

Optimal Commitments in Asymmetric Auctions with Incomplete Information

PINGZHONG TANG, Tsinghua University
ZIHE WANG, Tsinghua University
XIAOQUAN (MICHAEL) ZHANG, HKUST Business School

We solve the Bayesian sequential equilibrium of a general class of single-item first-price or all-pay auctions of incomplete information. Our main contribution is a general methodology for solving the optimal commitment problem, in closed form, for asymmetric continuous-type distributions.

Our approach consists of a number of innovations. We propose a modeling concept called equal-bid function to build a bridge between two players' strategies. Another concept called equal-utility curve transforms any commitment strategy into a weakly better continuous and everywhere directional-differentiable strategy.

The optimal commitment functions in these auctions reveal some important insights. When the player with commitment power (the leader) has low valuation, he bids passively. This is a credible way to alleviate competition and to enable collusion. We demonstrate, via concrete examples, that this is a credible way to threaten the follower so that the leader can secure a higher utility.

1. INTRODUCTION

The first-price auction has seen many applications because it is simple, intuitive and easy to implement. In first-price auctions, the highest bidder wins and pays her own bid. Theoretically, it enjoys many desirable modeling properties. For example, in symmetric settings, the auction has a unique efficient Bayes Nash Equilibrium (BNE) [Chawla and Hartline 2013]. In contrast, second-price auctions may have many inefficient equilibria.

At the same time, first-price auctions pose important challenges for both academics and practitioners. For one, in complete-information settings, the auction format, including the class of *generalized first-price auctions*, sometimes does not have a pure Nash equilibrium and is practically observed to be unstable [Börger et al. 2013; Edelman and Ostrovsky 2007; Edelman et al. 2007]. In incomplete-information settings where bidders have asymmetric type distributions, it is extremely difficult to solve or characterize its BNE. In fact, this has been one of the most elusive open problems in the literature of auction analysis [Fibich and Gavish 2011; Hartline et al. 2014; Lebrun 1999; Vickrey 1961]. To date, the problem has closed-form solutions only in very restrictive settings such as two-bidder asymmetric uniform distributions [Fibich and Gavish 2012; Kaplan and Zamir 2012].

In this paper, we examine a general class of first-price auction games in which a single leader is capable of making commitments to his strategy before a follower subsequently chooses her strategy to maximize her payoff. Different from complete-

Authors' addresses: Pingzhong Tang, IIIS, Tsinghua University; email: kenshinping@gmail.com; Zihe Wang, IIIS, Tsinghua University; email: wzh5858588@163.com; Xiaoquan (Michael) Zhang, HKUST Business School; email: zhang@ust.hk.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

EC'16, July 24–28, 2016, Maastricht, The Netherlands. Copyright © 2016 ACM 978-1-4503-3936-0/16/07 ...\$15.00.

<http://dx.doi.org/10.1145/2940716.2940739>

information symmetric first-price auction games, the players have asymmetric type distributions and the players' types may be random draws from period to period. Such a setting, known as Bayesian Stackelberg games [Fudenberg and Tirole 1991], is both theoretically interesting and practically meaningful. [Fudenberg and Levine 1989] propose a reputation game in which a long-run player can enjoy a premium by making up-front commitments. In flower auctions, some large players may have daily variations in their values, but they can make credible commitments by following a pre-announced strategy and making their past bids publicly verifiable. Such commitments can bring them extra payoffs compared with the case when daily auctions are treated independently. In a recent application of the theory [Paruchuri et al. 2007; Tambe 2011],¹ the Los Angeles International Airport adopted an algorithm called ARMOR (Assistant for Randomized Monitoring of Routes) to randomize the patrols [Pita et al. 2008]. At the core of the algorithm is a game of a leader (the law enforcement) committing to a publicly known strategy that achieves the highest payoff against the follower (the terrorist).

In a Stackelberg equilibrium, given that it is possible for a follower to know a leader's committed strategy in advance and best-responds to it, the leader solves for an optimal strategy to commit to. A Stackelberg equilibrium formulation is particularly useful when one player has credibility to commit. It is well known that commitment weakly increases the leader's utility compared to his payoff in a simultaneous-move Nash equilibrium. Furthermore, there are efficient algorithms to compute it in basic settings [Conitzer and Sandholm 2006; Letchford and Conitzer 2010]. Since the players' types are not always pre-determined (e.g., demand of flower on a particular day, potential terrorist attacks, etc.), we study a Bayesian Stackelberg game to accommodate such uncertainty.

1.1. Commitment

Consider the following example in a first-price auction.² There are two bidders. One of the bidders, we call him leader A , has the power to commit to a bidding strategy first. Then the other bidder, we call her follower B , best-responds to the leader's commitment. We next show that commitment can increase A 's payoff.

Example 1.1. Both players' types, x and y , are drawn uniformly from $[0, 1]$. Let the leader's strategy be $s_A(x) = x^2/2$. Clearly, the follower B must never bid more than 0.5 in this case. In fact, her utility when bidding $t \leq 0.5$ is $(y - t)\sqrt{2t}$. B 's best strategy is $s_B(y) = y/3$. The expected utility of A is:

$$\int_0^{\sqrt{2/3}} (x - \frac{1}{2}x^2) \cdot \frac{3}{2}x^2 dx + \int_{\sqrt{2/3}}^1 (x - \frac{1}{2}x^2) dx = 0.2029$$

The first term considers the case when x is in $[0, \sqrt{2/3}]$, while the second term considers the case when x is in $(\sqrt{2/3}, 1]$. On the other hand, in the unique symmetric BNE where each bidder bids half of the value, A 's expected revenue is $1/6 \approx 0.167$. In comparison, commitment to s_A increases A 's utility by 21%.

The concept of commitment has been observed in the domain of auction design, even though sometimes implicitly. Note that early bidding and sniping in online auctions (e.g., eBay auctions) can be regarded as two forms of commitment [Gray and Reiley

¹See also, <http://www.newsweek.com/random-security-laxs-armor-system-103885>, for a *Newsweek* report.

