appendix.pdf). □

Proposition 1 specifies an MPE as in Maskin and Tirole (1988b) and describes bidders' equilibrium bidding strategies characterized by $\underline{b}_1$ and $\bar{b}_2$, which are functions of θ_i, δ, and r. Proposition 1 shows that sequential bid adjustments can be supported as an equilibrium strategy. More specifically, bid patterns can be grouped into two phases: a price-escalating phase and a price-collapsing phase. Beginning with the reserve price r, bidders will wage a price war (outbidding each other by ϵ until price $\underline{b}_1$ is reached), then bidder 1 (the higher-valued bidder) will jump to bid $\bar{b}_2$, which is the highest bid that bidder 2 can afford to get the top slot. Now bidder 2 can no longer afford the costly competition and will be forced to remain at the second position. She can choose any price lower than $\bar{b}_2$ and keep the same position, but she is strictly better off by choosing to bid the reserve price r. When she does so, she switches from the escalating state to the collapsing state. Consequently, bidder 1 should follow this drop to bid $r + \epsilon$, at which price she remains at the first position and pays much less. Then a new round of the price war begins, and so on. This process is structurally similar to the price wars, or Edgeworth cycles, presented by Maskin and Tirole (1988b) and Noel (2007a, b). The driver behind the cycles in Proposition 1 is that bidders can outbid each other in each period to obtain more click-throughs while only paying slightly more than before. In Edgeworth cycles in a duopolistic price competition, the driver of the cycles is that firms undercut each other by a little bit, obtaining a higher market share through reducing the price only slightly.

Proposition 1 gives a general characterization when the bidders' per-click values are different. When the bidders' valuations are identical, their equilibrium strategy needs to be modified slightly:

Corollary 1. *When $\theta_1 = \theta_2$, with a sufficiently fine grid (small ϵ), there exist $\underline{b}$, $\bar{b}$, and probability $\mu \in [0, 1]$, such that bidder i's equilibrium strategy is*

$$b_i = \begin{cases} b_{-i} + \epsilon & \text{if } r \le b_{-i} < \underline{b}, \\ \bar{b} & \text{if } \underline{b} \le b_{-i} < \bar{b}, \\ \left.\begin{matrix} \bar{b} + \epsilon & \text{with probability } \mu \\ r & \text{with probability } 1-\mu \end{matrix}\right\} & \text{if } b_{-i} = \bar{b}, \\ r & \text{if } b_{-i} > \bar{b}. \end{cases} \tag{6}$$

This reduces to the Edgeworth cycle described by Maskin and Tirole (1988b, Proposition 7). The equilibrium strategy reported there can be understood as a special case of Proposition 1.[7] The major difference between the equilibrium strategy played by homogeneous bidders (as in Corollary 1) and that played by heterogeneous bidders (as in Proposition 1) is that, with heterogeneous bidders, on the equilibrium path, it is the higher-valued bidder (bidder 1) who jumps up to end the "escalating" phase, and it is always the lower-valued bidder (bidder 2) who first drops to bid r and restarts the cycle. In the case of homogeneous bidders, when the price reaches $\bar{b}$, it is equally costly for both bidders to bid higher. Whoever drops the price down first, however, does a favor to the other bidder, because the other bidder then can occupy the top position for two rounds.

The cycles we identify in Proposition 1 differ from Edgeworth cycles in the literature in a number of ways. First of all, we study price cycles in a different context (auction/bidding versus oligopolistic price competition). So the price cycles in our setting exhibit a reversed pattern (outbidding to raise prices in our setting versus undercutting to reduce prices in Edgeworth cycles). Second, and more importantly, because we relax the assumption of homogeneity of players, bidders' equilibrium bidding strategies display unique properties. In the first-price case, mixed strategy at the point of state switching is avoided with heterogeneous bidders. In the second-price case that we present next, the strategies of the two bidders are no longer symmetric.

We next turn to discuss the equilibrium of second-price auctions. Existing literature on "generalized second-price auctions" has predicted that there exists a locally "envy-free" equilibrium, in which advertisers' ranks, as well as bids, are constant in the equilibrium (Edelman et al. 2007, Varian 2007). The envy-free equilibrium concept is built upon the assumption that there exists a resting point such that "a player cannot improve his payoff by exchanging bids with the player ranked one position above him" (Edelman et al. 2007, p. 249). This assumption simplifies the model to a one-shot game and misses one important aspect of bidders' considerations: Although bid increases by the lower-valued bidder may not change the ranks or her own payoff, they could introduce a higher cost for the higher-valued bidder. For example, consider a situation in which advertiser 1 bids \$5 and advertiser 2 bids \$3. Even though an "envy-free" advertiser 2 may be happy to keep the bid unchanged at \$3, a "jealous" advertiser strictly prefers to bid \$4 to \$3: By increasing the bid by \$1, advertiser 2's cost remains unchanged at the reserve price, but advertiser 1's cost is increased by \$1. This observation is especially conspicuous in practice when advertisers have daily or weekly advertising budgets.[8] Even without budget constraints, ceteris paribus, a strategy that raises competitor's cost should be considered as strictly preferable to one that does not. We therefore explicitly assume that among the strategies that maximize their own discounted utilities, advertisers strictly prefer one that imposes the highest costs on their competitors. This assumption is not only intuitive and empirically valid, but also useful in eliminating many possible but unrealistic deviations from the equilibrium. The following proposition examines the bidding strategy in second-price auctions.

Proposition 2 (Equilibrium Bidding Strategy (Second Price)). *In a second-price auction, for heterogeneous bidders who prefer strategies that, ceteris paribus, impose the highest costs on competitors, there exists an equilibrium bidding strategy for a sufficiently fine grid (i.e., ϵ is sufficiently small). Specifically, when $\epsilon < \theta_1 - \theta_2 < (\tau_1/(\tau_1 - \tau_2))((3\delta - 1)/(1 - \delta))\epsilon$, define*

$$b_i = \begin{cases} b_{-i} + \epsilon & (\textit{Strategy I}), \\ b_{-i} - \epsilon & (\textit{Strategy II}); \end{cases} \tag{7}$$

then there exists an equilibrium with

$$b_2 = \begin{cases} b_1 - \epsilon & \text{if } b_1 \ge \underline{b}_2' \text{ or when bidder 1 plays Strategy II}, \\ b_1 + \epsilon & \text{if otherwise}; \end{cases}$$
$$b_1 = \begin{cases} b_2 + \epsilon & \text{if bidder 2 plays Strategy I}, \\ r & \text{if bidder 2 plays Strategy II}. \end{cases} \tag{8}$$

[7] In their setting, firms are homogeneous, and $\tau_2 = 0$ (because the firm with a higher price will get zero demand).

[8] If the lower-valued bidder can successfully push her competitor to a higher price level, she could enjoy the top position with a very low price once the higher-valued advertiser runs out of budget.

Proposition 2 indicates that the equilibrium bid-updating pattern under the second-price scheme can exhibit two phases, as in the first-price case. Starting from the lowest possible price r, bidders engage in a price war in which each advertiser outbids the other by the minimum increment until the price reaches $\underline{b}_2'$. At this price, the lower-valued advertiser (advertiser 2) can no longer afford to bid higher, and she would choose to stay at the second slot. To induce the highest cost for bidder 1, she chooses to bid just ϵ below bidder 1. If bidder 1's value is not significantly higher than that of bidder 2 (i.e., when $\theta_1 - \theta_2 < (\tau_1/(\tau_1-\tau_2))((3\delta-1)/(1-\delta))\epsilon$), bidder 1 would drop the bid to the reserve price and start the cycle again. It is important to note that although there is some similarity in the equilibrium strategies described in Propositions 1 and 2, the equilibrium concept in Proposition 2 is not an MPE. As a result, one caveat is that Proposition 2 is not as "robust" as the MPE discussed in Maskin and Tirole (1988b).

The threshold condition determines two possible scenarios. In the first scenario, when the difference in valuations is large, the higher-valued bidder can dominate the keyword by submitting a relatively high bid, b_1^*. The lower-valued bidder then adopts the bid-jamming strategy by submitting $b_1^* - \epsilon$. The high level of b_1^* is the the premium that bidder 1 pays to maintain the control of the top position. In the second scenario, if the difference in valuations is not large enough, then our model shows the existence of cyclical pattern in equilibrium. To illustrate, consider two bidders. If one bidder always bids r and takes the second position, her click-throughs will always be lower than her competitor's. Now if she bids a little higher (e.g., $r+2\epsilon$), she can easily take the first position. This creates the momentum for the cycle to start. Their competition will reach to the point that the lower-valued bidder finds it not profitable to raise the bid any more. The lower-valued bidder will keep the price to be ϵ lower than the bid of the higher-valued bidder. At this point, the bidder at the lower position only pays the reserve price, which is close to zero, but the bidder at the higher position has to pay a much higher price. Because the difference in valuation is not large enough, the higher-valued bidder finds it more profitable dropping her bid to the reserve price than maintaining the top position.

There is an important similarity in the above two scenarios: the higher-valued bidder has to pay a premium in equilibrium. In the stable-equilibrium scenario, the higher-valued bidder pays a much higher price to stay in the first position. In the cyclical-equilibrium scenario, the higher-valued bidder would not drop to r earlier in the cycle because she has to pay a premium to induce higher cost for the lower-valued bidder.

Figure 2 **Markov Chain Associated with Bidding War**

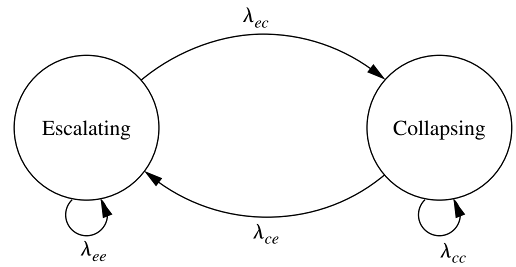

Figure 2 gives the Markov chain representation of Equations (5) and (8). Beginning from any bid, the bidders' dynamic bid adjustments will create a price pattern. Each bid is in one of two states: e (escalating) or c (collapsing). Escalating bids will beat the competitor's bid, and collapsing bids will drop the price to a lower level. The transition probabilities are indicated with λ_{ij}, with $i, j \in \{e, c\}$.

