On Cheating in Sealed-Bid Auctions

Ryan Porter *,1, Yoav Shoham **,1

Computer Science Department
Stanford University
Stanford, CA 94305

Abstract

Motivated by the rise of online auctions and their relative lack of security, this paper analyzes two forms of cheating in sealed-bid auctions. The first type of cheating we consider occurs when the seller examines the bids of a second-price auction before the auction clears and then submits a shill bid in order to increase the payment of the winning bidder. In the second type, a bidder cheats in a first-price auction by examining the competing bids before submitting his own bid. In both cases, we derive equilibrium strategies when bidders are aware of the possibility of cheating. These results provide insights into sealed-bid auctions even in the absence of cheating, including some counterintuitive results on the effects of overbidding in a first-price auction.

Key words: Game Theory, Sealed-Bid Auctions, Cheating
PACS: 89.65.Gh
1991 MSC: 91B26, 91A06, 91A10, 91A35

* Corresponding author. Email: rwporter@stanford.edu.
**Email: shoham@stanford.edu.
[1] This work was supported in part by DARPA grant F30602-00-2-0598. We also thank Eugene Nudelman for his help on several technical points of the proofs. A earlier version of this paper appeared as [17].

1 Introduction

Among the types of auctions commonly used in practice, sealed-bid auctions are a good practical choice because they require little communication or interaction between the participants. After each bidder submits a bid, the winner is immediately determined. However, sealed-bid auctions require that the bids be kept private until the auction clears. The increasing popularity of online auctions only makes this disadvantage more troublesome. At an auction house, with all participants present, it is difficult to examine a bid given directly to the auctioneer. However, in an online auction the auctioneer is often little more than a server with questionable security; and, since the seller and each of the bidders are in different locations, one can anonymously attempt to break into the server. In this paper, we present a game theoretic analysis of how bidders should behave when they are aware of the possibility of cheating that is based on knowledge of the bids prior to the closing of the auction.

We investigate this type of cheating along two dimensions: whether it is the seller or a bidder who cheats (we assume that the auctioneer is honest), and which variant (either first or second-price) of a sealed-bid auction is used. Note that two of these cases are trivial. In our setting (formally defined below), there is no incentive for the seller to submit a shill bid in a first price auction, because doing so would either cancel the auction or not affect the payment of the winning bidder. In a second-price auction, knowing the competing bids does not help a bidder because it is dominant strategy to bid his valuation. This leaves us with two cases that we examine in detail.

A seller can profitably cheat in a second-price auction by examining the bids before the auction clears and submitting an extra bid. This possibility was pointed out as early as the seminal paper [22] that introduced this type of auction. For example, if the bidders in an eBay auction each use a proxy bidder (essentially creating a second-price auction), then the seller may be able to break into eBay's server, observe the maximum price that a bidder is willing to pay, and then extract this price by submitting a shill bid just below it using a false identity. We assume that there is no chance that the seller will be caught when it cheats. However, not all sellers are willing to use this power (or, not all sellers can successfully cheat). We assume that each bidder knows the probability with which the seller will cheat. In this setting, we derive an equilibrium bidding strategy for the case in which each bidder's value for the good is independently drawn from an identical distribution (with no further assumptions except for continuity and differentiability). This result shows how first and second-price auctions can be viewed as the endpoints of a spectrum of auctions.

But why should the seller have all the fun? In a first-price auction, a bidder must bid below his value for the good (also called “shaving” his bid) in order to have positive utility if he wins. To decide how much to shave his bid, he must trade off the probability of winning the auction against how much he will pay if he does win. Of course, if he could simply examine the other bids before submitting his own, then his problem is solved: bid the minimum (up to his valuation) necessary to win the auction. In this setting, our goal is to derive an equilibrium bidding strategy for a non-cheating bidder who is aware of the possibility that he is competing against cheating bidders. When bidder values are drawn from the commonly-analyzed uniform distribution, we show the counterintuitive result that the possibility of other bidders cheating has no effect on the equilibrium strategy of an honest bidder. This result is then extended to show the robustness of the equilibrium of a first-price auction without the possibility of cheating. Finally, we explore other distributions, including some in which the presence of cheating bidders actually induces an honest bidder to lower its bid, and we provide sufficient conditions for this effect to occur.

The rest of the paper is structured as follows. In Section 2 we formalize the setting and present our results for the case of a seller possibly cheating in a second price auction. Section 3 covers the case of bidders possibly cheating in a first-price auction. In Section 4, we quantify the effects that the possibility of cheating has on the expected revenue of an honest seller in the two settings. We discuss related work, including other forms of cheating in auctions, in Section 5, before concluding with Section 6. All proofs and derivations appear in the appendix.

2 Second-Price Auction, Cheating Seller

In this section, we consider a second-price auction in which the seller may cheat by inserting a shill bid after observing all of the bids. The formulation for this section will be largely reused in the following section on bidders cheating in a first-price auction. While no prior knowledge of game theory or auction theory is assumed, good introductions can be found in [7] and [10], respectively.