²We will use first-price auction as a running example throughout. We extend our approach and results to general rank-and-bid based auctions [Chawla and Hartline 2013] in the full version.

2013; Roth and Ockenfels 2002]. An advertiser that has a “passive” image (i.e., rarely changes its bid, or always submits low bids) in sponsored search auctions can be seen as another form of commitment. [Abraham et al. 2013] consider a super bidder who has access to more information than others and study how this will affect seller’s revenue in an alternative solution concept called tremble-robust equilibrium. Their setting is similar to commitment but not the same. [Skreta 2006] considers another type of commitment where the auctioneer is lack of credibility to reserve the item and studies how this lack of commitment affects revenue.

A closely-related parallel work is [Xu and Ligett 2015]. They characterize optimal commitment for first price auctions with complete information. For the case with incomplete information, they assume that the bidders’ types are drawn from discrete distributions and prove a partial property that the commitment function can be divided into pieces. In comparison, we consider general continuous distributions and obtain closed-form characterizations for a more general class of auctions. Our technical approach also offers new insights in solving such models.

In this paper, we study the optimal commitment in first-price auctions with incomplete information (Bayesian setting). There is one item, and two bidders A and B . Two bidders’ valuation are independent and differentiable distributions $F_1(x)$ and $F_2(y)$.

1.2. Our contributions

Our main contribution is a general approach to solve and characterize optimal commitment in this class of auctions, for any continuous type distributions. In particular, applying our approach, we are able to solve optimal commitments for first-price and all-pay auctions in *closed-form* for fairly general distributions. Our approach and results on these auctions in a sense mitigate the difficulties of deriving a game-theoretical prediction in first price and all-pay auctions with asymmetric type distributions. Our approach consists of several nontrivial techniques. We dedicate Section 3 to introduce the technical contribution (our main contribution). Here, we focus on the economic interpretations of our results.

1.2.1. Equal-bid function and characterization. Equal bid function g is a mapping function from leader A ’s value to follower B ’s value. $g(x)$ is the minimum of follower’s value which has best response equal to the leader A ’s bid $s_A(x)$. Roughly speaking, when B ’s value is $g(x)$, B bids $s_A(x)$, same as when A has value x . $g(x)$ is the lowest such type. It is a convenient modeling trick that allows us to greatly simplify the calculations and representations. Based on equal bid function g , we derive our main theorem.

1.2.2. Structure and closed-form. Theorem 6.2 says when F_2 is uniformly distributed on $[b_1, b_2]$,

- if $\forall t, 2f_1^2(t) - F_1[t]f_1'(t) \geq 0$, then optimal $g(x)$ has at most 3 values.
- if $b_1 = 0$, then $g(x)$ has at most 2 pieces.

Theorem 6.3 gives the closed form of first price auction when two bidders’ type are drawn uniformly from $[0, 1]$.

$$s_A^*(x) = \begin{cases} 0 & x \in [0, t_0] \\ 1 - \frac{t_0}{x} & x \in (t_0, 1] \end{cases}, t_0 \approx 0.567$$

The optimal utility of the leader is 0.22, increasing by 32% than the utility in Bayesian Nash equilibrium.

Our results on first price auction reveal certain striking findings: the leader bids very passively when his type is low. Even worse, he bids 0 when his type is below a threshold (depending on type distributions). This is against the common sense in first




Fig. 1. The optimal commitment strategy keeps zero when type is low. It increases rapidly when type is large enough.

price and all pay auctions that bidding 0 has no chance of winning at all. However, a close scrutiny states otherwise: by committing to a passive image, the leader (credibly) ensures the follower that he has no intention to compete when he has a low type, thus effectively brings down the follower’s bid, since the follower does not know the leader’s actual type and views the leader as a mixed strategy. As a result, the leader eventually wins the auction with less competition when he has a high type. We also note that such passive bidding behaviors had been observed in major search engines such as Yahoo³ and Baidu (before they switch to GSP).

Furthermore, the commitment solutions are largely consistent with the collusive-seeming behaviors studied in first-price auction [Aryal and Gabrielli 2013; Lopomo et al. 2011; Marshall and Marx 2007; McAfee and McMillan 1992; Pesendorfer 2000]: players coordinate to bring down the prices. Our results further suggest that such collusive-seeming behaviors are stable: the trust between the players is built on the rationality of the follower, as well as the leader’s credibility to commit.

2. THE SETTING

We consider a single item auction with two bidders, one called the *leader A* (male) and the other called the *follower B* (female). Bidder *A* has a private valuation x drawn from distribution F_1 with support $[a_1, a_2]$, while bidder *B*’s private valuation y is drawn from distribution F_2 with support on $[b_1, b_2]$. We use f_1 and f_2 to denote the probability density of F_1 and F_2 , respectively. We also sometimes write $A = x$ to denote the case where *A*’s type is x . Similarly, we can write $B = y$.

Leader *A* commits to a Bayesian strategy $s_A : [a_1, a_2] \rightarrow \Delta R$, where ΔR denotes the set of bid distributions on R . He announces this strategy and the follower *B* best responds to the leader’s committed strategy via a single bid.⁴ We follow the standard definition by [Conitzer and Sandholm 2006] of Bayesian commitment that the leader only announces his strategy, i.e., the function $s_A(\cdot)$, without revealing his actual type. Compared to the utility from BNE, being able to commit increases the leader’s utility.

The timing of the game is summarized as follows:

- (1) Leader *A* announces her Bayesian strategy $s_A(\cdot)$ to follower *B*;
- (2) Leader *A* then draws a random type x from a type distribution F_1 , which is commonly known;

³<http://webscope.sandbox.yahoo.com/catalog.php?datatype=a>

⁴It is easy to see that it is never profitable for *B* to use a mixed strategy.

