PRODUCTION AND OPERATIONS MANAGEMENT

Vol. 19, No. 3, May–June 2010, pp. 305–321
ISSN 1059-1478 | EISSN 1937-5956 | 10 | 1903 | 0305

DOI 10.3401/poms.1080.01093

Keyword Auctions, Unit-Price Contracts, and the Role of Commitment

Jianqing Chen
Haskayne School of Business, University of Calgary, Calgary, AB T2N 1N4, Canada, jiachen@ucalgary.ca

Juan Feng
College of Business, City University of Hong Kong, Hong Kong, China, juafeng@cityu.edu.hk; Warrington College of Business, University of Florida, Gainesville, FL 32611, jane.feng@cba.ufl.edu

Andrew B. Whinston
Department of Information, Risk, and Operations Management, McCombs School of Business, University of Texas at Austin, Austin, Texas 78712, abw@uts.cc.utexas.edu

Motivated by the enormous growth of keyword advertising, this paper explores the design of performance-based unit-price contract auctions, in which bidders bid their unit prices and the winner is chosen based on both their bids and performance levels. The previous literature on unit-price contract auctions usually considers a static case where bidders' performance levels are fixed. This paper studies a dynamic setting in which bidders with a low performance level can improve their performance at a certain cost. We examine the effect of the performance-based allocation on overall bidder performance, auction efficiency, and the auctioneer's revenue, and derive the revenue-maximizing and efficient policies accordingly. Moreover, the possible upgrade in bidders' performance level gives the auctioneer an incentive to modify the auction rules over time, as is confirmed by the practice of Yahoo! and Google. We thus compare the auctioneer's revenue-maximizing policies when she is fully committed to the auction rule and when she is not, and show that the auctioneer should give less preferential treatment to low-performance bidders when she is fully committed.

Key words: performance-based pricing; unit-price auctions; keyword auctions; commitment
History: Received: October 2007; Accepted: June 2009, after 3 revisions.

1. Introduction

Unit-price contracts (UPCs) are widely used in competitive procurement auctions, such as highway contracting (Stark 1974) and national defense (Samuelson 1983, 1986). In such auctions, bidders submit their bids specifying the unit price for each of the input factors. The auctioneer then calculates a score for each bidder based on both the unit prices and the expected quantities needed. The bidder with the lowest score wins the auction. The final payment to the winner, however, is determined by the number of units that are finally consumed or needed in realization. Moreover, many other regular settings, such as the marketing of publishing rights for books (McAfee and McMillan 1986), can also be interpreted as UPC auctions.

In recent years, various formats of UPC auctions have been adopted by major search engines, such as Yahoo!, Google, and MSN, to sell keyword advertising slots on their websites. Keyword advertising is a form of targeted online advertising, in which advertisements are selectively displayed on web pages related to certain keywords. One popular form of keyword advertising is search-based advertising (also referred to as "sponsored search"). For example, when users search for the keyword "laptop" on a search engine, besides the organic results generated by its technical algorithm, the search engine may also display advertisements from the advertisers who are interested in showing up when "laptop" is searched. These keyword-related advertising slots are usually sold through UPC auctions. In these auctions, advertisers specify their per-click willingness to pay (and pay only if consumers click their advertisements), and the final payments to the search engine are determined by the actual amount of click-throughs that their advertisement generates. Keyword advertising has become a leading form of online advertising, accounting for 45 percent of the total online advertising revenue in 2008 (IAB 2008).

Previous literature on UPC procurement auctions reveals that to promote competition, it is beneficial for the auctioneer to give preferential treatment to bidders with a lower efficiency/performance level (McAfee and McMillan 1986, Ewerhart and Fieseler 2003, Rothkopf and Whinston 2007). In practice, UPC auctions often offer a subsidy to contractors with inferior production technologies or a lower performance level (Ewerhart and Fieseler 2003). Preferential government procurement policies also affect several hundred billion dollars' worth of trade worldwide each year (Graham 1983). For example, the U.S. government offers a 50 percent preference for domestic suppliers for military procurement (McAfee and McMillan 1989). These practices are interpreted as either a protectionist device or a way to increase bidding competition (McAfee and McMillan 1986).

There is also a fast-growing literature studying the "keyword auctions," which is one of the most popular places to apply the UPC auctions on the internet (Liu and Chen 2006, Feng et al. 2007a, b, Weber and Zheng 2007, Liu et al. 2008, Chen et al. 2009). This stream of literature usually focuses on the effect of various auction mechanisms on the advertisers' bidding behavior and the auctioneer's revenue. For example, Feng et al. (2007a) find that depending on the correlation between advertisers' willingness to pay and their relevance, Google's (without preferential treatment) and Yahoo's (with preferential treatment) ranking mechanism can outperform each other under different conditions. Weber and Zheng (2007) study a search model in which advertisers compete for positions in a search engine and show that the optimal winning rule should put non-zero weight on the advertisers' bids. Liu and Chen (2006) and Liu et al. (2008) consider a static weighted unit-price auction where bidders bid on unit prices, and the winner is determined by their bids as well as their past performance. This paper extends Liu and Chen's model by considering the dynamic effects of bidder performance evolution.

The extension is motivated as follows. Previous work on keyword auctions has assumed that bidders' performance levels are exogenously given and beyond firms' control, and does not consider the possible effect of the auction mechanism on bidders' performance level in the long term. It is not obvious, however, that the practice of preferential treatment works when bidders' performance may change over time for various reasons. For instance, advertisers may improve their performance for their own interest or in reaction to the performance-based ranking policy. In general, the possible upgrade in performance by some advertisers changes the distribution of performance levels among all advertisers. This change gives the auctioneer an incentive to adjust her preferential policy in time accordingly. In addition, preferential treatment may discourage low-performance bidders from investing in their performance levels, which could negatively affect the auctioneer's revenue. Moreover, under the preferential treatment policy, a less efficient bidder beats more efficient ones with positive probability (Ewerhart and Fieseler 2003). This efficiency loss, especially in the long term, might be significantly detrimental for an industry that is sensitive to its consumer response. As in the case of keyword advertising, the consumer base is crucial for the long-term revenue of a search engine.

In practice, the performance-based ranking mechanisms adopted by the search engines are constantly changing over time, which exhibits an experimental process and may reflect the need to adjust auction policies in light of dynamic features. For example, Yahoo! used to rank advertisers by their per-click willingness-to-pay and has now switched to a new mechanism that also considers click-through rates in its ranking. Google first introduced a design that ranks advertisers by the product of their per-click bidding prices and their historical click-through rates in 2002; later, it extended historical click-through rates to more comprehensive "Quality Scores" as the ranking factor, which also takes into account the quality of the advertisement text and other unannounced relevance factors. In addition, instead of announcing exactly how their ranking mechanisms incorporate the performance information, as they did earlier, many search engines now keep this ambiguous.

These observations indicate that it is not sufficient to study UPC auctions in a static setting, where the firms' performance levels remain unchanged. On the one hand, the UPC mechanism gives bidders with different performance levels different incentives to improve their performance, and this in turn affects the long-term overall bidder performance level. On the other hand, the updated distribution of firms' performance levels gives the auctioneer an incentive to modify her auction mechanism, so the auctioneer faces a commitment problem. What is the effect of the preferential ranking rules on bidders' performance choice? How should the auctioneer choose the performance-based allocation policy considering bidders' possible performance upgrade? Should the auctioneer commit to a certain mechanism or is it beneficial to modify its auction mechanism in time?

This paper tries to answer these questions. We consider a two-period model to capture the dynamic effects of bidder performance evolution. The winner of the auction is chosen based on both bidders' bids and performance levels in the second period. In the first period, the bidders may have low or high performance, and the low-performance bidders can invest in improving their performance level.

Under this performance-based unit-price auction framework, we study bidders' decisions on performance improvement, and then examine the effect of performance-based allocation on the overall bidder performance level, the auction efficiency, and the auctioneer's revenue. We find that in equilibrium low-performance bidders above a certain valuation level invest in their performance and the others do not. The overall bidder performance level is monotonic in the degree of preferential treatment given to those bidders with low performance: The more the auctioneer discriminates against low-performance bidders, the higher is the overall performance level. The efficient policy involves weighting bidders' unit-price bids by their expected performance, which happens to coincide with the efficient policy in a static case (where the performance level is fixed). This coincidence is surprising in that the efficiency in our dynamic setting concerns both allocative efficiency and investment efficiency, whereas the static one considers allocative efficiency only. We also compare the auctioneer's revenue-maximizing policies when she is fully committed to the auction rule and when she is not, and we show that she should give less preferential treatment to low-performance bidders when she is fully committed.

This paper also relates to literature on investment incentives in procurement auctions (Tan 1992, Arozamena and Cantillon 2004). Most of these studies concern the revenue equivalence under different auction formats (e.g., first-price sealed-bid auctions versus second-price auctions), while our paper focuses on performance-based unit-price auctions in the context of keyword advertising, and examines the effect and the design of the performance-based allocation policy. To our knowledge, this setup is mostly related to Branco (2002), which studies a procurement auction where two firms compete for a government project and the inefficient firm may improve its technology. In his setting, the technology choice is unobservable. In contrast, we study the case where the performance is observable and thus performance-based allocation is feasible, and we obtain different insights.

This paper is organized as follows. We describe our model in section 2. In section 3, we investigate bidders' bidding decisions and their performance evolution. In section 4, we study the effect of performance-based allocation on bidders' overall performance. We discuss social welfare in section 5 and explore the revenue-maximizing allocation for an auctioneer in section 6. Section 7 concludes the paper.

2. Model Setup

We consider an environment where the auctioneer sells a single object to n risk-neutral bidders, and the bidders have independent and private values for the object being sold. The object can represent a type of service (e.g., providing exposure for a company in the keyword auction) or a physical asset for rent, among others. We capture the long-term interactions between the auction mechanism and bidders' overall performance levels through a two-period model. In the first period, the auctioneer announces the auction mechanism, and the bidders decide whether to invest in improving their performance levels. In the second period, the auctioneer can either keep or modify the announced auction rules, and the bidders participate in the auction. We refer to the former as a full-commitment case and to the latter as a limited-commitment case.