Games with an infinite horizon often have a multiplicity of equilibria, and thus the two equilibria presented in our propositions may not be unique. Given our focus on demonstrating the existence of dynamic equilibrium outcomes, we refrain from discussing other possible static or dynamic outcomes. We want to emphasize that our results do not rule out the existence of other types of equilibria.

2.3. Bounds of Cyclical Bidding

Propositions 1 and 2 describe two equilibrium bidding strategies when the advertising slots are auctioned off using different pricing schemes. Although bidders' bidding strategies are different in these two auction mechanisms, the bounds of cyclical bidding in these two mechanisms share some common characteristics. Under both schemes, the bounds of cyclical bidding can be associated with (1) the gap between advertisers' valuations and (2) the differentiation between advertising slots.

Proposition 3. *Assuming $\theta_1 \geq \theta_2$, cyclical bidding is less likely to be sustained when*

1. *the difference between the per-click valuations of bidders become more significant (for given θ_1), and*
2. *the differentiation between the slots reduces (for given τ_1).*

Proposition 3 emphasizes that the cyclical equilibrium outcomes of bid adjustment exist only under specific conditions. Both advertiser- and visitor-specific characteristics can affect the sustainability of the cycles. The intuition behind this proposition is straightforward. First, when a bidder's valuation for a keyword is sufficiently higher than her competitor's, she will simply submit a very high bid and dominate the first position, without giving a chance for the lower-valued bidder to reach the first position. The

Table 1 **Summary Statistics**

Variable	Mean	Std. dev.	Median	Minimum	Maximum
Panel A: First-price auction (Yahoo!)					
Bid price (USD)	0.40	0.34	0.39	0.05	4.97
Time in market (days)	120.67	137.26	59.00	0.71	364.74
Number of bids submitted by firms (#)	36.70	60.87	9	1	245
Time between consecutive bids (hours)	4.77	8.65	1.63	1/60	88.20
Time between consecutive bids submitted by the same player (days)	2.33	6.73	0.71	$1/(24*60)$	98.08
Panel B: Second-price auction (Baidu)					
Bid price (RMB)	18.60	5.46	20.00	0.72	44.80
Time in market (days)	68.08	17.20	72.87	35.34	83.21
Number of bids submitted by firms (#)	50.91	93.44	9	1	265
Time between consecutive bids (hours)	3.60	6.56	1.00	1/60	44.20
Time between consecutive bids submitted by the same player (days)	1.31	5.63	0.15	$3/(24*60)$	79.89

higher her valuation is, the easier it is for her to find it profitable to submit a dominating high bid. One sufficient condition for this to happen is $(\theta_1-\theta_2)\tau_1 \geq (\theta_1-r)\tau_2$. Second, the ranking becomes less important when the two positions are equally attractive. In this case, both bidders can simply bid the reserve price r, and there will be no price war.

Based on comparative statics, Proposition 3 describes conditions for price cycles to exist. For a search engine, greater heterogeneity in bidders' valuations may lead to a stable equilibrium, with one bidder always paying a high price. Similarity of CTRs in adjacent positions can also lead to a stable equilibrium, but because of the lack of competition, both bidders stay close to the reserve price.

In summary, a profit-maximizing search engine should (1) differentiate advertising slots so that they generate significantly different click-throughs and (2) facilitate bid-updating behavior and try to reduce bidders' costs of changing their bids. This is consistent with the observation that search engines often (1) differentiate the ad positions on the top from those on the right-hand side and (2) offer bid research tools to advertisers.

3. Empirical Analysis

3.1. Data

To identify the existence of cyclical bidding behavior, we obtained data from two sources. The data set for first-price auction were obtained from Yahoo!. Before the Yahoo! acquisition, Overture adopted the first-price mechanism. Our data set therefore covers the period when Overture was operated independently. The data set for second-price auction were obtained from Baidu, the largest search engine in China.

Yahoo!'s data set contains year 2002's complete bidding history for one keyword, which for confidentiality reasons, we do not know. Each time an advertiser submitted a new bid, the system would record the bidder's ID, the date and time, and the bid value. A total of 1,800 bid adjustments have been recorded for this keyword. Forty-nine bidders submitted at least one bid during this period. Among these, seven bidders submitted more than 100 bids. Summary statistics are given in panel A of Table 1. *Time in market* is calculated by counting the number of days between the first and last observations of bid adjustments. *Time between consecutive bids* is measured by the number of hours between bid changes. *Time between consecutive bids submitted by the same player* measures the number of days between consecutive adjustments made by the same bidder.

Baidu was founded in 2000. It is now listed on the NASDAQ (symbol: BIDU) and serves more than 77% of Internet searches in China. Google and Baidu together account for more than 95% of China's search advertising market. Baidu has gained a greater market share because of a recent dispute between Google and the Chinese government. According to official numbers, Baidu can reach 95% of Internet users in China. It now indexes more than 10 billion webpages.[9] Similar to Google's AdWords and AdSense, Baidu offers search advertising and affiliate advertising. Advertisers on the search advertising platform compete through a second-price mechanism similar to that of Google's. A bidder's rank is determined by her own bid, the next-highest bid, and a quality score depending on the ad's historical performance. The second-price auction data from Baidu contain click events. Each time a visitor clicks on an ad, there is a record. The disadvantage of this data set is that we cannot observe all bid adjustments. If an advertiser makes several bid adjustments between two click events, we can only observe the latest change. Although this brings some inconvenience in data processing, the data set still satisfies our need to test for cyclical bidding patterns. First, an inability to observe

[9] See http://home.baidu.com/product/product.html (accessed July 2010).

some of the bid changes biases against our finding bid adjustments; our finding of bid adjustments would be strengthened had we been able to observe all bid changes. Second, this concern is alleviated for keywords with frequent clicks. Intuitively, an advertiser would not want to change bids more frequently than changes in clicks, especially if there is even a very small cost associated with bid adjustments. Third, the data set reports the next-highest bid for all click events. This frees us from the need to retrieve the history of all bids. The records we have span the duration from April 27 to July 19, 2009. In panel B of Table 1, we report the summary statistics for the keyword from Baidu's second-price auction.

For both data sets, bidders are very active in making adjustments. The median time between consecutive bids is about 1.6 hours on Yahoo! and 1 hour on Baidu.

Figures 3 and 4 show the complete bidding history of the Yahoo! and Baidu keywords, respectively. Each symbol represents a bidder. From these figures, bid adjustments can be clearly visualized. The lower parts of these figures show the estimated state probabilities calculated with our ensuing empirical model.

3.2. A Markov Switching Regression Model

We adopt the empirical strategy of the Markov switching regression to examine the patterns. The Markov switching regression was proposed by Goldfeld and Quandt (1973) to characterize changes in the parameters of an autoregressive process. From the cyclical trajectory of the bids shown in the figures, it is very tempting for us to assign one of two states (escalating state (e) or collapsing state (c)) to each observation of the bids and directly estimate the parameters with a discrete-choice model. Using the Markov switching regression gives us a few advantages. First, the result of the Markov switching regression matches nicely with the theoretical Markov chain depicted in Figure 2. Second, serial correlation is incorporated into the model. The parameters of an autoregression are viewed as the outcome of a two-state, first-order Markov process. Third, when the price trajectory is not as regular as that displayed by the data, the Markov switching regression can help identify the latent states. This eliminates the need for subjectively assigning dummy-variable values for the states, thus making it less dependent on subjective human judgment. Finally, from the estimation process, we can easily derive the Markov transition matrix, and the parameter estimates can be interpreted directly.

Formally, consider a two-state, ergodic Markov chain shown in Figure 2, with state space $\mathcal{S}(S_t) = \{e, c\}$, where $s_t = e$ represents the escalating phase, $s_t = c$ represents the collapsing phase, and $t = 1, \ldots, T$. These states are latent because we do not directly observe them. The process $\{S_t\}$ is a Markov chain with the stationary transition probability matrix $\Lambda = (\lambda_{ij})$, where

$$\lambda_{ij} = \mathrm{Prob}(s_t = j \mid s_{t-1} = i), \quad i, j \in \{e, c\}. \tag{9}$$

This matrix gives us a total of four transition probabilities: λ_{ee}, λ_{ec}, λ_{cc}, and λ_{ce}, for which we have $\lambda_{ij} = 1 - \lambda_{ii}$, $i \in \{e, c\}$ and $j \neq i$.

Denote the ergodic probabilities for this chain as ρ. This vector ρ is defined as the eigenvector of Λ associated with the unit eigenvalue; that is, the vector of ergodic probabilities ρ satisfies $\Lambda\rho = \rho$. The eigenvector ρ is normalized so that its elements sum to unity $1'\rho = 1$. For the two-state Markov chain studied here, we can derive the stationary probability vector as

$$\rho = \begin{bmatrix} \lambda_e \\ \lambda_c \end{bmatrix} = \begin{bmatrix} (1-\lambda_{cc})/(2-\lambda_{ee}-\lambda_{cc}) \\ (1-\lambda_{ee})/(2-\lambda_{ee}-\lambda_{cc}) \end{bmatrix}. \tag{10}$$

After defining the latent states, we write the following model:

$$b_t = \begin{cases} \alpha_e + X_t\beta_e + \epsilon_{et} & \text{if } s_t = e, \\ \alpha_c + X_t\beta_c + \epsilon_{ct} & \text{if } s_t = c, \end{cases} \quad t = 1, \ldots, T, \tag{11}$$

where b_t is the bid submitted at time t, and X_t is a vector of independent variables.[10] The error terms $\epsilon_{s_t t}$ are assumed to be independent of X_t. Following the standard approach, we assume $\epsilon_{et} \sim \mathrm{N}(0, \sigma_e^2)$, and $\epsilon_{ct} \sim \mathrm{N}(0, \sigma_c^2)$. Notice that for each period, the regime variable (Markov state s_t) is unobservable.

Because of a lack of information, such as the bidders' private per-click values and the CTR of each position, the theoretical model cannot be tested structurally. Many idealistic conditions assumed in the theoretical model, such as the absence of budget constraints, only two bidders, rational bids, no entry, complete information on competitor's bid in the previous round, and so on, are necessarily violated in the real world. Consequently, the main objective of this empirical model is to document and characterize observed cycles in the two markets. The result we obtain next is meant to offer a quantitative way to examine the cycles that we can observe and is by no means a proof of the theory. We discuss alternative explanations in the next subsection.