- (3) Leader A and B participate in a simultaneously auction, where A commits to bid $s_A(x)$, while B knows $s_A(\cdot)$ but not x and bids optimally to her available information.⁵
- (4) The payoffs of both players are then determined according to the auction rule.

There are two tie-breaking rules.

ASSUMPTION 2.1. *When B has multiple best responses, she will choose the one that maximizes A 's winning probability.*

ASSUMPTION 2.2. *When there is a tie, the good will be assigned to B .*

We make these assumptions for expositional simplicity. Our main results do not depend on these assumptions.⁶ Given s_A , B 's best response is fixed by assumption, and both players' winning probabilities are fixed.

Our goal is to solve for the optimal s_A in first price auctions. Finding the optimal strategy s_A is known to be difficult. [Conitzer and Sandholm 2006] show that computing optimal commitments in general Bayesian games is NP-hard.

Definition 2.3. $P_A[x, s_A]$ denotes A 's winning probability when he is at type x and adopts strategy s_A , while $P_B[t, s_A]$ is B 's winning probability when bidding t against A 's strategy s_A . In circumstances where there is no ambiguity, we use $P_A[t], P_B[t]$ instead.

Definition 2.4. We use u_A and u_B to denote A and B 's expected utility, respectively. Given s_A , we can derive P_B . Since B best responses to A , we have

$$u_B(y) = \sup_{t \geq 0} (y \cdot P_B[t] - t \cdot P_B[t])$$

where t represents B 's bid. After figuring out B 's response, we can compute A 's expected utility,

$$u_A(s_A) = E_{x \sim F_1} (x \cdot P_A[x] - s_A(x) \cdot P_A[x])$$

Although B 's valuation is on $[b_1, b_2]$, we extend its definition domain to $[0, b_2]$, and let $F_2[x] = 0$, for $x < b_1$.

Note that the follower cares about the local maximal utility on each type, and the input for u_B is a type. The leader cares about the maximal expected utility on the whole range and the input for u_A is a strategy function.

We use \sup rather than \max here because we need to show that \max is attainable, which is done in Lemma 4.1. We also note that bidding zero yields a nonnegative utility for B , so $u_B \geq 0$ and $u_B(0) = 0$.

Definition 2.5. Let $S_B(y, s_A)$ denote $B = y$'s best responses against A 's strategy s_A . In circumstances where there is no ambiguity, we use $S_B(y)$ instead.

We will prove this definition is well-defined in Lemma 4.1, i.e. $S_B(y)$ is nonempty.

3. AN OVERVIEW OF OUR APPROACH

The main contribution of this paper is to put forward a general approach for solving optimal commitments in first price auction. The approach can be sketched as follows. It is useful to understand the intuition behind this approach before the details.

⁵We use first-price auctions as a running example, these results remain to be valid for all-pay auctions.

⁶We will show formally in the full version that neither of the assumptions is necessary.

3.1. Difficulties of the problem

If the optimal commitment was monotone and differentiable, we can use the first-order condition to get the result: follower best responses to the leader. That is, (and if $g(x)$ is well-defined)

$$\begin{aligned}(g(x) - s_A(t))F_1[t] &\text{ is maximized at } t = x \\ (g(x) - s_A(t))f_1(t) - s'_A(t)F_1[t] &= 0 \\ g(x)f_1(t) &= [s_A(t)F_1[t]]' \\ s_A(x) &= \frac{1}{F_1(x)} \int_{a_1}^x g(t)f_1(t)dt\end{aligned}$$

However, Maskin and Riley [2003]’s approach to proving that player’s strategy is differentiable in a Nash equilibrium does not apply here. Their result is based on the fact that everyone best responds to others. In our setting, the leader’s optimal strategy is not necessarily continuous and differentiable. The main difficulty of the problem is that the follower’s best response cannot be represented as a closed-form function of the leader’s Bayesian strategy, which is also a function. This difficulty further prevents us from obtaining a closed-form representation of the leader’s winning probability and utility. As a result, standard functional optimization techniques cannot be applied.

To appreciate the difficulty, it is helpful to look at the problem of finding a Bayes Nash Equilibrium with two asymmetric bidders in a first-price auction — one of the most elusive open problems in the analysis of auctions [Hartline et al. 2014; Kaplan and Zamir 2012]. The main barrier in that literature is exactly the difficulty to represent one’s best response as a concise function of the other’s strategy.

Previous work on this problem focuses on different cases of finding optimal commitments in first price auctions with complete information [e.g., ?]. They do however consider the case of incomplete information but with discrete types. They obtain a partial characterization (the optimal commitment is a piecewise function). In fact, we consider our using equal utility curves to prove the continuity and differentiability to be one of our most innovative technical contributions. In addition, obtaining the closed form solution also allows us to derive richer economic insights.

3.2. Step one: sorting the leader’s strategy and making it monotone

As discussed above, one of the obstacles is that the optimal leader strategy may not be monotone or differentiable. Our first effort is to sort the leader’s strategy. We prove that, for any leader’s strategy (optimal or not), one can sort it into a monotone function that preserves the follower’s best response without hurting the leader’s utility. In other words, the leader always prefers the strategy after sorting.

3.3. Step two: smoothing the leader’s strategy and making it continuous

A more difficult task is to smooth the leader’s strategy, that is, to turn it into a continuous and everywhere directional-differentiable function. To achieve this goal, we introduce a methodological innovation called *equal-utility curve*. Roughly, given the follower’s type, the equal-utility curve represents a leader’s strategy such that the follower gets the same utility no matter what she bids. We can show that such a curve always exists. Furthermore, the supremum (over all follower’s types) of all such curves defines a new leader’s strategy that enjoys the following important properties: it is continuous, left and right differentiable, preserves the follower’s best response, and weakly improves the leader’s utility. To this point, one can truly focus on monotone, continuously everywhere directional-differentiable leader strategies. Actually after this step, the modified strategy has a nice structure.