Assume that each bidder is characterized by a performance level y and unit valuation v. A performance level y measures a bidder's productivity, yield, or efficiency of using the object being auctioned, and the unit valuation v is the bidder's valuation for each unit of y. In keyword advertising, the performance level measures the expected number of clicks that an advertiser can generate during a given period of time, and v measures the advertiser's per-click valuation. So a bidder with performance level y and unit valuation v has a total valuation of vy. For simplicity, we assume that bidders' performance levels are independent of their unit valuations. Therefore, a bidder with a higher performance level is stochastically more efficient than a bidder with a lower performance level.

For simplicity, we assume that the bidders can have either high or low performance; that is, the performance level y takes a discrete value: $y \in \{y_L, y_H\}$, where $y_H > y_L > 0$. A bidder with a low performance level (y_L) can invest in and improve his performance level to y_H at a cost c, where $c \geq 0$. In keyword advertising, for example, advertisers may engage in extensive marketing research or experiments to improve their website design and thus ultimately increase their click-through rates (Schlosser et al. 2006).

We assume that bidders receive their respective two-dimensional information at the beginning of the first period (v and y), but not that of others. The unit valuation v is independent and identically distributed on the interval $[0, 1]$, following a density function $f(v)$ that is positive and continuous everywhere on the interval. Let $F(v)$ denote the corresponding cumulative distribution function. By the convention, $F(v) = 1$ for $v > 1$. A bidder's performance level y in the first period is believed to be y_H with probability α, and y_L with probability $1 - \alpha$. Both $F(v)$ and α are commonly known. While it remains unknown to other bidders, each bidder's performance level is known to or can be verified by the auctioneer in the second period once the auction takes place. For example, in keyword auctions, the number of click-throughs that a certain advertisement generates is not observed by other advertisers but can be approximately predicted (based on past performance) and accurately recorded by search engines.

The object is sold through a first-price, sealed-bid unit-price auction. Each bidder places a unit-price bid b, $b \geq 0$. The auctioneer uses a scoring rule to evaluate the bids from bidders with different performance levels, as their performance levels can be observed and verified by the auctioneer. In particular, a weighting factor $w \in (0, 1]$ on the low-performance bidders' bids is introduced to measure this differential treatment, and a score is calculated for a bidder of performance level y according to the following formula:

$$s(y, b) = \begin{cases} b & \text{if } y = y_H, \\ wb & \text{if } y = y_L. \end{cases} \tag{1}$$

The bidder with the highest score wins and pays for all the realized performance at his unit-price bid b. There is no entry fee or reserve price.

We assume that a bidder's payoff from winning the object is additive in its total valuation for the object, yv, and its payment, yb. So the expected payoff for a bidder of performance level y and unit valuation v, who places a bid b is

$$U(y, v, b) = y(v - b) \Pr(\text{win}|y, b). \tag{2}$$

By allowing w to take different values, we can accommodate different auction formats. When $w = \frac{y_L}{y_H}$, bidders with low performance level and high performance level are treated "fairly," as the winner is chosen based on the ranking of the total revenue created. The original Google auction belongs to this category. When $w \in (0, y_L/y_H)$, the bid submitted by a bidder with a low performance level is treated unfavorably. When $w \in (y_L/y_H, 1]$, a bidder with a low performance level is treated favorably. Yahoo!'s earlier auction format determined the winner solely by the bidding amount, which is equivalent to the case where $w = 1$. We denote the case of $w = 1$ as a *standard unit-price auction* studied in literature, and the case of $w < 1$ as a *performance-based auction*.

We exclude the case of $w \leq 0$ because it implies that the low-performance bidders can never win over the high-performance ones, so that the model reduces to a trivial case with high-performance bidders competing with each other. We assume w is no greater than 1 for three reasons. First, as we shall show, any $w > 1$ can never be socially efficient. Social efficiency or the closeness to it is sometimes the auctioneer's concern for the sake of, for example, competition among auctioneers (Stegeman 1996). Second, $w > 1$ is practically non-sensible, because it lets a low-performance bidder always beat a high-performance bidder if they submit the same bid. Such an auction rule would be perceived to be very unfair and foster low performance (as some high-performance bidders may even wish to convert to low performance for the favorable auction rule). Third, auctions with $w > 1$ are seldom observed in reality. The case with $w \in (0, 1]$ can decently encompass the most popular ranking rules (such as Google's and earlier Yahoo!'s) without the perplexing technical details.

The timing is as follows. In the first period, the auctioneer announces a weighting factor w. Then bidders privately learn their performance levels and valuations. Bidders with a low performance level decide whether to improve their performance, without seeing the decisions of others. In the second period, the auctioneer may follow the preannounced auction rule (in the full-commitment case) or modify the auction rule (in the limited-commitment case), and then all bidders participate in the auction. Upon their (updated) performance, bidders' scores are calculated according to the scoring rule (1). The object is assigned to the bidder with the highest score, and payment is made to the auctioneer.

If the auctioneer always kept the preannounced auction rules and was believed to do so, the commitment issue could also be understood as the order of moves. In the full-commitment case, the auctioneer chooses auction rules first and then bidders make their investment decisions. In the limited-commitment case, the auctioneer and the bidders make their respective decisions simultaneously (so the bidders do not observe the chosen auction rules).

We are interested in the effects of the performance-based allocation on bidders' performance evolution, the overall bidder performance level, the resulting social welfare, and the auctioneer's expected revenue. The online supplement provides discussion on relaxing some model assumptions.

3. Bidding Strategy and Performance Conversion

In this section, we focus on the performance decisions of bidders with a low performance level. We conjecture that, in equilibrium, there exists a cutoff value v^* such that all low-performance bidders of unit valuations higher than v^* will improve their performance level and all low-performance bidders of lower unit valuations will remain at the low performance level. This conjecture is intuitive in that, in general, upon winning the object, the low-performance bidders with higher unit valuations gain more from improving their performance ($(y_H - y_L)v$) than those with lower unit valuations. As a result, low-performance bidders with higher valuations should be more likely to invest in their performance. We will verify later that this is an equilibrium strategy.

We start with the second period, where bidders participate in the auction with a given weighting factor w.

3.1. Second Period: The Bidding Strategy for a Given w

We denote P_H (P_L) as the probability that a bidder has a high (low) performance level in the second period.

According to the conjecture, a bidder has a low performance level in the second period only if he has a low performance in the first period (with probability $(1-\alpha)$) and remains at it (if his unit valuation is less than v^*). Therefore, in the second period, a bidder with probability $(1-\alpha)F(v^*)$ has a low performance level and with probability $[1-(1-\alpha)F(v^*)]$ has a high performance level; that is

$$P_L = (1-\alpha)F(v^*) \text{ and } P_H = 1-(1-\alpha)F(v^*). \quad (3)$$

We denote $F_H(v)$ as the distribution function of unit valuations for bidders with a high performance level in the second period, and $F_L(v)$ as that for bidders with a low performance level. $F_H(v)$ and $F_L(v)$ can be calculated by applying Bayes' rule (here we use $\Pr(\cdot)$ to represent the probability of an event in the second period):

$$F_L(v) = \frac{\Pr(V \le v, y = y_L)}{P_L} = \begin{cases} \dfrac{F(v)}{F(v^*)} & \text{if } v \le v^*, \\ 1 & \text{if } v > v^*, \end{cases} \quad (4)$$

and

$$F_H(v) = \frac{\Pr(V \le v, y = y_H)}{P_H} = \begin{cases} \dfrac{\alpha F(v)}{1-(1-\alpha)F(v^*)} & \text{if } v \le v^*, \\ \dfrac{F(v)-(1-\alpha)F(v^*)}{1-(1-\alpha)F(v^*)} & \text{if } v > v^*. \end{cases} \quad (5)$$

Equation (4) is due to that a bidder in the second period has a low performance level and unit valuation less than v only if in the first period he had a low performance and unit valuation less than v (which occurs with probability $(1-\alpha)F(v)$). Similarly, with probability $\alpha F(v)$, a bidder in the second period is of a high performance level and unit valuation less than v (if $v \le v^*$). For the case where $v > v^*$, the high-performance bidders in the second period also include bidders converted from a low performance level, which happens with probability $(1-\alpha)$ $(F(v)-F(v^*))$. Therefore, with probability $\alpha F(v)+(1-\alpha)$ $(F(v)-F(v^*))$, a bidder is of a high performance level with $v > v^*$ in the second period. This accounts for (5). We denote the corresponding density functions as $f_L(v)$ and $f_H(v)$ (specified in the Appendix).

We consider a symmetric, pure-strategy, perfect Bayesian Nash equilibrium (Fudenberg and Tirole 1991). By "symmetric," we mean that bidders with the same unit valuation and performance level will bid the same amount in equilibrium. Let $b_H(v)$ and $b_L(v)$ denote the bidding functions for bidders with a high and low performance level in the second period, respectively. We conjecture that the bidding functions are strictly increasing in bidders' unit valuations. (We will verify this later.) We obtain the following result regarding the bidding functions, which is similar to that in Liu and Chen (2006).

Lemma 1. *If any bidder of (v, y) has a unique optimal bid in equilibrium, then $b_H(wv) = wb_L(v)$ for $v \in [0, v^*]$.*

Proof. All proofs of lemmas and corollary are deferred to the online supplement. All proofs of propositions are deferred to the Appendix. □

Lemma 1 shows that, in equilibrium, a high-performance bidder of unit valuation wv and a low-performance bidder of unit valuation v will obtain the same score (recall the scoring rule [1]) and thus the same probability of winning. The intuition is as follows. Consider a high-performance bidder of unit valuation wv who bids wb and a low-performance bidder of unit valuation v who bids b. According to the scoring rule, the former has the same probability of winning as the latter. Then, the expected payoff of the former differs from that of the latter only by a scalar according to (2). Because multiplying a payoff function by a positive scalar does not alter the solution to an optimization problem, b maximizes the low-performance bidder's expected payoff if and only if wb maximizes the high-performance bidder's expected payoff.