Our empirical model would yield state probabilities for each bid based on the history of bids. The general derivation of the estimation procedure follows

[10] For expositional simplicity, we suppress subscript i here, which indicates the bidder. We later add the subscript to distinguish bids submitted by different bidders.

Figure 3 **Bidding History and Smoothed Probabilities (Yahoo!)**

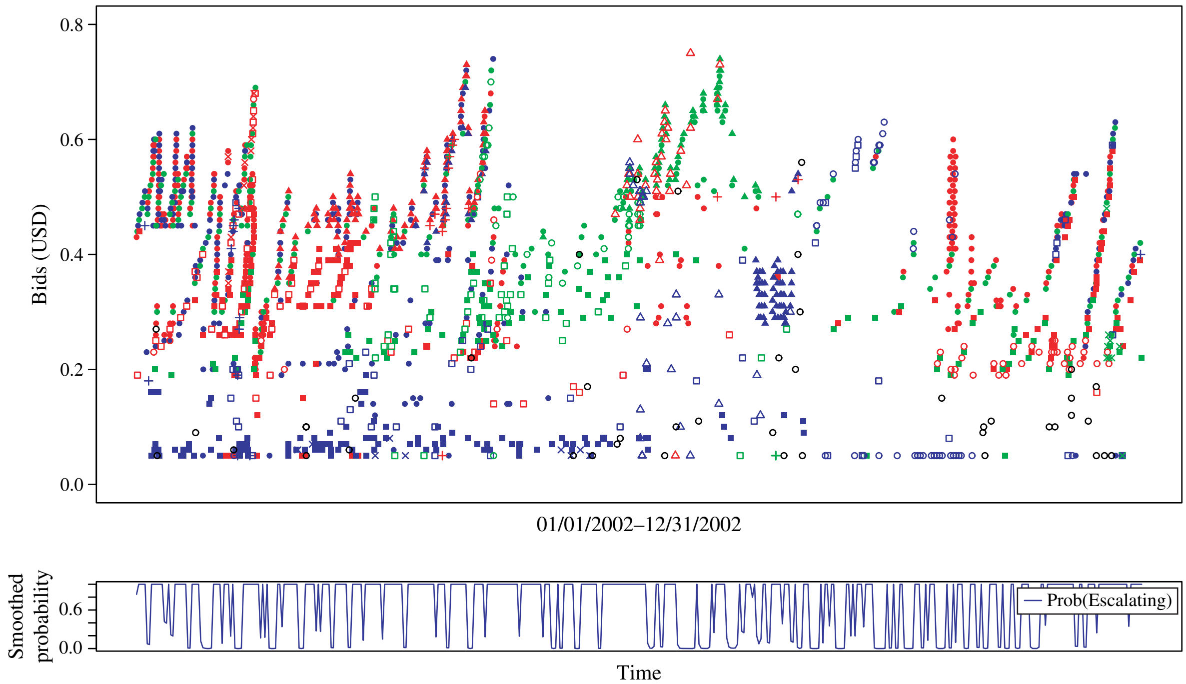

Cosslett and Lee (1985) and Hamilton (1989, 1990). We relegate a detailed description of the estimation procedure to the online appendix.

If transition probabilities are restricted only by the conditions that $\lambda_{kl} \geq 0$ and $\sum_{l=1}^{N} \lambda_{kl} = 1$, for all k and l, Hamilton (1990) shows that the maximum-likelihood estimates for the transition probabilities satisfy

$$\hat{\lambda}_{kl} = \frac{\sum_{t=2}^{T} \mathrm{Prob}(s_t = l, s_{t-1} = k \mid Y_T; \hat{\boldsymbol{\beta}})}{\sum_{t=2}^{T} \mathrm{Prob}(s_{t-1} = k \mid Y_T; \hat{\boldsymbol{\beta}})}, \tag{12}$$

Figure 4 **Bidding History and Smoothed Probabilities (Baidu)**

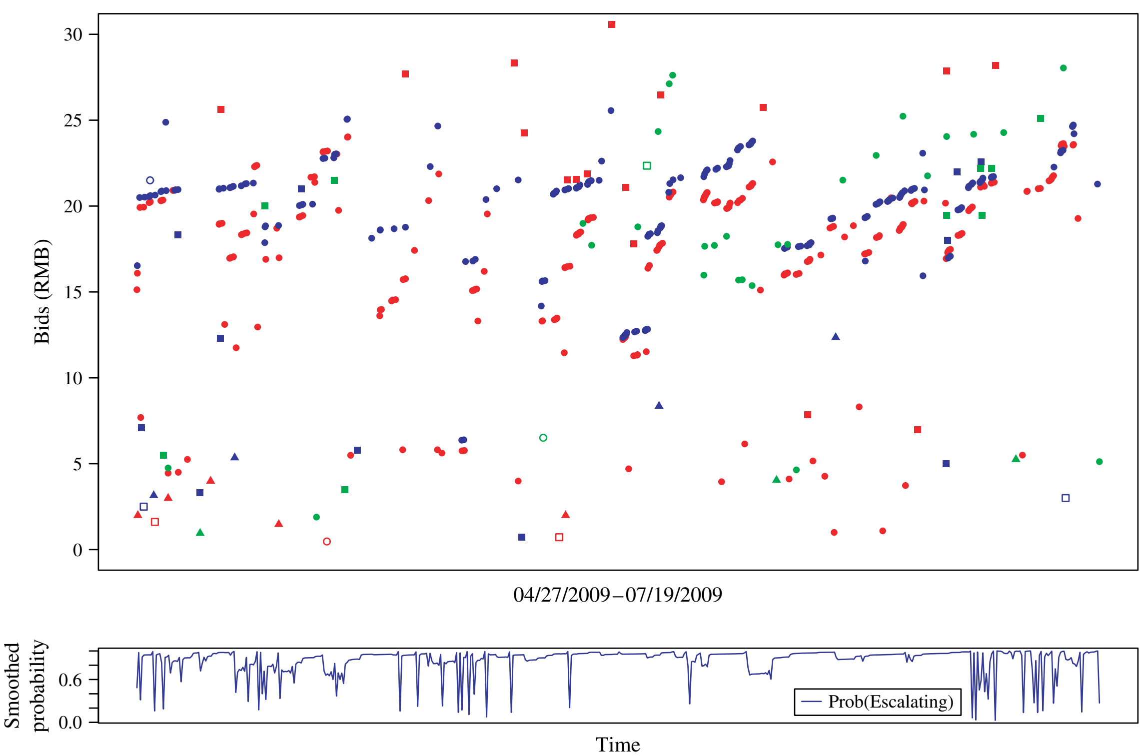

where Y_t is a vector containing all observations obtained through date t, and $\hat{\boldsymbol{\beta}}$ denotes the full vector of maximum-likelihood estimates. Equation (12) shows that the estimated transition probability $\hat{\lambda}_{kl}$ is the number of times state k has been followed by state l divided by the number of times the process is in state k.

Equipped with estimates of the transition probabilities, we can obtain other useful measures to characterize the cycles. The expected duration (ED) of a typical escalating or collapsing phase can be calculated with[11]

$$ED_k \equiv E(\text{duration of phase } s_t) = \frac{1}{1-\lambda_{s_t, s_t}}, \qquad (13)$$

where λ_{s_t, s_t} are the transition probabilities.

Thus, the expected duration of a cycle is

$$\begin{aligned} ED &\equiv E(\text{duration of a cycle}) \\ &= \sum_{s_t \in \{e, c\}} \frac{1}{1-\lambda_{s_t, s_t}} = \frac{1}{1-\lambda_{ee}} + \frac{1}{1-\lambda_{cc}}. \end{aligned}$$

We consider the following specification for model (11):

$$b_{i,t} = \begin{cases} \beta_{e0} + \beta_{e1} b_{i,t-1} + \beta_2 b_{-i,t} + \gamma_t + \epsilon_{et} & \text{if } s_t = e, \\ \beta_{c0} + \beta_{c1} b_{i,t-1} + \beta_2 b_{-i,t} + \gamma_t + \epsilon_{ct} & \text{if } s_t = c, \end{cases} \qquad (14)$$

where $b_{i,t}$ is the bid submitted by bidder i at time t, and $b_{-i,t}$ is the competitor's bid at time t. We also include bidder i's previous bid, $b_{i,t-1}$, as a proxy for the competitor's previous bid.[12] The variable γ_t represents a vector of day-of-week and month fixed effects.[13] The model fits an autoregressive process and estimates two sets of parameters corresponding to the two underlying states. We use a maximum-likelihood estimation to determine the parameters and the states simultaneously. The variable $b_{-i,t}$ is modeled as state independent, so the parameter estimate β_2 is the same for both states.[14] We have chosen Equation (14) as the specification because of its desirable feature of possibly being consistent with bidder i's decision-making process at time t. When a bidder determines her strategy in a period, she should look at her competitor's bid, compare it with her threshold value, and then decide whether she wants to outbid her competitor or drop the bid to some lower level. The bidder's state-switching decision is most likely triggered by her competitor's bid, so the parameter for $b_{i,t-1}$ is state dependent. After choosing the state, her decision would be related to the desired price level of her bid, which should be based on the competitor's bid, $b_{-i,t}$, and thus state independent.

Because our objective is to estimate price cycles to describe the bid-adjustment patterns we observe, we do not really rely on the form of Equation (14) to characterize the cycles. For each bidder, for example, we can run the following autoregressive AR(1) model to obtain individual-level cycle characteristics:

$$b_{i,t} = \begin{cases} \beta_{i,e0} + \beta_{i,e1} b_{i,t-1} + \gamma_t + \epsilon_{i,et} & \text{if } s_t = e, \\ \beta_{i,c0} + \beta_{i,c1} b_{i,t-1} + \gamma_t + \epsilon_{i,ct} & \text{if } s_t = c. \end{cases} \qquad (15)$$

What is important in these specifications is the derived state-probability vector with which we calculate transition probabilities and cycle durations. We next report our estimation results based on Equation (14).