3.4. Step three: representing the strategy by a everywhere directional-differentiable equal-bid function

A key insight here is to represent everything (the strategies, utilities and winning probabilities of both players) as a function of some g , coined the *equal-bid function*, that maps a leader's type to a follower's type. Intuitively, $g(t)$ is the follower's type at which she submits *the same* bid as the leader does when the leader has type t . In other words, $g(t)$, later to be proved as monotone, can be seen as a bifurcation type between winning and losing for the follower. As one can imagine, together with the cumulative distribution function of the follower's type, we can represent the leader's winning probability (hence utility) as a function of g . A similar but different idea has appeared in [Hafalir and Krishna 2008; Lebrun 1999] in which *inverse bid functions* are used to represent the best responses of players. However, *inverse bid functions* are in pairs and the two functions make the final optimization problem complicated. In contrast, our proposed equal-bid function represents everything with a single function.

3.5. Step four: optimizing A's expected utility in terms of equal-bid function

With the above transformations, it turns out that we can find a bijection between the set of monotone, continuously differentiable leader's strategies and the set of continuous and monotone equal-bid functions. As a result, we can focus on optimizing over equal-bid functions and the Lagrangian method applies.

We conclude with a characterization of the optimal commitment for general follower type distributions and compute the closed-form optimal commitments for both first-price and all-pay auctions when the follower has uniform type distributions.

4. SORTING AND SMOOTHING THE LEADER'S STRATEGY

We first transform an arbitrary leader strategy into a continuous weakly increasing strategy and show that this transformation does not hurt the leader's expected utility.

In the next section, we prove some additional properties, such as everywhere directional-differentiability, for the transformed strategy.

We first show that the notion of "best response" is well defined for the follower B .

LEMMA 4.1. *For any B 's valuation y , B has a best response, i.e. $u_B(y)$ can be attained by some bid and the smallest bid exists among the best responses.*

By Assumption 2.2, B always chooses the smallest bid among all best responses.

4.1. Sorting s_A

For an arbitrary strategy s_A , the support of $s_A(v)$ on value v may not be a single bid. Function s_A could also be nonmonotone. The following lemma says, to find the optimal commitment, it suffices to consider the strategies with the desirable properties below.

LEMMA 4.2. *We can sort any strategy s_A for A into a new strategy function \check{s}_A such that (1) $\check{s}_A(v)$ is a deterministic bid for any v and is weakly increasing in v (2) the best response of \check{s}_A remains the same as s_A (3) \check{s}_A yields at least the same utility for the leader as s_A .*

For example, suppose leader A has equal probability on three types 1, 2 and 3. Then \check{s}_A brings higher expected utility than s_A .

$$s_A(x) = \begin{cases} 0.8 & x = 1 \\ 0.6 & x = 2 \\ 1 & x = 3 \end{cases} \quad \check{s}_A(x) = \begin{cases} 0.6 & x = 1 \\ 0.8 & x = 2 \\ 1 & x = 3 \end{cases}$$

With this result, for ease of presentation, we can use s_A to denote \check{s}_A henceforth. So s_A is a weakly increasing, nonnegative strategy function.

The following example shows how to calculate u_B in a first-price auction.

Example 4.3. Both bidders' value distribution are uniform on $[0,1]$.

$$s_A(x) = \begin{cases} x/4 & x \leq 0.4 \\ x - 0.3 & x > 0.4 \end{cases}$$

By definition $u_B(y) = \max\{\max_{t \leq 0.1}(y-t) \cdot 4t, \max_{t > 0.1}(y-t)(t+0.3)\}$, we have

$$u_B(y) = \begin{cases} y^2 & y \in [0, 0.2) \\ 0.4y - 0.04 & y \in [0.2, 0.5] \\ (y+0.3)^2/4 & y \in (0.5, 1] \end{cases}$$

The following Lemma suggests the set of best responses of B at type y_1 generally does not intersect with the set of best responses at type y_2 and these sets are well sorted by the value of y , except for the following special case.

LEMMA 4.4. *For B 's valuations $y_1 < y_2$, if $\exists a \in S_B(y_1), b \in S_B(y_2)$ and $a > b$, then we have $S_B(y_1) \subseteq S_B(y_2)$, $u_B(y_1) = u_B(y_2) = 0$ and $P_B[b] = P_B[a] = 0$.*

We can now prove that the follower's utility is continuous and monotone.

LEMMA 4.5. *$u_B(y)$ is continuous and weakly increasing.*

4.2. Smoothing s_A

To smooth s_A into a continuous and everywhere directional-differentiable function, we now introduce an important innovation of our approach: the equal-utility curve.

Definition 4.6. Define equal-utility curve $eu_B(\cdot, \cdot) : [0, b_2] \times (a_1, a_2) \rightarrow \mathbb{R}$,

$$eu_B(y, x) = y - \frac{u_B(y)}{F_1[x]} \quad (1)$$

The interpretation of $eu_B(y, \cdot)$ is that, any value of $eu_B(y, \cdot)$ (as a function of x) is a best response of the follower at type y .

Consider Example 4.3, in a first-price auction, $F_1[x] = F_2[x] = x$, $\forall x \in [0, 1]$, the definition of eu_B above is simplified as

$$eu_B(y, x) = y - \frac{u_B(y)}{x}.$$

When B 's value $y = 0.5$, her best response is to bid 0.1 with utility 0.16. So $u_B(0.5) = 0.16$, and the equal-utility curve is $eu_B(0.5, x) = 0.5 - \frac{0.16}{x}$, shown in Fig 2. If A uses strategy

$$\max\{eu_B(0.5, x), 0\} = \begin{cases} 0 & x \in [0, 0.32) \\ 0.5 - \frac{0.16}{x} & x \in [0.32, 1] \end{cases}$$

then the utility of B when $y = 0.5$ is the same for any bid in $[0, 0.34]$.

Function eu_B represents the leader's strategy against which the follower will achieve the same largest utility no matter what the follower bids. It's easy to check that $eu_B(0, \cdot) = 0$.