By the monotonicity of bidding functions, a low-performance bidder of unit valuation v can beat a bidder j, if and only if j has a low performance level and unit valuation less than v, or j has a high performance level and a unit valuation less than wv (by Lemma 1). So the probability of winning for a low-performance bidder of unit valuation v can be represented by

$$\rho_L(v) \equiv [P_H F_H(wv) + P_L F_L(v)]^{n-1}. \quad (6)$$

Similarly,

$$\rho_H(v) \equiv \left[P_H F_H(v) + P_L F_L\left(\frac{v}{w}\right)\right]^{n-1}. \quad (7)$$

Lemma 2 presents the equilibrium bidding functions. To be rigorous, we let $b_H(v)|_{v=0} \equiv \lim_{v\to 0^+} b_H(v) = 0$ and $b_L(v)|_{v=0} \equiv \lim_{v\to 0^+} b_L(v) = 0$, respectively.

Lemma 2. *Given w $(0 < w \le 1)$ and the cutoff value v^*, the equilibrium bidding functions for bidders with low and high performance level are increasing in v, and can be represented as the following:*

$$\begin{cases} b_H(v) = v - \dfrac{\int_0^v \rho_H(t)dt}{\rho_H(v)} \text{ for } v \in [0,1], \\ b_L(v) = v - \dfrac{\int_0^v \rho_L(t)dt}{\rho_L(v)} \text{ for } v \in [0, v^*]. \end{cases} \quad (8)$$

It is easy to verify that the above bidding functions are indeed strictly increasing. (See the online supplement for the proof.) Also, as indicated by (8), a bidder's equilibrium unit-price bid is always less than his true unit valuation, which is common among first-price auctions. It is worth pointing out that in our setting the *symmetric* bidding function is unique, as in standard auctions. This is because the auction in our setting can be mapped to a standard first-price auction. We take a high-performance bidder as an example. The equilibrium tradeoff for the bidder is the same as if he were competing with other $n-1$ high-performance bidders, with the unit valuations following the cumulative distribution function $\left[P_H F_H(v) + P_L F_L(\frac{v}{w})\right]$ on the interval [0, 1]. As the latter scenario is in fact a standard independent-private-value auction and the equilibrium is unique (Lebrun 1999), so is the bidding function for high-performance bidders in our setting. The uniqueness of the bidding function for low-performance bidders follows due to the established mapping in Lemma 1 (i.e., $b_L(v) = b_H(wv)/w$).

Denote $V_H(v) \equiv U(y_H, v, b_H(v))$ as the equilibrium expected payoff of a bidder with performance level y_H and unit valuation v. Following standard derivation,

$$V_H(v) = y_H(v - b_H(v))\rho_H(v) = y_H \int_0^v \rho_H(t)dt, \qquad (9)$$

where the first equality is by noticing $\Pr(\text{win}|y_H, b_H(v)) = \rho_H(v)$ in payoff function (2) and the second equality is by substituting in the bidding functions (8). Similarly, $V_L(v) \equiv U(y_L, v, b_L(v)) = y_L \int_0^v \rho_L(t)dt$. Note that $\rho_H(v)$ and $\rho_L(v)$ represent the equilibrium probabilities of winning, so the above expressions of equilibrium expected payoffs $V_H(v)$ and $V_L(v)$ are consistent with the standard auction theory.

3.2. First Period: Performance Conversion for a Given w

Given the equilibrium bidding functions in the second period, we now solve for low-performance bidders' decisions on the performance upgrades in the first period. We show that such a cutoff value v^* does exist and is unique.

A bidder with a low performance level benefits from converting to a high performance level because of the increase in both his performance level and his probability of winning. The overall benefit is represented by the increase in the equilibrium payoffs from converting to a high performance level, that is, $V_H(v) - V_L(v)$. However, the investment in the performance incurs a cost c. Given w, whether a bidder with a low performance level converts depends on this tradeoff.

PROPOSITION 1. *Given w ($0 < w \leq 1$) and investment cost c, there exists an equilibrium cutoff value v^* such that a low-performance bidder with $v \geq v^*$ converts to high performance and a low-performance bidder with $v < v^*$ does not, where v^* is uniquely determined by*

$$v^* = \begin{cases} 0 & \text{if } c = 0, \\ \text{solution to } \Delta V(v^*) = c & \text{if } 0 < c < \Delta V(1), \\ 1 & \text{if } c \geq \Delta V(1), \end{cases} \qquad (10)$$

where

$$\Delta V(v^*) \equiv y_H \int_0^{v^*} \rho_H(t)dt - y_L \int_0^{v^*} \rho_L(t)dt. \qquad (11)$$

Intuitively, if a low-performance bidder can be better off converting to high performance, it is also true for a low-performance bidder of a higher unit valuation. This is because, on the one hand, a higher valuation bidder not only obtains more surplus than a lower valuation one but also gains more from improving performance. On the other hand, all bidders with a low performance level face the same investment cost. Therefore, low-performance bidders of higher valuations have greater incentives to improve their performance, which ensures that the performance choice pattern we conjecture is the only possible one.

Note that, given the equilibrium cutoff point v^*, $\Delta V(v^*)$ represents the benefit of improving performance for the low-performance bidder of unit valuation v^* (i.e., the marginal bidder). In equilibrium, the marginal bidder must be indifferent between converting and not converting, which explains $\Delta V(v^*) = c$. The uniqueness of the solution to $\Delta V(v^*) = c$ is due to the monotonicity of $\Delta V(v^*)$: As v^* increases, by (11) the benefit of improving performance for the marginal type also increases.

When it is cost-free for the low-performance bidders to convert, all low-performance bidders will convert to high performance (i.e., $v^* = 0$). Consequently, in the second period, all bidders in the auction have a high performance level. In such a case, the performance-based auction is identical to a standard unit-price auction, and the problem becomes straightforward. It is worth noting that as long as the performance improvement is associated with a positive cost, at least some low-performance bidders (of low valuations) will remain at a low performance because their benefit from conversion is too limited to compensate for the cost of conversion.

When the investment cost is prohibitive, the gain from improving performance cannot compensate for the cost, and no bidders with a low performance level will convert (i.e., $v^* = 1$). More precisely, when the investment cost is higher than $\Delta V(1)$ (the gain that the low-performance bidder with $v = 1$ can get from the conversion), no low-performance bidders will convert. By (11), the cutoff cost $\Delta V(1)$ increases in the high-performance value y_H and decreases in the low-performance value y_L and in the weighting factor w. Intuitively, as y_H increases, for example, the benefit

Figure 1 The Equilibrium Bidding Functions and Equilibrium Scores

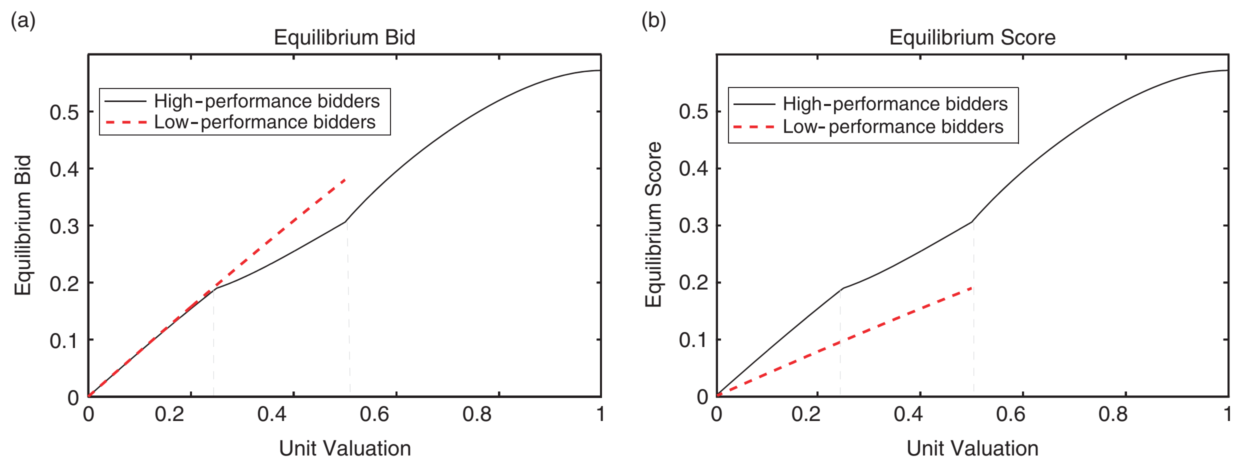

from converting to high performance increases, which results in a higher cutoff cost.

For the rest of this paper, we focus on the more interesting case, where the investment cost is modest and at least some bidders with a low performance level are interested in the conversion. We next use an example to illustrate the equilibrium performance choice and bidding functions.

EXAMPLE 1. *Let $\alpha = 0.5$, $y_L = 0.5$, $y_H = 1$, $n = 5$, $w = 0.5$, $F(v) = 2v - v^2$, and $c = 0.054$. By (10), we can derive $v^* = 0.5$, which means low-performance bidders of valuations higher than 0.5 convert to high performance and the other low-performance bidders do not. Figure 1 illustrates the equilibrium bidding functions and equilibrium scores. (By the scoring rule, the equilibrium scores for the high-performance bidders are their equilibrium bids and for the low-performance bidders are their equilibrium bids adjusted by the weighting factor w.)*

It is worth mentioning that there are two kinks (wv^* and v^*) in the bidding function of the high-performance bidders. This is due to the kinks in the probability of winning $\rho_H(t)$ (also see [A3] in the Appendix), as the competition situation faced by high-performance bidders is discontinuous. In the example above, the first kink in the high-performance bidding function is due to $F_L\left(\frac{t}{w}\right) = 1$ for all $t > 0.25$, and the second is due to the kink of $F_H(t)$ at $t = 0.5$. Intuitively, it is still possible for a high-performance bidder with a relatively low unit valuation ($v \leq 0.25$) to lose to a low-performance bidder, even though the bid from a low-performance bidder gets discounted by w. Once the unit valuation from a high-performance bidder exceeds a certain value (0.25), he can always beat a bidder with a low performance level. Therefore, there exists a kink at $v = 0.25$. Similarly, a high-performance bidder with an intermediate unit valuation ($0.25 < v \leq 0.5$) faces different competition pressure from a high-performance bidder with a higher valuation ($v > 0.5$), because the latter can be either a bidder who has a high performance level originally, or one who has converted from a low performance level. Therefore, the second kink happens at $v = 0.5$.