The empirical model yields the latent state for each of the bids, along with estimated probabilities. We can also get parameter estimates for β_e and β_c, where $\boldsymbol{\beta}_e \equiv (\beta_{e0}, \beta_{e1}, \beta_2)$ and $\boldsymbol{\beta}_c \equiv (\beta_{c0}, \beta_{c1}, \beta_2)$.

Estimation results for both markets are reported in Table 2.

Our next objective is to characterize the cycles by applying the results of Markov chain theory. We use the smoothed probabilities to obtain the latent state of each bid. By examining the change of states following Equation (12), we calculate the transition matrices, as shown in Tables 3 and 4.

The transition matrices for both auctions suggest that, starting from either state, the next bid is much more likely to end up as escalating than collapsing. The imbalance between the two states can be shown

Table 2 **Markov Switching Estimates**

	First price (Yahoo!)		Second price (Baidu)	
	Escalating	Collapsing	Escalating	Collapsing
$\hat{\beta}_0$	0.38	0.02	6.04	1.16
(Intercept)	(0.02)	(0.00)	(1.52)	(0.44)
$\hat{\beta}_1$	1.01	0.04	1.28	0.06
(Previous bid)	(0.00)	(0.06)	(0.07)	(0.02)
$\hat{\beta}_2$	0.80		0.90	
(Next-highest bid)	(0.02)		(0.02)	

Note. Standard errors are reported in parentheses.

[11] For details, refer to Gallager (1996, Chap. 4).

[12] The bid submitted immediately before $b_{i,t}$ may not be from the competitor because bidder i's bid at time t might be reacting to some earlier bid.

[13] We thank an anonymous reviewer for suggesting this effect. A model without time-fixed effects yields almost the same parameter estimates for the βs. Inclusion of these time-fixed effects slightly reduces the standard errors.

[14] The smoothed state probabilities remain qualitatively the same when we (1) estimate a marketwide AR(1) model (i.e., Equation (14) without including the term $b_{-i,t}$) or (2) relax the state independence on β_2 and estimate separate parameters for each state as β_{e2} and β_{c2}.

Table 3 Transition Matrix for Yahoo! Cycles

λ_{ij}	e	c
e	0.94	0.06
c	0.92	0.08

Table 4 Transition Matrix for Baidu Cycles

λ_{ij}	e	c
e	0.852	0.148
c	0.998	0.002

by calculating the limiting unconditional probabilities. From Equation (10),

$$\rho_{\text{Yahoo!}} = \begin{bmatrix} 0.9388 \\ 0.0612 \end{bmatrix}$$

and

$$\rho_{\text{Baidu}} = \begin{bmatrix} 0.8709 \\ 0.1291 \end{bmatrix}.$$

Therefore, in the long run, for the first-price (second-price) auction, about 93.88% (87.09%) of the states are in the escalating phase, and only about 6.12% (12.91%) of the states are in the collapsing phase.

With Equation (13), we can calculate the duration of the escalating (collapsing) phase. For Yahoo!, we have $ED_E^{\text{Yahoo!}} = 16.67$, and $ED_C^{\text{Yahoo!}} = 1.09$. For Baidu, we have $ED_E^{\text{Baidu}} = 6.76$, and $ED_C^{\text{Baidu}} = 1.002$. These results suggest that a typical escalating phase lasts about 16 and 7 periods for the first- and second-price auctions, respectively. A typical collapsing phase lasts about 1 period in both cases.

We also estimate the length of a cycle in terms of hours. For the recorded bids, the distribution of the duration between consecutive bids is heavily skewed to the right. We adopt the median time between bids for each auction type and determine that one cycle would last for 28.93 hours in Yahoo! and 7.76 hours in Baidu.

4. Discussions

4.1. Ads' Quality Score

To facilitate better matching between ads and searches, search engines often reward higher-quality ads with a higher position, even when their bids are lower than those of some other ads. "Quality-adjusted" bids are calculated by multiplying the raw bids by such a quality score. Our model can be adapted to incorporate this quality adjustment. Following Varian (2007), let e_i represent the quality score of the ad of bidder i. The actual CTR of bidder i at position j, or z_{ij}, is then determined by both e_i and τ_j, $j = 1, 2$, so we have $z_{ij} = e_i \tau_j$.

With the above setup, ads will be ranked by $e_i b_i$. Like in our analysis before, each advertiser needs to pay only the minimum amount to participate.[15] Let p_{ij} be the minimum amount that advertiser i needs to pay to take position j, and we have $p_{i1} e_i = e_{-i} b_{-i}$ and $p_{i2} = r$, $i = 1, 2$. So, $p_{i1} = e_{-i}/(e_i) b_2$. Then,

$$\pi_i(b_i; b_{-i}) = \begin{cases} \left(\theta_i - \dfrac{e_{-i}}{e_i} b_{-i}\right) z_{i1} & \text{if } b_i e_i > b_{-i} e_{-i}, \\ \dfrac{1}{2}\left(\theta_i - \dfrac{e_{-i}}{e_i} b_{-i}\right) z_{i1} + \dfrac{1}{2}(\theta_i - r) z_{i2} & \text{if } b_i e_i = b_{-i} e_{-i}, \\ (\theta_i - r) z_{i2} & \text{if } b_i e_i < b_{-i} e_{-i}. \end{cases}$$

Or equivalently, plugging in z_{ij},

$$\pi_i(b_i; b_{-i}) = \begin{cases} (\theta_i e_i - e_{-i} b_{-i}) \tau_1 & \text{if } b_i e_i > b_{-i} e_{-i}, \\ \frac{1}{2}(\theta_i e_i - e_{-i} b_{-i}) \tau_1 + \frac{1}{2}(\theta_i e_i - r e_i) \tau_2 & \text{if } b_i e_i = b_{-i} e_{-i}, \\ (\theta_i e_i - r e_i) \tau_2 & \text{if } b_i e_i < b_{-i} e_{-i}. \end{cases}$$

Let $\theta_i^* = \theta_i e_i$, $r^* = e_i r$, $b_i^* = e_i b_i$, and $b_{-i}^* = e_{-i} b_{-i}$. We then can rewrite the above payoff function as

$$\pi_i(b_i^*; b_{-i}^*) = \begin{cases} (\theta_i^* - b_{-i}^*) \tau_1 & \text{if } b_i^* > b_{-i}^*, \\ \frac{1}{2}(\theta_i^* - b_{-i}^*) \tau_1 + \frac{1}{2}(\theta_i^* - r) \tau_2 & \text{if } b_i^* = b_{-i}^*, \\ (\theta_i^* - r^*) \tau_2 & \text{if } b_i^* < b_{-i}^*. \end{cases}$$

We can see that with the slightly modified interpretation of quality-adjusted value, reserve price, and bids, the payoff function remains structurally the same as in §2. Hence, the analysis and equilibrium outcomes follow as before. The transformation is achieved by considering the *actual* click-through rate as the product of a "position effect" and a "quality effect." The "position effect" is equivalent to τ in our baseline model, and the "quality effect" is related only to the ad itself. Here one implicit assumption is that when an ad moves to a new position, it carries the ad-quality effect; that is, the quality score is independent from the ad position.

[15] In this section, we consider only the case of second-price auction, because it is now the most commonly adopted mechanism. We are not aware of a similar practice for first-price auctions, but the argument for the first-price mechanism follows the same logic.

Table 5 Payoffs of Matching Pennies

	Head	Tail
Head	+1, −1	−1, +1
Tail	−1, +1	+1, −1

4.2. The Analogy Between Dynamic Reactions and Mixed Strategy

Although the dynamic equilibrium in our framework is based on pure strategies, the connection between this type of cyclical equilibrium and a traditional mixed-strategy equilibrium can be illustrated with a simple example.

Consider the familiar "matching pennies" game in Table 5. There is no pure-strategy equilibrium for either player. When they move simultaneously, they must play the mixed-strategy equilibrium in which both of them randomize, even in a repeated game setting. If the market allows them to move sequentially and observe each other's actions, the game's equilibrium outcome would be a special case of the Edgeworth cycle. Specifically, suppose player 1 plays heads first, then player 2's best response is to play tails in the second period. In the next period, when player 1 moves, she would switch to play tails. Then player 2 would switch to play heads, and so on. The Markov chain of this outcome is exactly the same as the one depicted in Figure 2. We only need to change "escalating" for "heads," and "collapsing" for "tails."

The implications may not be limited to the search-advertising market. Our results suggest that advertisers adjust their bids quite frequently because of the low menu cost of changing bids and low monitoring cost of observing competitors' bids. Similar to the Edgeworth-cycle patterns of retail gasoline prices (Noel 2007a, b), these bids display a pattern of price dispersion. This form of price dispersion is yet to be understood. Varian (1980) distinguished between "spatial price dispersion," in which some firms persistently sell the product at a lower price than others, and "temporal price dispersion," in which firms play a mixed strategy and randomize the price. The price dispersion resulting from Edgeworth-like cycles is obviously not spatial, because no bidder consistently bids higher than the others. It is not temporal, either, because the advertisers are not necessarily randomizing from period to period. We can refer to this form of price dispersion as "reactive price dispersion," because the price pattern is a result of players reacting to each other's strategies. In oligopolistic markets, firms usually cannot perfectly make their decisions simultaneously. When players have to move sequentially, late movers often can observe the strategies played by early movers and adjust their strategy accordingly. Neither a one-shot nor a repeated game framework is appropriate to model this form of reaction.

4.3. Autobidders and Alternative Causes

Autobidders are software agents that an advertiser can use to automatically change the bids on advertisers' behalf. Advertisers may use such autobidders to reduce the cost of changing bids.

We cannot perfectly rule out the possibility of the use of autobidders, but a few observations make us believe that the cyclical patterns we document cannot be attributed entirely to the use of autobidders, even if they exist. First, autobidders typically revisit a keyword in a fixed time interval (e.g., every 10 minutes, every day, etc.). Hence, the time between consecutive bids by autobidders should be close to a multiple of a constant. We cannot identify such regularity in bids in either market. Second, even if autobidders are used, the advertisers still need to determine a strategy for the program, and these autobidders would merely carry out a predefined strategy.