LEMMA 4.7. (1) $eu_B(y, x)$ is weakly increasing and differentiable in x . $eu_B(y, x)$ is continuous in y .

(2) $eu_B(y, x) \leq s_A(x)$

We are now ready to introduce the smoothing method, by constructing an envelope s_A^* using eu_B .




Fig. 2. Equal-utility Curve. When $y = 0.5$, the utility of bidding 0 or 0.3 are always 0.16.

Definition 4.8. $s_A^*(x) = \sup_{y \in [0, b_2]} eu_B(y, x), \forall x \in (a_1, a_2]$

We will prove that strategy s_A^* yields at least the same revenue for A as strategy s_A . The outline is as follows. First, we prove that, although the leader's bid distribution changes, we keep the utility of the follower to be the same. Second, we prove the best response of B is still the best response after smoothing, for any follower's value (Lemma 4.10). Third, we prove the leader's winning probability does not change. Finally, since the leader's bid is weakly decreasing, we can prove that the leader's utility weakly increases after smoothing (Theorem 4.13).

A mathematical view of the motivation of $s_A^*(x)$ can be found in Remark 1. The idea is to suppress the bids of the leader while maintain his winning probability.

Consider Example 4.3, we can find $s_A^*(x)$ easily:

$$s_A^*(x) = \begin{cases} x/4 & x \in [0, 0.4) \\ x - 0.3 & x \in [0.4, 0.65] \\ 1 - \frac{0.65^2}{x} & x \in (0.65, 1] \end{cases}$$

We now prove some basic properties of s_A^* that will be used later.

LEMMA 4.9. (1) For any $x \in (a_1, a_2]$, $(x, s_A^*(x))$ must lie on some Equal-Utility Curve.

(2) When $s_A^*(x) > \lim_{t \rightarrow a_1} s_A^*(t)$, $s_A^*(x)$ strictly increases.

(3) When $s_A^*(x) = \lim_{t \rightarrow a_1} s_A^*(t)$, $(x, s_A^*(x))$ lies on $eu_B(s_A^*(x), \cdot)$, and $u_B(s_A^*(x)) = 0$.

(4) For any x , we have $s_A^*(x) \leq s_A(x)$.

(5) $s_A^*(x)$ is continuous.⁷

Up to now, the domain of s_A^* is defined as $(a_1, a_2]$. Since s_A^* is continuous, we can define $s_A^*(a_1) = \lim_{x \rightarrow a_1} s_A^*(x)$. For simplicity, we use $P_B^*[t]$ and $S_B^*(y)$ instead of $P_B[t, s_A^*]$ and $S_B(y, s_A^*)$. Next, we study how the follower's utility and best response would change in s_A and s_A^* .

LEMMA 4.10. (1) When A 's strategy s_A is changed to s_A^* , $u_B(y)$ remains the same.

(2) If $t \in S_B(y)$ then $t \in S_B^*(y)$, $S_B(y) \subseteq S_B^*(y)$. If $P_B^*[t] \neq P_B[t]$ then $u_B(y) = 0$ and $t = \lim_{x \rightarrow a_1} s_A^*(x)$

The lemma below draws connections between equal-utility curve and s_A^* .

⁷The limit of a continuous function may not be continuous, so this argument is not trivial. Consider $y_k(x) = kx, x \in [0, 1/k]$, constantly zero for $x \leq 0$ and constant one for $x \geq 1/k$. Clearly, y_k is continuous but $\sup_k y_k$ is not.

LEMMA 4.11. (1) If $(x_0, s_A^*(x_0))$ lies on equal-utility line $eu_B(y_0, \cdot)$, then $s_A^*(x_0) \in S_B^*(y_0)$. (2) If $s_A^*(x_0) \in S_B^*(y_0)$ and $s_A^*(x_0) \neq \lim_{x \rightarrow a_1} s_A^*(x)$ then $(x_0, s_A^*(x_0))$ lies on equal-utility line $eu_B(y_0, \cdot)$.

The best response set S_B^* of the follower in s_A^* is a superset of S_B . Since the best response of the follower is sorted, $S_B^*(y)$ is bounded by any element in $S_B^*(y - \epsilon)$ and $S_B^*(y + \epsilon)$. Therefore, it is bounded by $S_B(y - \epsilon)$ and $S_B(y + \epsilon)$. We then prove that for most of the types, the winning probability of the leader will not change.

Definition 4.12. If $\exists x$ such that $s_A^*(x) = \lim_{t \rightarrow a_1} s_A^*(t)$, let $\hat{x} = \sup\{x | \lim_{t \rightarrow a_1} s_A^*(t) = s_A^*(x)\}$.

Since s_A^* is continuous, $s_A^*(\hat{x}) = \lim_{t \rightarrow a_1} s_A^*(t)$. Combined with the fact that the bids decrease in s_A^* (thus lower payment) and the winning probability remains the same, we prove that the expected utility does not decrease.

THEOREM 4.13. By using s_A^* instead of s_A , the expected utility of A does not decrease.

The intuition is that both strategies lead to the same follower's best response, while s_A^* is lower and the leader has less payment. It may be possible that s_A^* is worse at countable number of breaking points, however, the effect of these points to the expected utility is 0.

5. BIJECTIVE MAPPING BETWEEN S_A^* AND G

The final step is to show that every s_A^* can be represented by a function g . Once we obtain such a function, we will be able to focus on optimizing such a g instead.




Fig. 3. M_1 is bijective between O_1 and O_2

Definition 5.1. $\hat{y} = \sup\{y | u_B(y) = 0\}$

Definition 5.2. $\forall x > a_1, Y(x) = \{y | eu_B(y, \cdot) \text{ passes through point } (x, s_A^*(x))\}$.