In general, the conversion decision is based on the tradeoff between the gain from conversion and the investment cost. Everything else being equal, a higher cost results in fewer low-performance bidder conversions. Similarly, the difference between the two performance levels affects the conversion decision for those bidders with a low performance level: Given y_L, a higher y_H, which means more gain from conversion, could induce more bidders with a low performance level to convert. As a summary,

COROLLARY 1. *For a given w, v^* increases in the investment cost (c) and decreases in the high performance level (y_H).*

It is worth noting that the equilibrium bids of the bidders with different valuations are affected differently by the change in the investment cost c. This is because, as indicated in Corollary 1, the change in the investment cost changes the equilibrium cutoff point (v^*), which changes the competition intensity faced by certain bidders and thus affects their bidding behavior. For example, Figure 2 plots the change of high-performance bidders' equilibrium bids when the investment cost in Example 1 is doubled. As the investment cost increases, fewer low-performance bidders with high valuations convert, which decreases the competition pressure for the high-performance bidders with high valuations. As a result, those high-performance bidders with high valuations bid less aggressively. The high-performance bidders with intermediate valuations bid more aggressively because some of them now face direct competition pressure from their low-perfor-

Figure 2 The Effect of the Investment Cost on Bidding

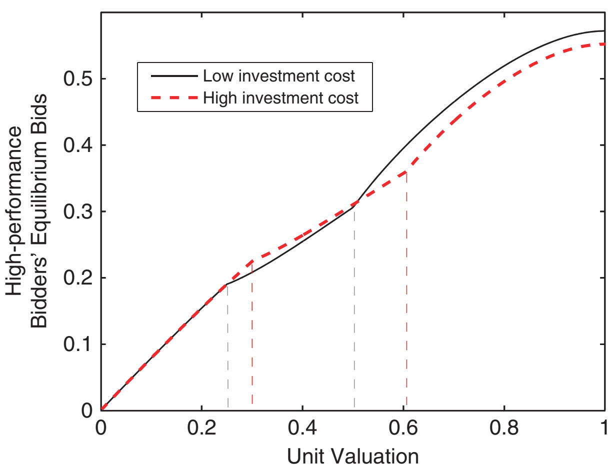

mance counterparts, which increases the competition intensity. The high-performance bidders with low valuations bid the same because they cannot win against high-valuation bidders anyway and the competition among the low-valuation bidders remains the same.

4. Bidders' Overall Performance

In this section, we examine the effect of differential treatment (w) on the winning bidder's overall performance level. First, we observe that the cutoff value is closely related to the weighting factor.

LEMMA 3. *The cutoff value v^* is increasing in the weighting factor w.*

Intuitively, as w increases, and hence the low-performance bidders are treated more favorably, bidders' benefit from improving their performance level declines. As a result, fewer low-performance bidders have incentive to invest in their performance level. In other words, the more the low-performance bidders are discriminated against in the auction (i.e., the smaller the w), the larger the incentive for the low-performance bidders to improve their performance level, because otherwise they have little chance to win.

Next, we explore the impact of the weighting factor on the expected winning performance level. The expected winning performance level is determined by the bidders' performance weighted by their winning probabilities, which can be expressed as

$$y_L P_L \int_0^{v^*} \rho_L(v) f_L(v) dv + y_H P_H \int_0^1 \rho_H(v) f_H(v) dv. \quad (12)$$

PROPOSITION 2. *The more the low-performance bidders are discriminated against, the higher the overall performance level. That is, the expected winning performance level is decreasing in w.*

To show this, from Lemma 3, we know that as w decreases, more low-performance bidders improve their performance level. Moreover, as it is harder for low-performance bidders to beat high-performance ones, the winning probabilities for low-performance bidders decrease (indicated by [6] or implied by Lemma 1). Both factors drive the expected winning performance level in the same direction: the lower the w, the higher the expected winning performance level. As a result, the performance-based auctions with $w \in (0,1)$, in general, enhance the expected winning performance level compared with the standard UPC auctions (where $w = 1$).

5. Auction Efficiency

In this section, we study the efficiency of the performance-based UPC auctions. We refer to social welfare as the total expected payoff of the auctioneer and bidders, or the total expected value created through the auction. The efficient weighting factor w_{eff} is defined as the one that maximizes social welfare, or the one that would be chosen by a social planner. Efficiency may be in the auctioneer's best interest in the long run. In keyword auctions, especially, the market for keyword advertising is relatively nascent, and using a UPC auction mechanism is still at an early stage. It is sensible for the auctioneer (i.e., the search engine) to choose an efficient design at this stage to maximize the "total pie." After all, unless advertisers feel that they get fair treatment in the auctions and see high returns, they may spend less in keyword advertising or may not even return for more business.

The social welfare can be expressed as the total expected valuation of the auctioneer and bidders minus the expected investment cost:

$$n\left[y_L P_L \int_0^{v^*} v\rho_L(v) f_L(v) dv + y_H P_H \int_0^1 v\rho_H(v) f_H(v) dv \right.$$
$$\left. -(1-\alpha)(1-F(v^*))c \right]. \quad (13)$$

So, w_{eff} maximizes (13).

In a standard static auction, an efficient policy considers allocative efficiency only. An efficient auction allocates the object to the bidder with the highest total valuation. In our dynamic setup, efficiency refers to both the allocative efficiency and the investment efficiency (whether there exists over-investment or under-investment in performance). As a result, the efficiency problem becomes more intricate when bidders can update their performance levels. Surprisingly, the prescribed policy is rather simple, as summarized below.

PROPOSITION 3. *The efficient weighting factor $w_{\mathit{eff}} = \frac{y_L}{y_H}$.*

Proposition 3 shows that the efficient policy in this two-period game involves weighting bidders' unit-price bids by their expected performance, which happens to coincide with the policy in the static setting (Liu and Chen 2006). The intuition for this surprising coincidence is as follows.

First, the weighting factor $w = y_L/y_H$ ensures the allocative efficiency in the second period. As we mentioned earlier, when $w = y_L/y_H$, bidders with low and high performance levels are treated "fairly" in the sense that the winner is chosen based on the ranking of their total payment. Therefore, this policy is equivalent to the standard first-price auction where bidders submit their total payments for the object, and is thus equivalent to the standard second-price auction by the Revenue Equivalence Theorem (Myerson 1981). As allocative efficiency is guaranteed in standard second-price auctions (since bidders bid their true value and the object is assigned to the bidder with the highest valuation in equilibrium), allocative efficiency is also guaranteed here.

Second, the weighting factor $w = y_L/y_H$ also leads to an efficient conversion. In general, whether an investment is *ex ante* efficient depends on the tradeoff between the social gain (increase in the expected social welfare) and the investment cost. Notice that a standard second-price auction corresponds to the Vickrey–Clarke–Groves (VCG) mechanism, for which a bidder's expected payoff equals the improvement of the social welfare brought in. The social gain from a bidder's conversion in the auction with $w = y_L/y_H$ here (equivalent to a standard first-price auction) equals the increase of the bidder's payoff $[V_H(v) - V_L(v)]$. Therefore, the bidder's decision on performance investment (balancing the payoff increase and investment cost) is aligned to the social planner's decision (balancing the social gain and investment cost). For example, a low-performance bidder with $v > v^*$ chooses to convert because the increase of the payoff is greater than the investment cost, which also ensures that the social gain associated with the conversion is greater than the cost and thus the conversion efficiency. Such a property of standard second-price auctions has been discussed under different settings. For instance Stegeman (1996) and Arozamena and Cantillon (2004) consider models that allow bidders to invest in information about their own valuations and to invest in cost reduction in procurement auctions, respectively. Both find that, under second-price auctions, bidders' optimization problems are aligned to the social planner's, which leads to efficient investment. So $w = y_L/y_H$ leads to investment efficiency.

For these reasons, we define the performance-based auction with $w = y_L/y_H$ as an *efficient performance-based auction*. In an efficient performance-based auction, the object is properly assigned to the bidder with the highest valuation in the second period. In the first period, low-performance bidders of unit valuations higher than the cutoff valuation (defined by [10] with $w = y_L/y_H$) are efficiently induced to convert to high performance. The following is an example of efficient performance-based auctions.

EXAMPLE 2. *Continue with Example 1. Notice that we used $w = 0.5$ in Example 1, which is the efficient weighting factor since $y_L/y_H = 0.5$. The allocative efficiency is guaranteed since a high-performance bidder with unit valuation v has the same probability of winning as a low-performance bidder with $2v$ (by Lemma 1 and as depicted in Figure 1(b)) who has the same total valuation as the former. For the investment efficiency, we can verify that the social gain from the conversion of the marginal bidder (i.e., the bidder with $v = 0.5$) breaks even with the investment cost. The low-performance bidders with higher unit valuations bring in higher social gains from their conversions, and those with lower unit valuations would have resulted in lower social gain if converted. So, bidders' conversion decisions are efficient.*

It is worth noting that the above result is independent of whether the auctioneer can fully commit to the preannounced auction rule. This is because, in the second period, it is always efficient to choose $w = y_L/y_H$, regardless of bidders' conversion patterns (as argued in the above). As a result, whether the auctioneer can fully commit does not influence the low-performance bidders' conversion decision in the first period.

6. The Revenue-Maximizing w

In this section, we focus on the weighting factor that maximizes the auctioneer's expected revenue within the class of unit-price auctions and compare the results in the full-commitment case to those in the limited-commitment case. We denote $w^{\mathrm{L}}_{\mathrm{opt}}$ and $w^{\mathrm{F}}_{\mathrm{opt}}$ as the revenue-maximizing weighting factors for the limited-commitment case and the full-commitment case, respectively.

Based on bidders' bidding strategy in the second period, and the performance conversion decisions for low-performance bidders in the first period, for a given w, the auctioneer's expected revenue can be expressed as

$$\pi = ny_LP_L \int_0^{v^*} \rho_L(v)\left(v - \frac{1 - F_L(v)}{f_L(v)}\right)f_L(v)dv + ny_HP_H \int_0^{1} \rho_H(v)\left(v - \frac{1 - F_H(v)}{f_H(v)}\right)f_H(v)dv. \tag{14}$$

(See the online supplement for the derivation.) Basically, the first term is the expected revenue from low-performance bidders in the second period (the *ex post* low-performance bidders), and the second term is the one from the *ex post* high-performance bidders. Denote $J_i(v) \equiv \left[v - \frac{1-F_i(v)}{f_i(v)}\right]$, $i \in \{H, L\}$. In the standard auction literature, $J_i(v)$ is called *virtual value* (Myerson 1981) or *marginal revenue* (Bulow and Roberts 1989), which in our case represents the unit revenue contribution to the auctioneer from a certain bidder. Notice that the revenue contribution by a bidder is different from the actual revenue generated by the same bidder, because revenue contribution also takes into account the impact of this bidder's participation on other bidders.