Can demand and cost shocks or date and time effects explain the observed cycles? We cannot fully rule out these possible causes. Summary statistics in Table 1 suggest, however, that advertisers change their bids every 2.3 days on Yahoo! and every 1.3 days on Baidu. If there are demand and cost shocks or date and time effects, bid changes should be closely related to daily or weekly human-activity patterns. Moreover, unless the shocks are extremely regular, such that the bidders are affected sequentially, it would also be unreasonable to observe sequential adjustments of bids in both search engines.

5. Conclusion

One distinctive feature of keyword auctions is that they are dynamic: Bids can be adjusted, ranks of ads can be updated, and payoffs are realized, all in real time. Recognizing the dynamic nature of these auctions, we find equilibrium bid adjustments that are not easily modeled and explained in static games. Our theoretical framework gives equilibrium conditions under which bidders engage in cyclical bid adjustments. In these cycles, the bids gradually rise up to a certain level, drop sharply, and then the cycles restart. We find that the cyclical pattern exists in both first- and second-price auctions. We also provide empirical evidence of cyclical bid adjustments with data from leading search engines. Our results show that advertisers not only change their bids, but they do so quite frequently.

A central message of this study is that there can be many types of equilibria in the search-engine advertising market. Although earlier studies offer important insights into the static properties of this market, we identify and emphasize dynamic interactions in this context. Data limitations preclude us from empirically studying the conditions under which cyclical bid adjustments are more easily observed. We nonetheless demonstrate the existence of such cycles and highlight the importance of adopting dynamic models to examine market participants' real-world actions.

One limitation of our study is that each data sample contains only one keyword. In this study, we look at each keyword as existing in a separate market. In reality, advertisers typically bid on multiple keywords. It is true that such issues as budgeting may affect bidding behavior, but within each individual keyword market, the advertisers still need to adopt a bidding strategy such as the one modeled and observed here and those in the literature. If the patterns we observe are rare, the generalizability of the study's findings could be significantly limited. To this end, the literature suggests that bid adjustments are far from isolated, idiosyncratic events. Edelman and Ostrovsky (2007, p. 194) document what they call a "saw-tooth" pattern in Yahoo!'s first-price data. From their description, it is obvious that the cyclical bidding pattern is quite prevalent. For second-price keyword auctions, the strongest support is from Ganchev et al. (2007), who show that (1) observed bid deviations from simple static models are likely to be the result of strategic bidding, and (2) there is evidence to support the practice of bid jamming.[16]

A dynamic perspective of the search-advertising market can potentially open a new door for important future works. There are a number of ways that our work can be improved upon to examine more realistic market conditions. For example, although this paper adopts a discrete-time framework to highlight bid reactions, the timing between consecutive bids can be endogenized in future studies. Richer strategic bidding patterns can also be studied if we explicitly consider issues that arise only in dynamic games, such as discount factors, budget constraints, menu costs, entry and exit, learning and heuristics, etc. In addition, future studies inevitably will have to deal with the issue of advertisers' bidding simultaneously in multiple keyword auctions. Finally, position- and advertiser-specific characteristics may affect the values of different slots, such that different slots may bring different sets of prospective customers, and different advertisers may value these slots differently. These nuances would be fruitful topics to examine in dynamic settings. We believe that further studies of dynamic reactions among advertisers will offer even more interesting insights about this important market.

[16] We thank an anonymous reviewer for pointing out this reference.

Acknowledgments

The authors are indebted to Tao Hong, David Pennock, Haoyu Shen, and Yunhong Zhou for graciously sharing data. For valuable comments and suggestions, the authors thank David Bell, Hemant Bhargava, Erik Brynjolfsson, Yiling Chen, Pradeep Chintagunta, Chris Dellarocas, Preyas Desai, Rob Fichman, Alok Gupta, Gary Koehler, John Little, Michael Noel, Hal Varian, Feng Zhu, the associate editor, three anonymous reviewers, and seminar participants at Boston College, Georgia Institute of Technology, Massachusetts Institute of Technology, National University of Singapore, University of Illinois at Urbana–Champaign, University of Southern California, the Wharton School, the 2005 International Conference on Information Systems, the 2006 INFORMS Conference, and the 2007 ACM Conference on Electronic Commerce.

Appendix

Proof of Proposition 1. We use the concept of MPE in this proof. First we show the conditions that the parameters $(\underline{b}_i, \bar{b}_i)$ must satisfy to constitute an MPE.

Let $V_i(b_{-i})$ denote advertiser i's valuation if (1) she is about to move, (2) the other firm just bid b_{-i}, and (3) both firms follow the equilibrium strategy thereafter. Define $V_1(b_2 = r) = \bar{V}_1$ and $V_2(b_1 = r+\epsilon) = \bar{V}_2$. Depending on whether $\underline{b}_1 - r$ is an odd or an even multiple of ϵ, $\bar{V}_1$ and $\bar{V}_2$ can be expressed as follows:

Case 1. $\underline{b}_1 - r = (2t+1)\epsilon$:

$$\begin{aligned}\bar{V}_1 &= V_1(b_2 = r)\\ &= (\theta_1 - r - \epsilon)(\tau_1 + \delta\tau_2)\\ &\quad + \delta^2(\theta_1 - r - 3\epsilon)(\tau_1+\delta\tau_2) + \cdots + \delta^{2t}(\theta_1 - \underline{b}_1)(\tau_1 + \delta\tau_2)\\ &\quad + \delta^{2t+2}(\theta_1 - \bar{b}_2)(1+\delta)\tau_1 + \delta^{2t+4}\bar{V}_1; \qquad (16)\end{aligned}$$

$$\begin{aligned}\bar{V}_2 &= V_2(b_1 = r+\epsilon)\\ &= (\theta_2 - r - 2\epsilon)(\tau_1 + \delta\tau_2)\\ &\quad + \delta^2(\theta_2 - r - 4\epsilon)(\tau_1+\delta\tau_2) + \cdots + \delta^{2t}(\theta_2 - \underline{b}_1 - \epsilon)\\ &\quad \cdot (\tau_1 + \delta\tau_2) + \delta^{2t+2}(\theta_2 - r)(1+\delta)\tau_2 + \delta^{2t+4}\bar{V}_2. \qquad (17)\end{aligned}$$

Case 2. $\underline{b}_1 - r = 2t\epsilon$:

$$\begin{aligned}\bar{V}_1 &= V_1(b_2 = r)\\ &= (\theta_1 - r - \epsilon)(\tau_1 + \delta\tau_2)\\ &\quad + \delta^2(\theta_1 - r - 3\epsilon)(\tau_1+\delta\tau_2) + \cdots + \delta^{2t-2}(\theta_1 - \underline{b}_1 + \epsilon)\\ &\quad \cdot (\tau_1 + \delta\tau_2) + \delta^{2t}(\theta_1 - \bar{b}_2)(1+\delta)\tau_1 + \delta^{2t+2}\bar{V}_1; \qquad (18)\end{aligned}$$

$$\begin{aligned}\bar{V}_2 &= V_2(b_1 = r+\epsilon)\\ &= (\theta_2 - r - 2\epsilon)(\tau_1 + \delta\tau_2)\\ &\quad + \delta^2(\theta_2 - r - 4\epsilon)(\tau_1+\delta\tau_2) + \cdots + \delta^{2t-2}(\theta_2 - \underline{b}_1)\\ &\quad \cdot (\tau_1 + \delta\tau_2) + \delta^{2t}(\theta_2 - r)(1+\delta)\tau_2 + \delta^{2t+2}\bar{V}_2. \qquad (19)\end{aligned}$$

Now given $\bar{V}_1$ and $\bar{V}_2$, how is $\bar{b}_2$ determined? Threshold $\bar{b}_2$ is the highest bid that advertiser 2 will submit for position 1. Given $\theta_1 > \theta_2$ and that advertiser 1 will bid ϵ more than advertiser 2 when the current price is lower than $\bar{b}_1$, if the current price is $\bar{b}_2 - \epsilon$, advertiser 2 should be better off outbidding advertiser 1 by ϵ than dropping back to r. If the current price is $\bar{b}_2$, advertiser 2 should be weakly better off by dropping back to r than outbidding advertiser 1 by ϵ, so

$$\begin{aligned}&(\theta_2-\bar{b}_2)(\tau_1+\delta\tau_2)+\delta^2(\theta_2-r)(1+\delta)\tau_2+\delta^4\bar{V}_2\\&\quad>(\theta_2-r)(1+\delta)\tau_2+\delta^2\bar{V}_2\\&\quad\geq(\theta_2-\bar{b}_2-\epsilon)(\tau_1+\delta\tau_2)+\delta^2(\theta_2-r)(1+\delta)\tau_2+\delta^4\bar{V}_2.\end{aligned}\quad(20)$$

More specifically, $\bar{b}_2$ is determined when the the second inequality is binding

$$\begin{aligned}&(\theta_2-r)(1+\delta)\tau_2+\delta^2\bar{V}_2\\&\quad=(\theta_2-\bar{b}_2-\epsilon)(\tau_1+\delta\tau_2)+\delta^2(\theta_2-r)(1+\delta)\tau_2+\delta^4\bar{V}_2.\end{aligned}\quad(21)$$

It is obvious that $\bar{b}_2<\theta_2$, because by dropping to bid r, advertiser 2 can receive the second position and make a positive profit. This can be shown by reorganizing the second inequality of Equation (20) as $(\theta_2-\bar{b}_2-\epsilon)(\tau_1+\delta\tau_2)=(1-\delta^2)(\theta_2-r)(1+\delta)\tau_2+\delta^2\bar{V}_2>0$; thus $\theta_2-\bar{b}_2-\epsilon>0$.