LEMMA 5.3. (1) $Y(x)$ is closed. (2) $Y(x) \geq \hat{y}$. (3) For all $x_1 < x_2, Y(x_1) \leq Y(x_2)$. (4) $(\hat{y}, b_2] \subseteq \cup_x Y(x)$ (5) $Y(x)$ is an interval or is a unique number. $Y(x)$ contains only one element for almost all x .

Definition 5.4. Equal-bid function $g(x) = \min Y(x), \forall x \in (a_1, a_2]$

The intuitive explanation of g is: when the follower's type $y = g(x)$, one of her best response is equal to the bid $s_A^*(x)$ by the leader. In most of the time except for countable types, if the follower gives the same bid as the leader, then the lowest type of the follower must be $g(x)$. In other words, the follower with value $g(x)$ submits the same bid as the leader who has value x .

With equal-bid function, the winning probability of the leader has a surprisingly concise form.

Consider Example 4.3, we have

$$g(x) = \begin{cases} x/2 & x \in [0, 0.4) \\ 2x - 0.3 & x \in [0.4, 0.65] \\ 1 & x \in (0.65, 1] \end{cases}$$

It's easy to check that $s_A(0.45) = 0.15$, $g(0.45) = 0.6$, $S_B(0.6) = \{0.15\}$. The leader with value 0.45 bids 0.15, same as the follower with value $g(0.45) = 0.6$.

By lemma 5.3, $\{y | eu_B(y, x) = s_A^*(x)\}$ is closed, so $g(x)$ is well defined. When $y < g(x)$, the leader $A = x$ beats the follower y by Lemma 4.4. When $y \geq g(x)$, the follower y beats the leader $A = x$ by the tie-breaking rule. So we can calculate the winning probability of the leader $A = x$ using $g(x)$.

Note $s_A^*(x)$ is not yet defined on a_1 . In fact, bidding with zero probability does not affect overall utility. We can define $s_A^*(a_1) = \lim_{t \rightarrow a_1} s_A^*(t)$ for convenience.

Now we prove the last a few desirable properties of s_A^* : s_A^* is differentiable on both sides. Based on the derivatives, we find the relationship between g and s_A^* .

THEOREM 5.5.

- (1) $s_A^*(x)$ is left-hand differentiable and right-hand differentiable.
- (2) $s_A^*(x) = \frac{1}{F_1[x]} \int_{a_1}^x f_1(t)g(t)dt$.

Up to now, we have developed a new strategy s_A^* for A based on s_A , with at least 2 desirable properties: it yields at least as much utility as s_A and is left-hand differentiable and right-hand differentiable. In the following, we will calculate the winning probability and find out the s_A^* with the optimal utility.

Remark 1. From s_A (we only need the weakly increasing condition), we can define g directly, but s_A cannot be calculated by g . To see this, the follower bids $s_A(t)$ and when $t = x$, the follower achieves the highest utility. Taking first-price auctions for example, we have $(g(x) - s_A(x))F_1[x] \geq (g(x) - s_A(t))F_1[t] \forall t$, then $s_A(t) \geq g(x) - \frac{(g(x) - s_A(x))F_1[x]}{F_1[t]}$, equality can be achieved by setting t to be x . Moreover $s_A(t) \geq \sup_x (g(x) - \frac{u_B(x)}{F_1[t]}) = \sup_x eu_B(g(x), t)$, equality may not be achieved, because if we fix t first, there might be no corresponding x . Thus we do not have the exact formula of $s_A(t)$. If s_A is optimal, for any leader's type, his bid should be as small as possible without changing the follower's behavior. To do this, when $s_A(t) > \sup_x eu_B(g(x), t)$, we can decrease his bid to $\epsilon + \sup_x eu_B(g(x), t)$, without letting the follower match his new bid. This is the nature of the smoothing method. So in the optimal strategy, we should have $s_A(t) = \sup_x eu_B(g(x), t)$, which is exactly the new strategy out of the smoothing method, s_A^* .

We can check the correctness of relationship between g and s_A^* in Example 4.2.

When $x_0 \in [0, 0.4]$,

$$\frac{1}{F_1[x_0]} \int_{a_1}^{x_0} g(x)f_1(x)dx = \frac{1}{x_0} \left[\int_0^{x_0} x/2dx \right] = x_0/4 = s_A^*(x_0)$$

When $x_0 \in [0.4, 0.65]$,

$$\frac{1}{F_1[x_0]} \int_{a_1}^{x_0} g(x)f_1(x)dx = \frac{1}{x_0} \left[\int_0^{0.4} x/2dx + \int_{0.4}^{x_0} (2x - 0.3)dx \right] = x_0 - 0.3 = s_A^*(x_0)$$

When $x_0 \in [0.65, 1]$,

$$\begin{aligned} \frac{1}{F_1[x_0]} \int_{a_1}^{x_0} g(x) f_1(x) dx &= \frac{1}{x_0} \left[\int_0^{0.4} x/2 dx + \int_{0.4}^{0.65} (2x - 0.3) dx + \int_{0.65}^{x_0} 1 dx \right] \\ &= \frac{1}{x_0} [0.04 + 0.65 * 0.35 - 0.04 + x_0 - 0.65] = 1 - \frac{0.65^2}{x} = s_A^*(x_0) \end{aligned}$$

Definition 5.6.

$O_1 = \{s_A \mid \text{strategies resulted from any nonnegative strategy after smoothing}\}$

$O_2 = \{(g, s_A(a_1)) \mid g \text{ is weakly increasing and left continuous and in } [0, b_2]\}$

Since the domain of y in $eu_B(y, x)$ is $[0, b_2]$, we have $Y(x) \subset [0, b_2]$ and $g(x) \in [0, b_2]$. Thus Definition 5.4 gives a mapping $M_1 : O_1 \rightarrow O_2$. In fact, we will prove that there is a bijective mapping between the two sets. The idea is that we construct a mapping $M_2 : O_2 \rightarrow O_1$ and prove $M_1 \circ M_2 = I$.

THEOREM 5.7. *There is a bijective mapping between O_1 and O_2 .*

Thus finding the optimal strategy is equivalent to finding the optimal function g such that $(g, 0) \in O_2$.