6.1. Auctioneer with Limited Commitment

Since some higher valued, low-performance bidders improve their performance level in the first period, the distribution of the bidders' performance levels changes at the beginning of the second period. As a result, the original w, which is preannounced at the beginning of the first period, may no longer be optimal, and the auctioneer has incentive to host the auction following a different rule. In practice, both Yahoo! and Google experience a migration of their ranking rules over time. Moreover, neither of them announces the exact w being used, which gives them more flexibility in altering their ranking mechanisms. This corresponds to the case where the auctioneer has limited commitment to the preannounced policy, and the actual auction rule could differ from an earlier announcement (before the low-performance bidders make their conversion decision).

When the auctioneer has limited commitment, the low-performance bidders make their conversion decisions knowing that their conversions will be anticipated and that the weighting factor will be chosen to maximize the auctioneer's revenue based on the *ex post* distribution of firm performance. The revenue-maximizing weighting factor and the cutoff value are determined by the equation system below:

$$\Delta V(v^*) = c \text{ and } \frac{\partial \pi}{\partial w} = 0, \tag{15}$$

where the first equation is the break-even condition for the marginal low-performance bidder, and the second equation comes from the first-order condition of the expected revenue (14) with respect to w. Specifically,

$$\begin{aligned}\frac{\partial \pi}{\partial w} &= n(1-\alpha)y_L \int_0^{v^*} \frac{d\rho_L(v)}{dw} J_L(v) f(v)dv \\ &\quad + n\alpha y_H \int_0^{wv^*} \frac{d\rho_H(v)}{dw} J_H(v) f(v)dv.\end{aligned} \tag{16}$$

Proposition 4 discusses the revenue-maximizing w.

We say a cumulative distribution function $F(v)$ has an *increasing hazard rate* (IHR) if its hazard rate $\frac{f(v)}{1-F(v)}$ increases in v within the support. Many distributions, including uniform, normal, and exponential distributions, satisfy the condition of IHR. Notice that if the distribution $F_i(v)$ satisfies the IHR condition, the virtual value $J_i(v) = \left[v - \frac{1-F_i(v)}{f_i(v)}\right]$ is monotonically increasing: For bidders with the same performance, a bidder with a higher unit valuation contributes more to the auctioneer. Many studies in standard auctions assume monotonic virtual value, which is often called the regularity condition (Myerson 1981).

PROPOSITION 4. *In the limited commitment case, the revenue-maximizing weighting factor* w^{L}_{opt} *is determined by (15). Moreover, if F(v) satisfies the IHR condition, the revenue-maximizing weighting factor is greater than the efficient weighting factor, that is,* $w^{L}_{opt} > \frac{y_L}{y_H}$.

Proposition 4 shows that it is beneficial for the auctioneer to give more preferential treatment to the low-performance bidders rather than simply using the efficient level treatment. It indicates that the recommendation of preferential treatment to low-performance bidders in a static setting also applies to the dynamic setting. This is because, in the absence of commitment, the auctioneer chooses her revenue-maximizing weighting factor based on the *ex post* (after conversion) performance level distribution in the second period only, which is the same as making the choice in a static case.

Favoring bidders with a low performance level has two opposing effects on the auctioneer's expected revenue: First, it raises the winning probabilities for low-performance bidders, which tends to lower the auctioneer's expected revenue because of the efficiency loss; second, it may increase the competitive pressure on some high-performance bidders and thus induce them to bid more aggressively, which tends to enhance the auctioneer's expected revenue. Under the IHR condition, the resolution of such a tradeoff involves preferential treatment to the low-performance bidders. This is because for any weighting factor less than w_{eff}, if the unit valuation distribution satisfies the IHR condition,

$$\begin{aligned} y_L J_L(v) &= y_L\left(v - \frac{F(v^*) - F(v)}{f(v)}\right) \\ &> y_H\left(wv - \frac{\frac{1-(1-\alpha)F(v^*)}{\alpha} - F(wv)}{f(wv)}\right) \\ &= y_H J_H(wv). \end{aligned} \tag{17}$$

(See more details about the above inequality in the Appendix.) That is, the revenue contribution of a low-performance bidder of valuation v is always higher than that of a high-performance bidder of valuation wv, although the two have the same probability of winning. Thus, the auctioneer can always earn a higher revenue by raising w to allocate the object more often to low-performance bidders. Such a result holds regardless of the original performance distribution (α) and the exact performance levels (y_H and y_L).

The revenue-maximizing weighting factor, however, is different from that in a static case. The static case corresponds to a special case of $v^* = 1$ in our setting. Fixing $v^* = 1$ and solving the second equation in (15) give us the revenue-maximizing weighting factor in the static case. Proposition 5 compares the revenue-maximizing w in the limited commitment case to that in the static case.

PROPOSITION 5. *When $v \sim U[0,1]$, in the dynamic case (where bidders' performance levels can be improved) the auctioneer should give more preferential treatment to low-performance bidders than in a static case (where bidders' performance levels are fixed).*

In the above proposition, we use uniform distribution, which allows us to derive a closed-form solution for the revenue-maximizing weighting factor and to conduct further comparisons. However, we conjecture that the result holds under a more general class of distributions (e.g., the distributions with IHR, supported by numerical examples). This is because, after the conversion of some higher valued, low-performance bidders, the remaining bidders with a low performance level are in a more disadvantaged situation. Thus, the auctioneer should give them more preferential treatment to promote competition when the distribution of bidders' unit valuations satisfies the IHR condition.

How does the revenue-maximizing weighting factor affect the equilibrium performance level compared with an efficient weighting factor? Recall that Proposition 2 indicates that a higher weighting factor leads to a lower expected winning performance. So, we have

COROLLARY 2. *If F(v) satisfies the IHR condition, the revenue-maximizing weighting factor results in a higher cutoff v^* and a lower expected winning performance level than the efficient weighting factor does.*

In some instances, the expected winning performance is an important factor for the auctioneer to consider. For example, the performance in keyword auctions measures the relevance of the advertisements to consumers, which could affect consumers' search costs for the information of interest and hence the long-term user base. In this case, subsidizing bidders with a low performance level is not only at the cost of efficiency, as discussed earlier, but also at the cost of lowering the performance; in other words, it is very costly.

6.2. Auctioneer with Full Commitment

When the auctioneer is fully committed to the preannounced auction rules, her decision on w directly affects the performance choices of the low-performance bidders. So the auctioneer has to consider the effect of w on bidders' performance conversion patterns, as well as on the intensity of competition. The effect on competition is similar to that in the limited commitment case. The effect on the performance conversion is reflected by the term of $\frac{dv^*}{dw}$ in the first-order condition below:

$$n\left[(1-\alpha)[(y_L\rho_L(v^*) - y_H\rho_H(v^*))v^* + c]f(v^*).\right.$$
$$\left.+\alpha y_H \int_{wv^*}^{v^*} \frac{d\rho_H(v)}{dv^*} J_H(v) f(v)dv\right]\frac{dv^*(w)}{dw} + \frac{\partial\pi}{\partial w} = 0, \quad (18)$$

where $\frac{\partial\pi}{\partial w}$ is defined as in (16). The revenue-maximizing weighting factor $w_{\mathrm{opt}}^{\mathrm{F}}$ is jointly determined by (10) and (18).

Technically, corner solutions may occur in (18) since we restrict $w \leq 1$. Such a possibility does not affect the results so far. Propositions 1, 2, and 5 are not related to this possibility. Proposition 3 continues to hold, as any $w > 1$ cannot lead to an efficient allocation in the second period and thus cannot be efficient overall. Under some special cases, one condition for Proposition 4 in (15), $\partial\pi/\partial w = 0$, may have a corner solution. However, the insight from Proposition 4 continues to hold (i.e., the revenue-maximizing weighting factor is greater than the efficient weighting factor as $y_L/y_H < 1$). The only result that might be affected is Proposition 6: If (18) has corner solutions, then $w_{\mathrm{opt}}^{\mathrm{F}} = w_{\mathrm{opt}}^{\mathrm{L}} = 1$ and the strict inequality in the proposition should become weak inequality.

Considering the *ex post* competition in the second period, the auctioneer has incentive to subsidize low-performance bidders (by setting $w > y_L/y_H$), which induces some high-performance bidders to bid more aggressively. We call this effect the *competition effect*, which is captured by the term $\frac{\partial\pi}{\partial w}$ in (18). However, lower w, or less subsidy to low-performance bidders, makes more low-performance bidders convert, and this could presumably enhance the auctioneer's revenue. We call this effect the *conversion effect*, which is captured by the term in the first square bracket in (18). In general, the revenue-maximizing weighting factor is determined by the tradeoff between the competition effect and the conversion effect, which depends on the proportion of bidders with low and high performance levels in the distribution.

How does the equilibrium in the full-commitment case differ from that in the limited-commitment case?