Then what condition should $\underline{b}_1$ satisfy? Threshold $\underline{b}_1$ is the threshold value that advertiser 1 prefers to jump to bid advertiser 2's maximum bid to guarantee the first position and end the price war. Therefore, given $\theta_1>\theta_2$ and that advertiser 2 follows the strategy to outbid advertiser 1 by ϵ when the current price is lower than $\underline{b}_1$, it should be the case that if the current price is $\underline{b}_1-\epsilon$, advertiser 1 is better off by bidding $\underline{b}_1$ than jumping to $\bar{b}_2$; whereas if the current price is $\underline{b}_1$, advertiser 1 is weakly better off by jumping to $\bar{b}_2$ in the current period than bidding $\underline{b}_1+\epsilon$. Thus, the following needs to be satisfied:

$$\begin{aligned}&(\theta_1-\underline{b}_1)(\tau_1+\delta\tau_2)+\delta^2(\theta_1-\bar{b}_2)(1+\delta)\tau_1+\delta^4\bar{V}_1\\&\quad>(\theta_1-\bar{b}_2)(1+\delta)\tau_1+\delta^2\bar{V}_1\\&\quad\geq(\theta_1-\underline{b}_1-\epsilon)(\tau_1+\delta\tau_2)+\delta^2(\theta_1-\bar{b}_2)(1+\delta)\tau_1+\delta^4\bar{V}_1.\end{aligned}\quad(22)$$

Then $\underline{b}_1$ is determined by the second inequality when it is binding

$$\begin{aligned}&(\theta_1-\bar{b}_2)(1+\delta)\tau_1+\delta^2\bar{V}_1\\&\quad=(\theta_1-\underline{b}_1-\epsilon)(\tau_1+\delta\tau_2)+\delta^2(\theta_1-\bar{b}_2)(1+\delta)\tau_1+\delta^4\bar{V}_1.\end{aligned}\quad(23)$$

It can be shown that $\underline{b}_1<\theta_1$ by reorganizing the second inequality as $(\theta_1-\underline{b}_1-\epsilon)(\tau_1+\delta\tau_2)=(1-\delta^2)((\theta_1-\bar{b}_1)\cdot(1+\delta)\tau_1+\delta^2\bar{V}_1)>0$. It can also be shown that $\underline{b}_1$ should be lower than $\bar{b}_2$, otherwise the first inequality can not be satisfied, because $\bar{V}_1\geq(\theta_1-\bar{b}_2)(1+\delta)\tau_1+\delta^2\bar{V}_1$.

Solving Equations (21) and (23) gives $\underline{b}_1$ and $\bar{b}_2$.

Next we show that (R_1,R_2) constitutes an MPE. We have the following claims:

1. Both advertisers will never bid lower than r, because r is the smallest allowed bid.

2. Advertiser 2 will never raise her bid to more than $\bar{b}_2+\epsilon$. To show this, compare advertiser 2's payoff when bidding $\bar{b}_2+2\epsilon$ to that when she drops to bid r:

$$\begin{aligned}&(\theta_2-\bar{b}_2-2\epsilon)(\tau_1+\delta\tau_2)+\delta^2(\theta_2-r)(1+\delta)\tau_2+\delta^4\bar{V}_2\\&\quad<(\theta_2-r)(1+\delta)\tau_2+\delta^2\bar{V}_2.\end{aligned}$$

This inequality follows from Equation (20), because $\theta-\bar{b}_2-2\epsilon<\theta-\bar{b}_2-\epsilon$.

3. Advertiser 1 will never bid between $[\underline{b}_1+\epsilon,\bar{b}_2-\epsilon]$. This can be seen from (22), and the fact that $\theta_1-\underline{b}_1-2\epsilon<\theta_1-\underline{b}_1-\epsilon$.

4. For $r<b_{-i}<\underline{b}_1$, no advertiser is willing to unilaterally deviate from bidding $b_{-i}+\epsilon$ to $b_{-i}+2\epsilon$ (in turn, $b_{-i}+k\epsilon$ for $k>2$). To show this, let the current bid price be p. We consider only the case of $\underline{b}_1-p=(2t+1)\epsilon$. The case when $\underline{b}_1-p=2t\epsilon$ can be worked out similarly:

$$\begin{aligned}V_1(b_2=p)&=(\theta_1-p-\epsilon)(\tau_1+\delta\tau_2)\\&\quad+\delta^2(\theta_1-p-3\epsilon)(\tau_1+\delta\tau_2)+\cdots+\delta^{2t}(\theta_1-\underline{b}_1)\\&\quad\cdot(\tau_1+\delta\tau_2)+\delta^{2t+2}(\theta_1-\bar{b}_2)(1+\delta)\tau_1+\delta^{2t+4}\bar{V}_1;\\V_2(b_1=p)&=(\theta_2-p-\epsilon)(\tau_1+\delta\tau_2)\\&\quad+\delta^2(\theta_2-p-3\epsilon)(\tau_1+\delta\tau_2)+\cdots+\delta^{2t}(\theta_2-\underline{b}_1)\\&\quad\cdot(\tau_1+\delta\tau_2)+\delta^{2t+2}(\theta_2-r)(1+\delta)\tau_2+\delta^{2t+4}\bar{V}_2.\end{aligned}$$

If advertiser 1 switches to bid $b_2+2\epsilon$ for only one period when advertiser 2 keeps the original strategy, then

$$\begin{aligned}\tilde{V}_1(b_2=p)&=(\theta_1-p-2\epsilon)(\tau_1+\delta\tau_2)\\&\quad+\delta^2(\theta_1-p-4\epsilon)(\tau_1+\delta\tau_2)+\cdots+\delta^{2t}(\theta_1-\underline{b}_1-\epsilon)\\&\quad\cdot(\tau_1+\delta\tau_2)+\delta^{2t+2}(\theta_1-\bar{b}_2)(1+\delta)\tau_1+\delta^{2t+4}\bar{V}_1\\&<V_1(b_2=p).\end{aligned}$$

This inequality is obvious because the first t terms in $V_1(p)$ are greater than those in $\tilde{V}_1(p)$. The same reasoning works for advertiser 2:

$$\begin{aligned}\tilde{V}_2(b_1=p)&=(\theta_2-p-2\epsilon)(\tau_1+\delta\tau_2)\\&\quad+\delta^2(\theta_2-p-4\epsilon)(\tau_1+\delta\tau_2)+\cdots+\delta^{2t}(\theta_2-\underline{b}_2-\epsilon)\\&\quad\cdot(\tau_1+\delta\tau_2)+\delta^{2t+2}(\theta_2-r)(1+\delta)\tau_2+\delta^{2t+4}\bar{V}_2\\&<V_2(b_2=p).\end{aligned}$$

This inequality is straightforward because the first t terms in $V_2(p)$ are greater than those in $\tilde{V}_2(p)$. Thus, neither advertiser has incentives to deviate to bid $b_i+2\epsilon$.

5. For $r<b_{-i}<\underline{b}_2$, no advertiser is willing to deviate to bid $p<b_i+\epsilon$. This follows the same logic as (4) thus is omitted.

6. For $\underline{b}_1<b_{-i}<\bar{b}_2$, advertiser 1 prefers to jumping to bid $\bar{b}_{-i}$. This is straightforward from (22).

By all of these claims, (R_1,R_2) constitutes an MPE for this game.

Finally we show the existence of the MPE. Consider an arbitrary $\bar{V}_1$ and $\bar{V}_2\in(0,((\theta_i-r-\epsilon)\tau_1)/(1-\delta))$,[17] and $\bar{b}_2(\bar{V}_2)$ is defined by Equation (21). It can be shown that $\bar{b}_2(\bar{V}_2)$ is continuous in $\bar{V}_2$ and is in $[r,\theta_2]$.

Similarly, given $\bar{V}_1$, $\bar{V}_2$, and $\bar{b}_1(\bar{V}_2)$, $\underline{b}_2(\bar{V}_1)$ is defined by Equation (23). It is continuous in $\bar{V}_1$ and is in $[r,\theta_1]$.

[17] From Equations (16)–(19), it is straightforward to obtain that both $\bar{V}_1$ and $\bar{V}_2$ are in $[0,((\theta_i-r-\epsilon)\tau_1)/(1-\delta)]$.

Consider $\bar{V}_1$ first. Define

$$U_1(p, \bar{V}_1) = (\theta_1 - p - \epsilon)(\tau_1 + \delta\tau_2) + \cdots + \begin{cases} \delta^{2t}(\theta_1 - \underline{b}_1)(\tau_1 + \delta\tau_2) + \delta^{2t+2}(\theta_1 - \bar{b}_2)(1+\delta)\tau_1 + \delta^{2t+4}\bar{V}_1 \\ \qquad \text{if } \underline{b}_2 - p = (2t+1)\epsilon, \\ \delta^{2t-2}(\theta_1 - \underline{b}_1 + \epsilon)(\tau_1 + \delta\tau_2) + \delta^{2t}(\theta_1 - \bar{b}_2)(1+\delta)\tau_1 + \delta^{2t+2}\bar{V}_1 \\ \qquad \text{if } \underline{b}_2 - p = 2t\epsilon. \end{cases} \tag{24}$$

We must show that there exists a $\bar{V}_1$ such that $\bar{V}_1 = U_1(\underline{b}_1(\bar{V}_1), \bar{b}_2(\bar{V}_2))$, $\bar{V}_1 = \tilde{U}_1(\bar{V}_1, \bar{V}_2)$. It is easy to show that $U_1(\bar{V}_1, \bar{V}_2)$ is continuous in $\bar{V}_1$. It can be shown $U_1(\bar{V}_1, \bar{V}_2)$ is continuous in $\bar{V}_2$, because U_1 is continuous in $\bar{b}_2(\bar{V}_2)$, and $\bar{b}_2(\bar{V}_2)$ is continuous in $\bar{V}_2$. To show that $\tilde{U}_1$ has a fixed point, we only need to show that $\tilde{U}_1$ maps $[0, ((\theta_1 - r - \epsilon)\tau_1)/(1-\delta)]$ into $[0, ((\theta_1 - r - \epsilon)\tau_1)/(1-\delta)]$. Obviously, $\tilde{U}_1 \geq 0$ if $\bar{V}_1 \geq 0$. To show $\tilde{U}_1 \leq ((\theta_1 - r - \epsilon)\tau_1)/(1-\delta)$ if $\bar{V}_1 \leq ((\theta_1 - r - \epsilon)\tau_1)/(1-\delta)$, note that from Equation (24), when $\underline{b}_1 - r = (2t+1)\epsilon$,

$$\begin{aligned} \tilde{U}_1(\bar{V}_1) &\leq (\theta_1 - r - \epsilon)(1 + \delta + \cdots + \delta^{2t+3})\tau_1 + \delta^{2t+4}\frac{(\theta_1 - r - \epsilon)}{(1-\delta)}\tau_1 \\ &= \frac{(\theta_1 - r - \epsilon)(1 - \delta^{2t+4})}{(1-\delta)}\tau_1 + \delta^{2t+4}\frac{(\theta_1 - r - \epsilon)}{(1-\delta)}\tau_1 \\ &= \frac{(\theta_1 - r - \epsilon)}{(1-\delta)}\tau_1. \end{aligned}$$