6. OPTIMIZING EQUAL-BID FUNCTION G

In this section, we solve for the optimal g in order to derive the final form of s_A^* .

6.1. General optimization

Since $s_A^*(x) = \frac{1}{F_1[x]} \int_{a_1}^x f_1(t)g(t)dt$, the expected utility becomes

$$u_A(s_A^*) = \int_{a_1}^{a_2} \left[x - \frac{1}{F_1[x]} \int_{a_1}^x f_1(t)g(t)dt \right] F_2[g(x)] f_1(x) dx$$

This is a function of g , denoted by $M(g)$. For any admissible function $g + \epsilon j$, i.e. $(g + \epsilon j, 0) \in O_2$, we have $M(g) \geq M(g + \epsilon j)$. Consider the marginal loss in direction j ,

$$\begin{aligned} 0 &\geq \lim_{\epsilon \rightarrow 0} \frac{M(g + \epsilon j) - M(g)}{\epsilon} \\ &= \int_{a_1}^{a_2} \left[x - \frac{1}{F_1[x]} \int_{a_1}^x f_1(t)g(t)dt \right] f_2(g(x)) j(x) f_1(x) dx \\ &\quad + \int_{a_1}^{a_2} -\frac{1}{F_1[x]} \int_{a_1}^x f_1(t)g(t)dt F_2[g(x)] f_1(x) dx \\ &= \int_{a_1}^{a_2} j(x) f_1(x) [f_2(g(x)) \left(x - \frac{1}{F_1[x]} \int_{a_1}^x f_1(t)g(t)dt \right) - \int_x^{a_2} \frac{F_2[g(t)] f_1(t)}{F_1[t]} dt] \quad (2) \end{aligned}$$

Let $h(x)$ denote the coefficient of $j(x)$:

$$h(x) = f_1(x) \left(x f_2(g(x)) - \frac{f_2(g(x))}{F_1[x]} \int_{a_1}^x f_1(t)g(t)dt - \int_x^{a_2} \frac{F_2[g(t)] f_1(t)}{F_1[t]} dt \right)$$

We can deduce the optimal g when $h(x)$ has some good property.

THEOREM 6.1. (1) *For an interval L , if $h(x) > 0$, $x \in L$, we have*

$$g(x) = \lim_{t \rightarrow (\sup L)^+} g(t) \quad x \in L$$

Moreover, if $\sup L = a_2$, then $g(x) = b_2$, $x \in L$.

Similarly, if $h(x) < 0, x \in L$, we have

$$g(x) = \lim_{t \rightarrow \inf L} g(t), x \in L$$

If $\inf L = a_1$, then $g(x) = 0, x \in L$.

(2) There is an optimal g such that $g(x) \in 0 \cup [b_1, b_2]$.

In fact, g can be derived explicitly in fairly general settings, as we show below.

6.2. Some specific optimization results

THEOREM 6.2. *When F_2 is uniformly distributed on $[b_1, b_2]$,*

(1) *if $\forall t, 2f_1^2(t) - F_1[t]f_1'(t) \geq 0$, then optimal $g(x)$ consists of at most 3 values.*

(2) *if $b_1 = 0$, then optimal $g(x)$ consists of 2 pieces. When $t_0 = a_2 \int_{t_0}^{a_2} \frac{f_1(t)}{F_1[t]} dt$ has a solution,*

$$g(x) = \begin{cases} 0 & x \in (a_1, t_0] \\ b_2 & x \in (t_0, a_2) \end{cases}, \text{ where } t_0 = a_2 \int_{t_0}^{a_2} \frac{f_1(t)}{F_1[t]} dt$$

Otherwise, $g(x) = b_2, \forall x \in [a_1, a_2]$.

THEOREM 6.3. *In Example 4.3, the optimal utility of the leader is 0.22. The closed-form representation of the optimal $g(x)$ and $s_A^*(x)$ are:*

$$g(x) = \begin{cases} 0 & x \in [0, t_0] \\ 1 & x \in (t_0, 1] \end{cases} \quad s_A^*(x) = \begin{cases} 0 & x \in [0, t_0] \\ 1 - \frac{t_0}{x} & x \in (t_0, 1] \end{cases}, t_0 \approx 0.567$$

Here t_0 is the solution of $t_0 = b_2 \int_{t_0}^{a_2} \frac{f_1(x)}{F_1[x]} dx$.

We should notice that the leader bids zero 56.7% of the time.

7. CONCLUSION

In many institutional settings, out of many participants, one dominant player may exert a significant influence on the potential equilibrium that will be played by all. Commitment is one such way. When a dominant player can establish a reputation, commitment to a publicly known strategy can result in higher profits for the player.

This paper offers a framework to derive the closed-form solution to a class of games including first-price and all-pay auctions when one player can use commitment as a viable strategy. Different from works in the literature, our model allows for incomplete information with asymmetric value distributions.

To solve this problem, we propose some innovations in the modeling approach. Specifically, we propose a modeling concept called equal-bid function to build a bridge between two players' strategies. Another concept called equal-utility curve transforms any commitment strategy into a weakly better continuous and everywhere directional-differentiable strategy.

With closed-form solutions, our findings offer some interesting insights. With commitment, a relatively low-valued leader will be less aggressive so that the two sides can collude and reduce competition. Overall, we show that commitment allows the leader to obtain higher payoffs compared to the case when commitment is not allowed. Our methodological approach can be adopted in similar settings to find optimal commitment strategies.

8. ACKNOWLEDGEMENT

We are grateful to the colleagues in Baidu for discussions on the passive bidding behaviors in sponsored search auctions. We thank Vince Conitzer and Jason Hartline for helpful discussions on Bayes Commitment.