PROPOSITION 6. *If the expected revenue in (14) is concave with respect to w (i.e., $\partial^2\pi/\partial w^2 < 0$) and F(v) satisfies the IHR condition, $w_{\mathrm{opt}}^{\mathrm{F}} < w_{\mathrm{opt}}^{\mathrm{L}}$.*

Proposition 6 shows that an auctioneer who can commit to the auction rule should give less preferential treatment to the low-performance bidders than she does when she cannot fully commit. The intuition is as follows. In the full-commitment case, the preferential treatment affects not only the competition

Figure 3 Comparison of the Full-Commitment Case and the Limited-Commitment Case

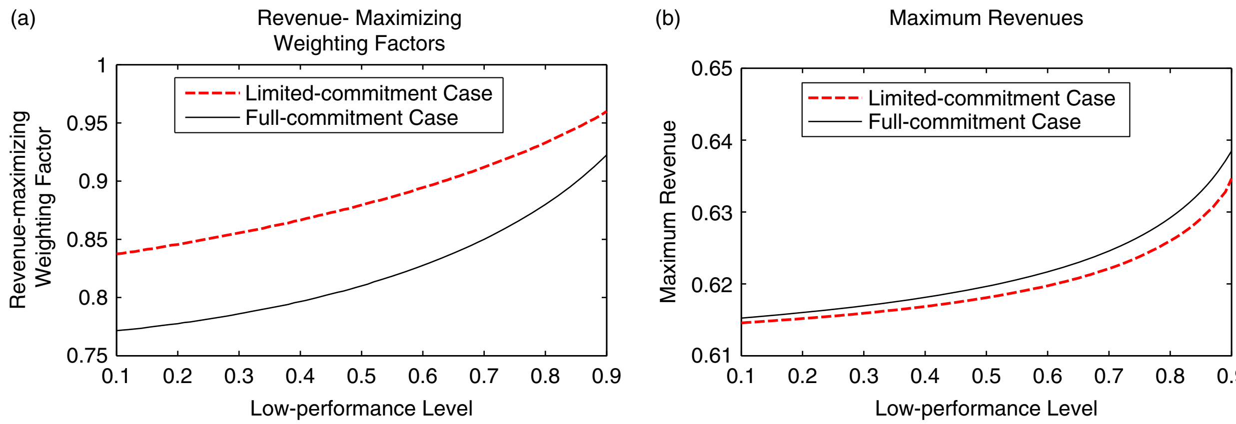

in the second period (as in the limited-commitment case), but also the conversion decisions of the low-performance bidders. A less preferential treatment to low-performance bidders gives them stronger incentive to improve their performance levels, which is beneficial for the auctioneer. However, this objective cannot be achieved when the auctioneer cannot commit to an auction rule. This is because any announcement made before the second period is not credible; thus, bidders make their performance upgrade decisions based on their own expectations.

The concavity condition *per se* does not drive the result in Proposition 6; rather, it ensures that the optimization problem in the limited commitment case is well behaved. In fact, as long as $\partial\pi/\partial w = 0$ (the second equation in [15]) has a unique solution, the result in Proposition 6 holds. (See this in the proof of Proposition 6.) Also, we can show that the concavity assumption on the revenue function can be relaxed to a weaker notion of quasiconcave functions. The concavity condition is satisfied by many common distributions, for instance, uniform distribution.

The following is one numerical example.

EXAMPLE 3. *Let $F(v) = v$ (i.e., uniform distribution at [0,1]), $n = 5$, $\alpha = 0.5$, $c = 0.051$, $y_H = 1$, and y_L change from 0.1 to 0.9. We can derive the revenue-maximizing weighting factors and the maximum revenue with the changes in the y_L in the limited-commitment case and the full-commitment case, respectively.*

Figure 3 shows that the revenue-maximizing weighting factor in the full-commitment case is less than that in the limited-commitment case. Also, Figure 3 shows that a fully committed auctioneer makes more expected revenue than an auctioneer with limited commitment.

It is worth noting that in general a fully committed auctioneer makes more revenue than an auctioneer with limited commitment. This is because, with full commitment, the auctioneer can always commit to the revenue-maximizing weighting factor chosen in the limited-commitment case and achieve at least the same revenue as in the limited-commitment case (but the reverse is not true). However, bidders may question whether the auctioneer can be trusted to fully commit to her preannounced policy in the absence of a credible communication mechanism. This is because, in the second period, it is in the auctioneer's interest to deviate to an *ex post* optimal policy (different from the committed one), given that bidders have already made their performance investment. In practice, some mechanisms that affect firms' long-term profitability, such as reputation, may alleviate this trust issue to some degree. For example, firms take the risk of losing reputation, and thus losing customers, by failing to follow the policies to which they have previously committed. The online supplement provides discussion on the case with a probabilistic commitment.

7. Conclusion

We study a performance-based unit-price auction model where low-performance bidders can improve their performance level at a cost. We find that in equilibrium low-performance bidders of unit valuations higher than a cutoff value convert to a high performance level, and the other low-performance bidders choose not to do so. The efficient allocation involves weighting bidders' bids by their performance levels, which happens to coincide with the policy in the static case. The revenue-maximizing weighting factor, however, is different from the one in the static case. In addition, whether the auctioneer can commit to an auction policy also affects the revenue-maximizing factor. Under many commonly used distributions with an IHR property (such as uniform, normal, and exponential), the auctioneer in the limited-commitment case continues to have incentive to subsidize low-performance bidders to promote competition as in a static case, and the incentive is even stronger. Moreover, the revenue-maximizing weighting factor in the full-commitment case is less than that in the limited-commitment case, which indicates that the auctioneer should give less preferential treatment to the low-performance bidders to encourage them to improve their performance levels.

Our research generates several managerial implications. First, allocation policies critically affect bidders' performance choices, and auctioneers should handicap low-performance bidders for better overall performance. Improving overall performance should be an important concern for markets such as two-sided networks. For example, in keyword auctions hosted by the search engines, both user traffic and advertising revenue are important for the search engine, and the performance and relevance of the advertisement affect the search engine's long-term user base.

To achieve long-run allocation efficiency, which is probably more important for a start-up company trying to establish a good reputation, it is sufficient to follow the efficient policy developed in a static setting. In particular, the auctioneer needs to estimate the performance levels of the bidders and adjust their unit-price bids by the estimated performance in allocating objects.

If performance improvement takes a long time, and the short-run profit is also important for the auctioneer, it is beneficial for the auctioneer to be biased toward disadvantaged bidders to promote overall competition. Taking into account bidders' possible upgrades in their performance, auctioneers with limited commitment should be even more biased than suggested in a static case. It is important to note that the benefit of preferential treatment to low-performance bidders is at the cost of lower efficiency and the market's overall performance. Moreover, auctioneers may be better off if they can fully commit to a preannounced revenue-maximizing policy. The lack of commitment in keyword auctions may be due to the fact that the market for keyword advertising is still nascent, and both search engines and advertisers are learning from practice. In a well-understood market, it is better for an auctioneer to credibly communicate its commitment (e.g., explicitly announce the weighting factor). In doing so, the auctioneer should show less bias for low-performance bidders to encourage them to improve performance.

Future extension of this research includes introducing competition among multiple auctioneers. In the traditional procurement setting, it is natural to assume that auctioneers are monopolists, since the objects requested differ from one another. In keyword auctions, however, the competition among keyword advertising providers (mainly Google, Yahoo!, and Microsoft) is commonly observed and has become an important feature. Open questions remain: How do bidders make the performance choice with an outside option? And how is the revenue-maximizing weighting factor determined in competing auctions?

Acknowledgments
We thank the departmental editor, the senior editor, three anonymous referees, Barrie R. Nault, and Randal B. Watson for their generous inputs to this manuscript. We also thank seminar participants at the University of Calgary for their helpful feedback.

Appendix: Proof of Propositions

Proof of Proposition 1. We prove the existence and uniqueness in two steps.

First, in any equilibrium, if a low-performance bidder of a certain unit valuation chooses to convert to high performance, so does any low-performance bidder of a higher unit valuation. This is because, for any low-performance bidder of unit valuation v who chooses to remain at the low performance level, his best bidding strategy is to bid $b_H(wv)/w$ by a similar argument to the one for Lemma 1. Therefore, noticing $U(y_H, wv, wb) = \frac{wy_H}{y_L} U(y_L, v, b)$ (also see (W3) in the online supplement), his maximum expected payoff could be

$$U(y_L, v, b_H(wv)/w) = \frac{y_L}{wy_H} U(y_H, wv, b_H(wv)) = \frac{y_L}{wy_H} V_H(wv). \quad \text{(A1)}$$

On the other hand, the bidder earns $V_H(v)$ if he chooses to convert. Because its first-order derivative $[y_H \rho_H(v) - y_L \rho_H(wv)]$ is positive, the payoff difference from conversion $[V_H(v) - \frac{y_L}{wy_H} V_H(wv)]$ increases in v, which means that if a low-performance bidder of unit valuation v has incentive to convert, all low-performance bidders of higher unit valuations have incentive to convert. So the only possible equilibrium pattern is that a cutoff value v^* exists such that all low-performance bidders of unit valuations higher than v^* convert to high performance, and of lower unit valuations do not. The low-performance bidder of unit valuation v^* is indifferent about converting or not, which leads to (10).

Second, v^* is unique. Based on the equilibrium pattern identified above, $F_H(v)$ and $F_L(v)$ must be in the form of (5) and (4). Substituting P_H, P_L (defined by [3]), $F_H(v)$, and $F_L(v)$ into (6) and (7), we can further specify $\rho_L(v)$ and $\rho_H(v)$ as

$$\rho_L(v) = [\alpha F(wv) + (1-\alpha)F(v)]^{n-1}, \text{ for } v \in [0, v^*], \quad \text{(A2)}$$

and

$$\rho_H(v) = \begin{cases} \left[\alpha F(v) + (1-\alpha)F\left(\frac{v}{w}\right)\right]^{n-1}, & \text{for } v \in [0, wv^*], \\ [\alpha F(v) + (1-\alpha)F(v^*)]^{n-1}, & \text{for } v \in (wv^*, v^*], \\ [F(v)]^{n-1}, & \text{for } v \in (v^*, 1]. \end{cases} \tag{A3}$$

If $c = 0$, all low-performance bidders convert to high performance for the associated benefit and thus $v^* = 0$. If $0 < c < \Delta V(1)$, then $\Delta V(v^*) = c$ has a unique solution $v^* \in (0, 1)$. This is because $\Delta V(0) = 0$, and $\Delta V(v^*)$ is continuous and monotonically increasing in v^* by noting the first-order derivative

$$y_H\rho_H(v^*) - y_L\rho_L(v^*) + y_H\int_{wv^*}^{v^*}(n-1)[\alpha F(v) + (1-\alpha)F(v^*)]^{n-2}(1-\alpha)f(v^*)dv > 0. \tag{A4}$$