The same works for the case when $\underline{b}_1 - r = 2t\epsilon$. It can also be shown, if we define $\tilde{U}_2$ in the same way, that $\tilde{U}_2 \leq ((\theta_i - r - \epsilon)\tau_1)/(1-\delta)$ if $\bar{V}_2 \leq ((\theta_i - r - \epsilon)\tau_1)/(1-\delta)$. Applying the fixed-point theorem, we complete the proof of the existence of the equilibrium. □

Proof of Proposition 2. The proof follows the same logic as in the proof of Proposition 1. The difference is that we embed the bidders' relative positions of the last period into the strategy profile (R_1, R_2) so that the response depends on both the other player's last bid and the bidders' relative positions. In other words, the bidders' payoff-relevant states are determined by both the last bid and how the bid was reached (i.e., from an escalating state or a collapsing state). Because the same bid price can be in different states, the responses can be different even when the last bids were the same. The equilibrium here is not an MPE and therefore is not "robust" in the sense discussed by Maskin and Tirole (1988b). Because the objective is to show existence, it is sufficient to identify one equilibrium out of many possible equilibria.

First we establish the necessary conditions under which there does not exist price cycles. If there exists an equilibrium in which both advertisers have no incentives to deviate by updating their bids, assuming that $\theta_1 > \theta_2$, it must be the case that bidder 1 bids some b_1^* and bidder 2 bids $b_1^* - \epsilon$, and bidders 1 and 2 obtain payoffs of $(1/(1-\delta))(\theta_1 - b_1^* + \epsilon)\tau_1$ and $(1/(1-\delta))(\theta_2 - r)\tau_2$, respectively. Consider such a pair of bids $(b_1^*, b_1^* - \epsilon)$; because bidder 1 has no incentive to raise her bid as she already occupies the first position, we only need to consider her deviation of lowering her bid. Bidder 1's best option is to bid $b_1^* - 2\epsilon$ because this induces the highest cost for bidder 2. In response, bidder 2 will drop the bid to $b_1^* - 3\epsilon$. In this scenario, bidder 1's deviation can be prevented if

$$(\theta_1 - r)\tau_2 + \frac{\delta}{1-\delta}(\theta_1 - b_1^* + 3\epsilon)\tau_1 \leq \frac{1}{1-\delta}(\theta_1 - b_1^* + \epsilon)\tau_1. \tag{25}$$

Similarly, bidder 2 should have no incentive to lower her bid because it only lowers the cost of bidder 1. Thus, we only need to consider bidder 2's deviation of raising her bid. The best deviation option for bidder 2 is to bid $b_1^* + \epsilon$. When bidder 2 deviates, bidder 1 responds to outbid bidder 2 with $b_1^* + 2\epsilon$. In this case, the following inequality should be satisfied to prevent bidder 2 from deviating:

$$(\theta_2 - b_1^*)\tau_1 + \delta\frac{1}{1-\delta}(\theta_2 - r)\tau_2 \leq \frac{1}{1-\delta}(\theta_2 - r)\tau_2. \tag{26}$$

From Equations (25) and (26), we have

$$\theta_2 - \frac{\tau_2}{\tau_1}(\theta_2 - r) \leq b_1^* \leq \theta_1 - \frac{\tau_2}{\tau_1}(\theta_1 - r) - \frac{3\delta - 1}{1-\delta}\epsilon.$$

Because $\epsilon > -((3\delta - 1)/(1-\delta))\epsilon$, we obtain that the necessary condition for the existence of a stable equilibrium is that there exists a b_i^* such that

$$\theta_2 - \frac{\tau_2}{\tau_1}(\theta_2 - r) \leq b_1^* \leq \theta_1 - \frac{\tau_2}{\tau_1}(\theta_1 - r) - \frac{3\delta - 1}{1-\delta}\epsilon.$$

So if a stable equilibrium exists such that bidder 1 and 2 bid b_1^* and $b_1^* - \epsilon$, respectively, the following must be satisfied:

$$\theta_2 - \frac{\tau_2}{\tau_1}(\theta_2 - r) \leq \theta_1 - \frac{\tau_2}{\tau_1}(\theta_1 - r) - \frac{3\delta - 1}{1-\delta}\epsilon.$$

Or equivalently,

$$\theta_1 - \theta_2 \geq \frac{\tau_1}{\tau_1 - \tau_2}\frac{3\delta - 1}{1-\delta}\epsilon. \tag{27}$$

Thus, a necessary condition for the existence of cyclical equilibrium is

$$\theta_1 - \theta_2 < \frac{\tau_1}{\tau_1 - \tau_2}\frac{3\delta - 1}{1-\delta}\epsilon. \tag{28}$$

Now we consider the equilibrium with price cycles. We first show the conditions that $\underline{b}_2'$ must satisfy to constitute a perfect dynamic equilibrium.

Define $V_1(b_2 = r + \epsilon) = \bar{V}_1$ and $V_2(b_1 = r) = \bar{V}_2$.

1. If $\underline{b}_2' = r + 2t\epsilon$, then $\bar{V}_1$ can be written as

$$\begin{aligned} \bar{V}_1 &= V_1(b_2 = r + \epsilon) \\ &= (\theta_1 - r - \epsilon)\tau_1 + \delta(\theta_1 - r)\tau_2 \\ &\quad + \delta^2[(\theta_1 - r - 3\epsilon)\tau_1 + \delta(\theta_1 - r)\tau_2] \\ &\quad + \cdots + \delta^{2t-2}(\theta_1 - \underline{b}_2' + \epsilon)\tau_1(1+\delta) \\ &\quad + \delta^{2t}(\theta_1 - r)\tau_2(1+\delta)\delta^{2t+2}\bar{V}_1, \end{aligned} \tag{29}$$

and $\bar{V}_2$ can be written as

$$\begin{aligned} \bar{V}_2 &= V_2(b_1 = r) \\ &= (\theta_2 - r)\tau_1 + \delta(\theta_2 - r)\tau_2 + \delta^2[(\theta_2 - r - 2\epsilon)\tau_1 + \delta(\theta_2 - r)\tau_2] \\ &\quad + \cdots + \delta^{2t-2}[(\theta_2 - \underline{b}_2' + 2\epsilon)\tau_1 + \delta(\theta_2 - r)\tau_2] \\ &\quad + \delta^{2t}[(\theta_2 - r)\tau_2 + \delta(\theta_2 - r)\tau_1] + \delta^{2t+2}\bar{V}_2. \end{aligned} \tag{30}$$

2. If $\underline{b}_2' = r + (2t+1)\epsilon$, then $\bar{V}_1$ can be written as

$$\begin{aligned}\bar{V}_1 &= V_1(b_2 = r+\epsilon)\\ &= (\theta_1 - r - \epsilon)\tau_1 + \delta(\theta_1 - r)\tau_2 + \delta^2[(\theta_1 - r - 3\epsilon)\tau_1\\ &\quad + \delta(\theta_1 - r)\tau_2] + \cdots + \delta^{2t-2}(\theta_1 - \underline{b}_2' + 2\epsilon)\tau_1 + \delta(\theta_1 - r)\tau_2\\ &\quad + \delta^{2t}(\theta_1 - \underline{b}_2')\tau_1(1+\delta)\\ &\quad + \delta^{2t+2}(\theta_1 - r)\tau_2(1+\delta)\delta^{2t+4}\bar{V}_1, \end{aligned} \tag{31}$$

and $\bar{V}_2$ can be written as

$$\begin{aligned}\bar{V}_2 &= V_2(b_1 = r)\\ &= (\theta_2 - r)\tau_1 + \delta(\theta_2 - r)\tau_2 + \delta^2[(\theta_2 - r - 2\epsilon)\tau_1 + \delta(\theta_2 - r)\tau_2]\\ &\quad + \cdots + \delta^{2t}[(\theta_2 - \underline{b}_2' + \epsilon)\tau_1 + \delta(\theta_2 - r)\tau_2]\\ &\quad + \delta^{2t+2}[(\theta_2 - r)\tau_2 + \delta(\theta_2 - r)\tau_1] + \delta^{2t+4}\bar{V}_2. \end{aligned} \tag{32}$$

Now how is $\underline{b}_2'$ determined? It is a threshold value above which bidder 2 is not willing to pay for staying at the top position. So when $b_1 = \underline{b}_2' - \epsilon$, bidder 2 is strictly better off outbidding bidder 1; when $b_1 = \underline{b}_2'$, bidder 2 is weakly better off by bidding $b_1 - \epsilon$ and staying at the bottom position than outbidding bidder 1. So we have

$$\begin{aligned}&(\theta_2 - \underline{b}_2' + \epsilon)\tau_1 + \delta(\theta_2 - r)\tau_2 + \delta^2[(\theta_2 - r)\tau_2 + \delta(\theta_2 - r)\tau_1] + \delta^4\bar{V}_2\\ &\quad > (\theta_2 - r)\tau_2 + \delta(\theta_2 - r)\tau_1 + \delta^2\bar{V}_2, \end{aligned} \tag{33}$$

and

$$\begin{aligned}&(\theta_2 - \underline{b}_2')\tau_1 + \delta(\theta_2 - r)\tau_2 + \delta^2[(\theta_2 - r)\tau_2 + \delta(\theta_2 - r)\tau_1] + \delta^4\bar{V}_2\\ &\quad \le (\theta_2 - r)\tau_2 + \delta(\theta_2 - r)\tau_1 + \delta^2\bar{V}_2. \end{aligned} \tag{34}$$