This work was supported by the National Basic Research Program of China Grant 2011CBA00300, 2011CBA00301, the Natural Science Foundation of China Grant 61033001, 61361136003, 61303077, 61561146398, a Tsinghua Initiative Scientific Research Grant and a China Youth 1000-talent program. Michael Zhang is grateful to generous support from Hong Kong Research Grant Council grants 694213, 11504815, and 16504614.

REFERENCES

- Ittai Abraham, Susan Athey, Moshe Babaioff, and Michael Grubb. 2013. Peaches, lemons, and cookies: designing auction markets with dispersed information. In *Proceedings of the fourteenth ACM conference on Electronic commerce*. ACM, 7–8.
- Gaurab Aryal and Maria F Gabrielli. 2013. Testing for collusion in asymmetric first-price auctions. *International Journal of Industrial Organization* 31, 1 (2013), 26–35.
- Tilman B"orgers, Ingemar Cox, Martin Pesendorfer, and Vaclav Petricek. 2013. Equilibrium bids in sponsored search auctions: Theory and evidence. *American economic Journal: microeconomics* 5, 4 (2013), 163–187.
- Shuchi Chawla and Jason D Hartline. 2013. Auctions with unique equilibria. In *Proceedings of the fourteenth ACM conference on Electronic commerce*. ACM, 181–196.
- Vincent Conitzer and Tuomas Sandholm. 2006. Computing the Optimal Strategy to Commit to. In *Proceedings of the 7th ACM conference on Electronic commerce*. ACM, 621–641.
- Benjamin Edelman and Michael Ostrovsky. 2007. Strategic bidder behavior in sponsored search auctions. *Decision support systems* 43, 1 (2007), 192–198.
- Benjamin Edelman, Michael Ostrovsky, and Michael Sshwartz. 2007. Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords. *The American economic review* 97, 1 (2007), 242–259.
- Gadi Fibich and Nir Gavish. 2011. Numerical simulations of asymmetric first-price auctions. *Games and Economic Behavior* 73, 2 (2011), 479–495.
- Gadi Fibich and Nir Gavish. 2012. Asymmetric First-Price Auctions-A Dynamical-Systems Approach. *Mathematics of Operations Research* 37, 2 (2012), 219–243.
- Drew Fudenberg and David K. Levine. 1989. Reputation and Equilibrium Selection in Games with a Patient Player. *Econometrica* 57, 4 (1989), 759–778.
- Drew Fudenberg and Jean Tirole. 1991. *Game Theory*. MIT Press.
- Sean Gray and David H. Reiley. 2013. Measuring the Benefits to Sniping on eBay: Evidence from a Field Experiment. *Journal of Economics and Management* 9, 2 (July 2013), 137–152.
- Isa Hafalir and Vijay Krishna. 2008. Asymmetric Auctions with Resale. *American Economic Review* 98, 1 (2008), 87–112.
- Jason Hartline, Darrell Hoy, and Sam Taggart. 2014. Price of anarchy for auction revenue. In *Proceedings of the fifteenth ACM conference on Economics and computation*. ACM, 693–710.
- Todd R Kaplan and Shmuel Zamir. 2012. Asymmetric first-price auctions with uniform distributions: analytic solutions to the general case. *Economic Theory* 50, 2 (2012), 269–302.
- Bernard Lebrun. 1999. First price auctions in the asymmetric N bidder case. *International Economic Review* 40, 1 (1999), 125–142.
- Joshua Letchford and Vincent Conitzer. 2010. Computing Optimal Strategies to Commit to in Extensive-Form Game. In *Proceedings of the 11th ACM conference on Electronic commerce*. ACM, 83–92.
- Giuseppe Lopomo, Leslie M Marx, and Peng Sun. 2011. Bidder collusion at first-price auctions. *Review of Economic Design* 15, 3 (2011), 177–211.

- Robert C Marshall and Leslie M Marx. 2007. Bidder collusion. *Journal of Economic Theory* 133, 1 (2007), 374–402.
- Eric Maskin and John Riley. 2003. Uniqueness of equilibrium in sealed high-bid auctions. *Games and Economic Behavior* 45, 2 (2003), 395–409.
- R Preston McAfee and John McMillan. 1992. Bidding rings. *The American Economic Review* (1992), 579–599.
- Praveen Paruchuri, Jonathan P. Pearce, Milind Tambe, Fernando Ordonez, and Sarit Kraus. 2007. An efficient heuristic approach for security against multiple adversaries. In *AAMAS 07: Proceedings of the 6th international joint conference on Autonomous agents and multiagent systems*. International Foundation for Autonomous Agents and Multiagent Systems.
- Martin Pesendorfer. 2000. A study of collusion in first-price auctions. *The Review of Economic Studies* 67, 3 (2000), 381–411.
- James Pita, Manish Jain, Janusz Marecki, Fernando Ordóñez, Christopher Portway, Milind Tambe, Craig Western, Praveen Paruchuri, and Sarit Kraus. 2008. Deployed ARMOR protection: the application of a game theoretic model for security at the Los Angeles International Airport. In *Proceedings of the 7th international joint conference on Autonomous agents and multiagent systems: industrial track*. International Foundation for Autonomous Agents and Multiagent Systems, 125–132.
- Alvin E. Roth and Axel Ockenfels. 2002. Last-Minute Bidding and the Rules for Ending Second-Price Auctions: Evidence from eBay and Amazon Auctions on the Internet. *American Economic Review* 92, 4 (September 2002), 1093–1103.
- Vasiliki Skreta. 2006. Sequentially Optimal Mechanisms. *Review of Economic Studies* (2006), 1085–1111.
- Milind Tambe. 2011. *Security and Game Theory: Algorithms, Deployed Systems, Lessons Learned* (1st ed.). Cambridge University Press, New York, NY, USA.
- W. Vickrey. 1961. Counterspeculation, Auctions and Competitive Sealed Tenders. *Journal of Finance* (1961), 8–37.
- Yunjian Xu and Katrina Ligett. 2015. Commitment in First-Price Auctions. In *Algorithmic Game Theory*. Springer, 299.