If $\Delta V(1) \le c$, then $v^* = 1$. □

Proof of Proposition 2. First, by (4) and (5), we can specify $f_L(v)$ and $f_H(v)$ as

$$f_L(v) = \frac{f(v)}{F(v^*)}, \text{ if } v \in [0, v^*]. \tag{A5}$$

$$f_H(v) = \begin{cases} \frac{\alpha f(v)}{1-(1-\alpha)F(v^*)}, & \text{if } v \in [0, v^*], \\ \frac{f(v)}{1-(1-\alpha)F(v^*)}, & \text{if } v \in (v^*, 1]. \end{cases} \tag{A6}$$

Substituting (3), (A5), and (A6) into (12), we can reorganize the expected winning performance as

$$(1-\alpha)y_L\int_0^{v^*}\rho_L(v)f(v)dv + \alpha y_H\int_0^{v^*}\rho_H(v)f(v)dv + y_H\int_{v^*}^1\rho_H(v)f(v)dv. \tag{A7}$$

Taking the first-order derivative of the above with respect to w (recall $\rho_H(v)$ is a step function as specified in [A3]), we have

$$\begin{aligned}&(1-\alpha)[y_L\rho_L(v^*) - y_H\rho_H(v^*)]f(v^*)\frac{dv^*(w)}{dw} \\ &+ (1-\alpha)y_L\int_0^{v^*}\frac{d\rho_L(v)}{dw}f(v)dv \\ &+ \alpha y_H\int_0^{wv^*}\frac{d\rho_H(v)}{dw}f(v)dv \\ &+ \alpha y_H\frac{dv^*(w)}{dw}\int_{wv^*}^{v^*}\frac{d\rho_H(v)}{dv^*}f(v)dv.\end{aligned} \tag{A8}$$

Notice that by (A2) and (A3), for $v \in [0, v^*]$,

$$\rho_L(v) = \rho_H(wv), \tag{A9}$$

and

$$\frac{d\rho_L(v)}{dw} = -w\frac{\alpha}{1-\alpha}\frac{f(wv)}{f(v)}\frac{d\rho_H(v)}{dw}|_{wv}. \tag{A10}$$

Using integration by substitution and then applying (A10),

$$\begin{aligned}\int_0^{wv^*}\frac{d\rho_H(v)}{dw}f(v)dv &= \int_0^{v^*}\frac{d\rho_H(v)}{dw}|_{wz}f(wz)wdz \\ &= -\frac{1-\alpha}{\alpha}\int_0^{v^*}\frac{d\rho_L(v)}{dw}f(v)dv.\end{aligned} \tag{A11}$$

Substituting (A11) into (A8) and integrating $\int_{wv^*}^{v^*}\frac{d\rho_H(v)}{dv^*}f(v)dv$, we can rewrite the first-order derivative as

$$\begin{aligned}&(1-\alpha)[y_L\rho_L(v^*) - y_H\rho_H(v^*)]f(v^*)\frac{dv^*(w)}{dw} \\ &+ (1-\alpha)y_L\int_0^{v^*}\frac{d\rho_L(v)}{dw}f(v)dv \\ &- (1-\alpha)y_H\int_0^{v^*}\frac{d\rho_L(v)}{dw}f(v)dv \\ &+ (1-\alpha)y_H(\rho_H(v^*) - \rho_H(wv^*))f(v^*)\frac{dv^*(w)}{dw} \\ &= (1-\alpha)(y_L - y_H)\rho_L(v^*)f(v^*)\frac{dv^*(w)}{dw} \\ &+ (1-\alpha)(y_L - y_H)\int_0^{v^*}\frac{d\rho_L(v)}{dw}f(v)dv,\end{aligned} \tag{A12}$$

where the equality is due to $\rho_H(wv^*) = \rho_L(v^*)$ from (A9). Since $\frac{dv^*(w)}{dw} > 0$ (by Lemma 3) and $\frac{d\rho_L(v)}{dw} > 0$, the first-order derivative is negative, which implies that the expected performance decreases in w. □

Proof of Proposition 3. Substituting (3), (A5), and (A6) into (13), we can reorganize the social welfare as

$$\begin{aligned}&n\left[(1-\alpha)y_L\int_0^{v^*}v\rho_L(v)f(v)dv \right.\\ &\left.+\alpha y_H\int_0^{v^*}v\rho_H(v)f(v)dv + y_H\int_{v^*}^1 v\rho_H(v)f(v)dv\right] \\ &- n(1-\alpha)(1-F(v^*))c.\end{aligned} \tag{A13}$$

To maximize (A13), we take the first-order derivative with respect to w (and remove the constant n):

$$\begin{aligned}&(1-\alpha)y_L\int_0^{v^*} v\frac{d\rho_L(v)}{dw}f(v)dv\\&+\alpha y_H\int_0^{wv^*} v\frac{d\rho_H(v)}{dw}f(v)dv\\&+\alpha y_H\frac{dv^*(w)}{dw}\int_{wv^*}^{v^*} v\frac{d\rho_H(v)}{dv^*}f(v)dv\\&+(1-\alpha)[y_L\rho_L(v^*)-y_H\rho_H(v^*)]v^*f(v^*)\frac{dv^*(w)}{dw}\\&+(1-\alpha)f(v^*)\frac{dv^*(w)}{dw}c.\end{aligned} \quad (A14)$$

We have

$$\begin{aligned}\alpha y_H\int_{wv^*}^{v^*} v\frac{d\rho_H(v)}{dv^*}f(v)dv &= \alpha y_H\int_{wv^*}^{v^*} v(n-1)\\&\quad\times[\alpha F(v)+(1-\alpha)F(v^*)]^{n-2}(1-\alpha)f(v^*)f(v)dv\\&=(1-\alpha)y_Hf(v^*)\int_{wv^*}^{v^*} vd\rho_H(v)\\&=(1-\alpha)y_Hf(v^*)\left(\rho_H(v)v|_{wv^*}^{v^*}-\int_{wv^*}^{v^*}\rho_H(v)dv\right)\\&=(1-\alpha)y_Hf(v^*)\\&\quad\times\left(\rho_H(v^*)v^*-\rho_L(v^*)wv^*-\int_{wv^*}^{v^*}\rho_H(v)dv\right),\end{aligned} \quad (A15)$$

where the third equality is due to integration by parts and last equality is due to $\rho_H(wv^*)=\rho_L(v^*)$ from (A9). Using integration by substitution and then applying (A10), we also have

$$\begin{aligned}\int_0^{wv^*} v\frac{d\rho_H(v)}{dw}f(v)dv &= w^2\int_0^{v^*} v\frac{d\rho_H(v)}{dw}|_{wz}f(wz)dz\\&=-w\frac{1-\alpha}{\alpha}\int_0^{v^*} v\frac{d\rho_L(v)}{dw}f(v)dv.\end{aligned} \quad (A16)$$

Substituting (A15) and (A16) into (A14), the first-order derivative can be reorganized as

$$\begin{aligned}&(1-\alpha)(y_L-wy_H)\int_0^{v^*} v\frac{d\rho_L(v)}{dw}f(v)dv\\&+(1-\alpha)(y_L-wy_H)\rho_L(v^*)v^*f(v^*)\frac{dv^*(w)}{dw}\\&+(1-\alpha)f(v^*)\frac{dv^*(w)}{dw}\left[c-y_H\int_{wv^*}^{v^*}\rho_H(v)dv\right].\end{aligned} \quad (A17)$$

Using $\rho_H(wv)=\rho_L(v)$ by (A9), we can rewrite the equilibrium condition $\Delta V(v^*)=c$ as

$$(y_Hw-y_L)\int_0^{v^*}\rho_L(x)dx+y_H\int_{wv^*}^{v^*}\rho_H(x)dx=c. \quad (A18)$$

By substituting (A18) into (A17), the first-order derivative can be reorganized as

$$\begin{aligned}&(1-\alpha)\left[f(v^*)\frac{dv^*(w)}{dw}\left(v^*\rho_L(v^*)-\int_0^{v^*}\rho_L(v)dv\right)\right.\\&\left.+\int_0^{v^*} v\frac{d\rho_L(v)}{dw}f(v)dv\right](y_L-wy_H).\end{aligned} \quad (A19)$$

Since the coefficient of (y_L-wy_H) is positive, $w=y_L/y_H$ is the only solution making (A19) zero. Therefore, $w^*=y_L/y_H$. □

Proof of Proposition 4. Given v^*, the first-order derivative of π in (14) with respect to w can be organized as (substituting [3], [A5], and [A6] into [14])

$$\begin{aligned}\frac{\partial\pi}{\partial w}&=n(1-\alpha)y_L\int_0^{v^*}\frac{d\rho_L(v)}{dw}J_L(v)f(v)dv\\&\quad+n\alpha y_H\int_0^{wv^*}\frac{d\rho_H(v)}{dw}J_H(v)f(v)dv\\&=n(1-\alpha)y_L\int_0^{v^*}\frac{d\rho_L(v)}{dw}J_L(v)f(v)dv\\&\quad-n(1-\alpha)y_H\int_0^{v^*}\frac{d\rho_L(v)}{dw}J_H(wv)f(v)dv,\end{aligned} \quad (A20)$$

where the second equality is obtained by applying integration by substitution similar to the one in (A11). Substituting $J_L(v)$ and $J_H(v)$ into (A20), we have

$$\begin{aligned}\frac{\partial\pi}{\partial w}&=n(1-\alpha)\int_0^{v^*}\frac{d\rho_L(v)}{dw}\left[(y_L-y_Hw)v.\right.\\&\quad\left.+\left(y_H\frac{1-F_H(wv)}{f_H(wv)}-y_L\frac{1-F_L(v)}{f_L(v)}\right)\right]f(v)dv.\end{aligned} \quad (A21)$$

Notice that

$$\frac{1-F_L(v)}{f_L(v)}=\frac{F(v^*)-F(v)}{f(v)}\le\frac{1-F(v)}{f(v)} \text{ and}$$

$$\frac{1-F(wv)}{f(wv)}\le\frac{\frac{1-(1-\alpha)F(v^*)}{\alpha}-F(wv)}{f(wv)}=\frac{1-F_H(wv)}{f_H(wv)}.$$