In summary,

$$\begin{aligned}&(\theta_2 - \underline{b}_2')\tau_1 + \delta(\theta_2 - r)\tau_2\\ &\quad \le (1-\delta^2)[(\theta_2 - r)\tau_2 + \delta(\theta_2 - r)\tau_1 + \delta^2\bar{V}_2]\\ &\quad < (\theta_2 - \underline{b}_2' + \epsilon)\tau_1 + \delta(\theta_2 - r)\tau_2, \end{aligned} \tag{35}$$

and $\underline{b}_2'$ is determined by the first inequality:

$$\begin{aligned}&(\theta_2 - \underline{b}_2')\tau_1 + \delta(\theta_2 - r)\tau_2\\ &\quad = (1-\delta^2)[(\theta_2 - r)\tau_2 + \delta(\theta_2 - r)\tau_1 + \delta^2\bar{V}_2]. \end{aligned} \tag{36}$$

We next show that (R_1, R_2) constitutes a perfect dynamic equilibrium by checking the following claims:

1. Neither advertiser will bid lower than r.
2. Advertiser 2 will never bid more than $\underline{b}_2'$, according to (35).
3. Advertiser 1 will never bid more than $\underline{b}_2' + \epsilon$, because given advertiser 2's strategy to bid $b_1 - \epsilon$ when $b_1 \ge \underline{b}_2'$, bidding more will only harm bidder 1's own per-click profit.
4. Neither advertiser has incentive to bid more than $b_{-i} + \epsilon$. The argument is similar to that in the first-price case (point 4).
5. When advertiser 2 plays strategy $b_1 - \epsilon$, and when $b_2 = \underline{b}_2'$, advertiser 1 has no incentive to deviate from bidding r. This can be seen from (36) and the violation of (27).
6. When $b_2 = \underline{b}_2'$, advertiser 1 has no incentive to deviate from bidding r to any $r + k\epsilon < \underline{b}_2' + \epsilon$, where $k > 0$ is an integer. This is because the payoffs of the periods following the deviation will be strictly lower than the payoff that bidder 1 can achieve when she bids r.
7. In Strategy II, bidder 2 has no incentive to bid $b_1 - k\epsilon$, with any integer $k > 1$. This is because the bid-jamming strategy ($k = 1$) can impose the highest cost for her competitor.

So (R_1, R_2) constitutes a perfect equilibrium. The existence of the equilibrium could be derived from applying the fixed-point theorem as in the first-price case. □

References

Animesh, A., V. Ramachandran, S. Viswanathan. 2010. Quality uncertainty and the performance of online sponsored search markets: An empirical investigation. *Inform. Systems Res.* **21**(1) 190–201.

Animesh, A., S. Viswanathan, R. Agarwal. 2011. Competing "creatively" in online markets: Evidence from sponsored search. *Inform. Systems Res.* **22**(1) 153–169.

Arnold, M., E. Darmon, T. Penard. 2011. To sponsor or not to sponsor: Sponsored search auctions with organic links and firm dependent clickthrough rates. Working paper, University of Delaware, Newark.

Athey, S., K. Bagwell. 2001. Optimal collusion with private information. *RAND J. Econom.* **32**(3) 428–465.

Athey, S., K. Bagwell. 2004. Dynamic auctions with persistent private information. Working paper, Harvard University, Cambridge, MA.

Athey, S., K. Bagwell. 2008. Collusion with persistent cost shocks. *Econometrica* **76**(3) 493–540.

Athey, S., K. Bagwell, C. Sanchirico. 2004. Collusion and price rigidity. *Rev. Econom. Stud.* **71**(2) 317–349.

Borgers, T., I. J. Cox, M. Pesendorfer, V. Petricek. 2007. Equilibrium bids in auctions of sponsored links: Theory and evidence. Working paper, Department of Economics, University of Michigan, Ann Arbor.

Breese, J. S., D. Heckerman, C. Kadie. 1998. Empirical analysis of predictive algorithms for collaborative filtering. G. Cooper, S. Moral, eds. *Proc. 14th Conf. Uncertainty in Artificial Intelligence*, Morgan Kaufmann, San Francisco, 43–52.

Cary, M., A. Das, B. Edelman, I. Giotis, K. Heimerl, A. R. Karlin, C. Mathieu, M. Schwarz. 2008. On best-response bidding in GSP auctions. NBER Working Paper 13788, National Bureau of Economic Research, Cambridge, MA.

Chen, J., J. Feng, A. Whinston. 2009a. Keyword auctions, unit-price contracts, and the role of committment. *Production Oper. Management* **19**(3) 305–321.

Chen, J., D. Liu, A. Whinston. 2009b. Auctioning keywords in online search. *J. Marketing* **73**(4) 125–141.

Cosslett, S. R., L.-F. Lee. 1985. Serial correlation in discrete variable models. *J. Econometrics* **27**(1) 79–97.

Edelman, B., M. Ostrovsky. 2007. Strategic bidder behavior in sponsored search auctions. *Decision Support Systems* **43**(1) 192–198.

Edelman, B., M. Ostrovsky, M. Schwarz. 2007. Internet advertising and the generalized second price auction: Selling billions of dollars worth of keywords. *Amer. Econom. Rev.* **97**(1) 242–259.

Edgeworth, F. Y., ed. 1925. The pure theory of monopoly. *Papers Relating to Political Economy*, Vol. 1. Macmillan, London, 111–142.

Ellison, G. 1994. Theories of cartel stability and the joint executive committee. *RAND J. Econom.* **25**(1) 37–57.

Feng, J. 2008. Optimal mechanism for selling a set of commonly ranked objects. *Marketing Sci.* **27**(3) 501–512.

Feng, J., X. M. Zhang. 2007. Dynamic price competition on the Internet: Advertising auctions. *Proc. 8th ACM Conf. Electronic Commerce*, Association for Computing Machinery, New York, 57–58.

Feng, J., H. Bhargava, D. Pennock. 2007. Implementing sponsored search in Web search engines: Computational evaluation of alternative mechanisms. *INFORMS J. Comput.* **19**(1) 137–148.

Gallager, R. 1996. *Discrete Stochastic Processes*. Kluwer Academic Publishers, Norwell, MA.

Ganchev, K., A. Kulesza, J. Tan, R. Gabbard, Q. Liu, M. Kearns. 2007. Empirical price modeling for sponsored search. X. Deng, F. C. Graham, eds. *Proc*. *3rd Internat*. *Conf*. *Internet and Network Econ*. (*WINE*), Springer-Verlag, Berlin, 541–548.

Ghose, A., S. Yang. 2009. An empirical analysis of search engine advertising: Sponsored search in electronic markets. *Management Sci.* **55**(10) 1605–1622.

Goldfarb, A., C. Tucker. 2011. Search engine advertising: Channel substitution when pricing ads to context. *Management Sci.* **57**(3) 458–470.

Goldfeld, S. M., R. E. Quandt. 1973. A Markov model for switching regressions. *J. Econometrics* **1**(1) 3–16.

Green, E. J., R. H. Porter. 1984. Noncooperative collusion under imperfect price competition. *Econometrica* **52**(1) 87–100.

Hamilton, J. D. 1989. A new approach to the economic analysis of nonstationary time series and the business cycle. *Econometrica* **57**(2) 357–384.

Hamilton, J. D. 1990. Analysis of time series subject to changes in regime. *J. Econometrics* **45**(1–2) 39–70.

Katona, Z., M. Sarvary. 2010. The race for sponsored links: Bidding patterns for search advertising. *Marketing Sci.* **29**(2) 199–215.

Lee, L.-F., R. H. Porter. 1984. Switching regression models with imperfect sample separation information—With an application on cartel stability. *Econometrica* **52**(2) 391–418.

Liu, D., J. Chen. 2006. Designing online auctions with past performance information. *Decision Support Systems* **42**(3) 1307–1320.

Liu, D., J. Chen, A. B. Whinston. 2010. Ex ante information and design of keyword auctions. *Inform. Systems Res.* **21**(1) 133–153.

Maskin, E., J. Tirole. 1988a. A theory of dynamic oligopoly, I: Overview and quantity competition with large fixed costs. *Econometrica* **56**(3) 549–569.

Maskin, E., J. Tirole. 1988b. A theory of dynamic oligopoly, II: Price competition, kinked demand curves, and Edgeworth cycles. *Econometrica* **56**(3) 571–599.

Noel, M. 2007a. Edgeworth price cycles, cost-based pricing and sticky pricing in retail gasoline markets. *Rev. Econom. Statist.* **89**(2) 324–334.

Noel, M. 2007b. Edgeworth price cycles: Evidence from the Toronto retail gasoline market. *J. Indust. Econom.* **55**(1) 69–92.

Porter, R. H. 1983. A study of cartel stability: The joint executive committee, 1880–1886. *Bell J. Econom.* **14**(2) 301–314.

Rotemberg, J. J., G. Saloner. 1986. A supergame-theoretic model of price wars during booms. *Amer. Econom. Rev.* **76**(3) 390–407.

Rutz, O., R. Bucklin. 2007. A model of individual keyword performance in paid search advertising. Working paper, Yale School of Management, Yale University, New Haven, CT.

Rutz, O., R. Bucklin. 2008. From generic to branded: A model of spillover in paid search advertising. Working paper, Yale School of Management, Yale University, New Haven, CT.

Stigler, G. J. 1964. A theory of oligopoly. *J. Political Econom.* **72**(1) 44–61.

Stokes, R. 2010. *Ultimate Guide to Pay-Per-Click Advertising*. Entrepreneur Press, Newburgh, NY.

Varian, H. R. 1980. A model of sales. *Amer. Econom. Rev.* **7**(4) 651–659.

Varian, H. R. 2007. Position auctions. *Internat. J. Indust. Organ.* **25**(6) 1163–1178.

Wilbur, K., Y. Zhu. 2009. Click fraud. *Marketing Sci.* **28**(2) 293–308.

Yang, S., A. Ghose. 2009. Analyzing the relationship between organic and sponsored search advertising: Positive, negative, or zero interdependence? *Marketing Sci.* **29**(4) 602–623.

Zhang, X. M., J. Feng. 2005. Price cycles in online advertising auctions. *Internat*. *Conf*. *Inform*. *Systems* (*ICIS*), *Las Vegas*, *NV*, 769–781.