If $\frac{f(v)}{1-F(v)}$ is increasing in v, then $\frac{1-F(v)}{f(v)} \leq \frac{1-F(wv)}{f(wv)}$ and thus $\frac{1-F_L(v)}{f_L(v)} \leq \frac{1-F_H(wv)}{f_H(wv)}$. Therefore, for all $w \in [0, \frac{y_L}{y_H}]$, we have $\frac{\partial\pi}{\partial w} > 0$, which implies $w^{\mathrm{L}}_{\mathrm{opt}} > \frac{y_L}{y_H}$. □

Proof of Proposition 5. Substituting $F(v) = v$ and $f(v) = 1$ into (A21),

$$\frac{\partial\pi}{\partial w} = n(1-\alpha)(n-1)(\alpha w + 1 - \alpha)^{n-2}\alpha \int_0^{v^*} v^{n-1} \times \left[2(y_L - wy_H)v + \left(\frac{1-(1-\alpha)v^*}{\alpha} y_H - v^* y_L\right)\right] dv. \tag{A22}$$

By integration, we can obtain the solution to $\partial\pi/\partial w = 0$:

$$w^{\mathrm{L}}_{\mathrm{opt}} = \frac{(n+1)y_H \frac{1-(1-\alpha)v^*}{\alpha} + (n-1)y_L v^*}{2ny_H v^*}.$$

It can be verified that $w^{\mathrm{L}}_{\mathrm{opt}}$ decreases in v^*. Notice that in the static case $v^* = 1$ and in our dynamic (nontrivial) case with limited commitment $v^* < 1$. So the revenue-maximizing weighting factor in a dynamic limited-commitment case is greater than in a static case. □

Proof of Proposition 6. In the full-commitment case, we have

$$\frac{d\pi}{dw} = \frac{\partial\pi}{\partial w} + \frac{\partial\pi}{\partial v^*}\frac{dv^*(w)}{dw}, \tag{A23}$$

where $\partial\pi/\partial w$ is the same as specified in (A20). To get $\partial\pi/\partial v^*$, by substituting P_L, P_H, $f_L(v)$, and $f_H(v)$ (from [3], [A5], and [A6]) into (14), we can reorganize the expected revenue π as

$$\begin{aligned} & n(1-\alpha)y_L \int_0^{v^*} \rho_L(v) J_L(v) f(v) dv \\ & + ny_H\left[\int_0^{v^*} \rho_H(v) J_H(v)\alpha f(v) dv + \int_{v^*}^1 \rho_H(v) J_H(v) f(v) dv\right] \\ & = n(1-\alpha)y_L \int_0^{v^*} \rho_L(v)[vf(v) - (F(v^*) - F(v))] dv \\ & + \ ny_H\left[\int_0^{v^*} \rho_H(v)[\alpha v f(v) - (1-(1-\alpha)F(v^*)) \right. \\ & \left. + \alpha F(v)] dv + \int_{v^*}^1 \rho_H(v)[vf(v) - (1-F(v))] dv\right], \end{aligned}$$

where the equality is obtained by substituting in $J_L(v)$, $J_H(v)$, $F_L(v)$, and $F_H(v)$ (from [4] and [5]). Then, we have

$$\begin{aligned} \frac{\partial\pi}{\partial v^*} & = n(1-\alpha)[y_L\rho_L(v^*) - y_H\rho_H(v^*)]v^* f(v^*) \\ & - n(1-\alpha)f(v^*)y_L \int_0^{v^*} \rho_L(v) dv \\ & + n(1-\alpha)f(v^*)y_H \int_0^{v^*} \rho_H(v) dv \\ & + ny_H \int_{wv^*}^{v^*} \frac{d\rho_H(v)}{dv^*} J_H(v)\alpha f(v) dv \\ & = n\left[(1-\alpha)[(y_L\rho_L(v^*) - y_H\rho_H(v^*))v^* + c]f(v^*). \right. \\ & \left. + \alpha y_H \int_{wv^*}^{v^*} \frac{d\rho_H(v)}{dv^*} J_H(v) f(v) dv\right], \end{aligned} \tag{A24}$$

where the last equality is obtained by substituting in $y_H \int_0^{v^*} \rho_H(v)dv - y_L \int_0^{v^*} \rho_L(v)dv = c$. This explains (18).

Notice that $J_H(v)\alpha f(v) = [\alpha v f(v) - (1-(1-\alpha) F(v^*)) + \alpha F(v)]$ for $v \in [wv^*, v^*]$. By substituting (A15) into (A24),

$$\begin{aligned} \frac{\partial\pi}{\partial v^*} & = n(1-\alpha)f(v^*)\left[(y_L\rho_L(v^*) - wy_H\rho_H(wv^*))v^*. \right. \\ & \left. + y_H \int_0^{wv^*} \rho_H(v) dv - y_L \int_0^{v^*} \rho_L(v) dv\right] \\ & + n\left[-y_H \int_{wv^*}^{v^*} \frac{d\rho_H(v)}{dv^*}(1-(1-\alpha)F(v^*) - \alpha F(v)) dv\right]. \end{aligned}$$

Notice that the term in the second square bracket is negative. The term in the first square bracket, by integration by substitution for $\int_0^{wv^*} \rho_H(v)dv$ and applying $\rho_H(wv) = \rho_L(v)$ from (A9), can be simplified to $[(y_L - wy_H)(\rho_L(v^*)v^* - \int_0^{v^*} \rho_L(v)dv)]$, which is negative for $w > y_L/y_H$. Therefore, $\partial\pi/\partial v^* < 0$ for all $w \geq y_L/y_H$.

Notice that $\frac{dv^*(w)}{dw} > 0$ by Lemma 3. Therefore, for a certain $\tilde{w} \geq y_L/y_H$, we have $\frac{d\pi}{dw}|_{\tilde{w}} < 0$ by (A23) as long as $\frac{\partial\pi}{\partial w}|_{\tilde{w}} \leq 0$. In particular, if $F(v)$ satisfies the IHR condition, $w^{\mathrm{L}}_{\mathrm{opt}} > y_L/y_H$; if $\frac{\partial^2\pi}{\partial w^2} < 0$, then $\frac{\partial\pi}{\partial w}|_{w^{\mathrm{L}}_{\mathrm{opt}}} = 0$ leads to $\frac{\partial\pi}{\partial w} < 0$ for $w \in (w^{\mathrm{L}}_{\mathrm{opt}}, 1]$. So we have $\frac{d\pi}{dw} < 0$ for $w \in (w^{\mathrm{L}}_{\mathrm{opt}}, 1]$, which implies $w^{\mathrm{F}}_{\mathrm{opt}} < w^{\mathrm{L}}_{\mathrm{opt}}$. □

References

Arozamena, L., E. Cantillon. 2004. Investment incentives in procurement auctions. *Rev. Econ. Stud.* **71**(1): 1–18.

Branco, F. 2002. Procurement favouritism and technology adoption. *Eur. Econ. Rev.* **46**(1): 73–91.

Bulow, J., J. Roberts. 1989. The simple economics of optimal auctions. *J. Pol. Econ.* **9**(5): 1060–1090.

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

Ewerhart, C., K. Fieseler. 2003. Procurement auctions and unit-price contracts. *RAND J. Econ.* **34**(3): 569–581.

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

Feng, J., Z. Shen, R. Zhan. 2007b. Ranked items auctions and online advertisement. *Prod. Operat. Manag.* **16**(4): 510–522.

Fudenberg, D., J. Tirole. 1991. *Game Theory.* The MIT Press, Cambridge, MA.

Graham, W. C. 1983. Government procurement policies: GATT, the EEC, and the United States. Trebilcock, M. J., J. R. S. Prichard, T. J. Courchene, J. Whalley eds. *Federalism and the Canadian Economic Union*. University of Toronto Press, Toronto, 355–393.

IAB. 2008. IAB internet advertising revenue report. Tech. rep., An Industry Survey Conducted by PricewaterhouseCoopers and Sponsored by the Interactive Advertising Bureau.

Lebrun, B. 1999. First price auctions in the asymmetric n bidder case. *Int. Econ. Rev.* **40**(1): 125–142.

Liu, D., J. Chen. 2006. Designing online auctions with past performance information. *Dec. Supp. Syst.* **42**(3): 1307–1320.

Liu, D., J. Chen, A. B. Whinston. 2008. Ex-ante information and the design of keyword auctions. *Info. Syst. Res.* [DOI: 10.1287/isre.1080.0225].

McAfee, R. P., J. McMillan. 1986. Bidding for contracts: A principal-agent analysis. *Rand J. Econ.* **17**(3): 326–338.

McAfee, R. P., J. McMillan. 1989. Government procurement and international trade. *J. Int. Econ.* **26**(3–4): 291–308.

Myerson, R. B. 1981. Optimal auction design. *Math. Operat. Res.* **6**(1): 58–73.

Rothkopf, M. H., A. B. Whinston. 2007. On e-auctions for procurement operations. *Prod. Operat. Manag.* **16**(4): 404–408.

Samuelson, W. F. 1983. Competitive bidding in defense contracting. Engelbrecht-Wiggans, R., M. Shubik, R. M. Stark eds. *Auctions, Bidding, and Contracting: Uses and Theory.* New York University Press, New York, 389–419.

Samuelson, W. F. 1986. Bidding for contracts. *Manag. Sci.* **32**(12): 1533–1550.

Schlosser, A. E., T. B. White, S. M. Lloyd. 2006. Converting web site visitors into buyers: How web site investment increases consumer trusting beliefs and online purchase intentions. *J. Market.* **70**(3): 133–148.

Stark, R. M. 1974. Unbalanced highway contract tendering. *Operat. Res. Quart.* **25**(3): 373–388.

Stegeman, M. 1996. Participation costs and efficient auctions. *J. Econ. Theory* **71**(1): 228–259.

Tan, G. 1992. Entry and R&D costs in procurement contracting. *J. Econ. Theory* **58**(1): 41–60.

Weber, T. A., Z. Zheng. 2007. A model of search intermediaries and paid referrals. *Info. Syst. Res.* **18**(4): 414–436.

Supporting Information

Additional supporting information may be found in the online version of this article:

Appendix S1: Proof of Lemmas and Corollary.
Appendix S2: Derivation of Expected Revenue.
Appendix S3: The Case with a Probabilistic Commitment.
Appendix S4: Extensions.

Please note: Wiley-Blackwell are not responsible for the content or functionality of any supporting materials supplied by the authors. Any queries (other than missing material) should be directed to the corresponding author for the article.