2.1 Formulation

The setting consists of N bidders, or agents, (indexed by $i = 1, \ldots, n$) and a seller. Each agent has a type $\theta_i \in [0, 1]$, drawn from a continuous range, which represents the agent’s independent private value for the good being auctioned.[2] Each agent's type is independently drawn from a cumulative distribution function (*cdf*) F over $[0, 1]$, where $F(0) = 0$ and $F(1) = 1$. We assume that $F(\cdot)$ is strictly increasing and differentiable over the interval $[0, 1]$. Call the probability density function (*pdf*) $f(\theta_i) = F'(\theta_i)$, which is the derivative of the *cdf*.

Each agent knows its own type θ_i, but knows only the distribution $F(\cdot)$ over the possible types of the other agents. A bidding strategy for an agent $b_i : [0, 1] \to [0, 1]$ maps its type to its bid.[3] Let $\theta = (\theta_1, \cdots, \theta_n)$ be the vector of types for all agents, and $\theta_{-i} = (\theta_1, \cdots, \theta_{i-1}, \theta_{i+1}, \cdots \theta_n)$ be the vector of all types except for that of agent i. We can then combine the vectors so that $\theta = (\theta_i, \theta_{-i})$. We also define the vector of bids as $b(\theta) = (b_1(\theta_1), \ldots, b_n(\theta_n))$, and this vector without the bid of agent i as $b_{-i}(\theta_{-i})$. Let $b_{[1]}(\theta)$ be the value of the highest bid of the vector $b(\theta)$, with a corresponding definition for $b_{[1]}(\theta_{-i})$.

An agent obviously wins the auction if its bid is greater than all other bids; however, ties complicate the formulation. In our analysis, we can ignore the case of ties because our continuity assumption makes them a zero probability event in equilibrium. Also, we assume that the seller does not set a reserve price.[4]

If the seller does not cheat, then the winning agent pays the highest bid by another agent. On the other hand, if the seller does cheat, then the winning agent pays its bid, since we assume that a cheating seller would take full advantage of its power. Let the indicator variable μ^c be 1 if the seller cheats, and 0 otherwise. The probability that the seller cheats, P^c, is known by all agents. We can then write the payment of the winning agent as follows.

$$p_i(b(\theta), \mu^c) = \mu^c \cdot b_i(\theta_i) + (1 - \mu^c) \cdot b_{[1]}(\theta_{-i}) \tag{1}$$

Let $\mu(\cdot)$ be an indicator function that takes an inequality as an argument and returns 1 if it is true, and 0 otherwise. The utility for agent i is zero if it does not win the auction, and the difference between its valuation and its payment if it does.

$$u_i(b(\theta), \mu^c, \theta_i) = \mu\Big(b_i(\theta_i) > b_{[1]}(\theta_{-i})\Big) \cdot \Big(\theta_i - p_i(b(\theta), \mu^c)\Big) \tag{2}$$

[2] We can restrict the types to the range $[0, 1]$ without loss of generality because any distribution over a different range can be normalized to this range.
[3] We thus limit agents to deterministic bidding strategies. However, because of our continuity assumption, there always exists a pure strategy equilibrium.
[4] This simplifies the analysis, but all of our results can be applied to the case in which the seller announces a reserve price before the auction begins.

We assume that each agent is risk-neutral, and thus aims to maximize its expected utility, with the expectation taken over the types of the other agents and over whether or not the seller cheats. By pushing the expectation inward (and conditioning the payment term on the agent winning the auction), we can write the expected utility as:

$$E_{\theta_{-i},\mu^c}[u_i(b(\theta),\mu^c,\theta_i)] = Prob\Big(b_i(\theta_i) > b_{[1]}(\theta_{-i})\Big)\cdot \\ \Big(\theta_i - E_{\theta_{-i},\mu^c}\Big[p_i(b(\theta),\mu^c) \mid \Big(b_i(\theta_i) > b_{[1]}(\theta_{-i})\Big)\Big]\Big) \quad (3)$$

Because of the uncertainty over the types of the other agents, we will be looking for a Bayes-Nash equilibrium. A vector of bidding strategies b^* is a Bayes-Nash equilibrium if for each agent i and each possible type θ_i, agent i cannot increase its expected utility by using an alternate bidding strategy b'_i, holding the bidding strategies for all other agents fixed. Formally, b^* is a Bayes-Nash equilibrium if the following expression holds.

$$\forall i,\theta_i,b'_i \;\; E_{\theta_{-i},\mu^c}\Big[u_i\Big(\Big(b^*_i(\theta_i),b^*_{-i}(\theta_{-i})\Big),\mu^c,\theta_i\Big)\Big] \geq \\ E_{\theta_{-i},\mu^c}\Big[u_i\Big(\Big(b'_i(\theta_i),b^*_{-i}(\theta_{-i})\Big),\mu^c,\theta_i\Big)\Big] \quad (4)$$

2.2 Equilibrium

We first present the Bayes-Nash equilibrium for an arbitrary distribution $F(\cdot)$.

Theorem 1 *In a second-price auction in which the seller cheats with probability P^c, it is a Bayes-Nash equilibrium for each agent to bid according to the following strategy:*

$$b^*_i(\theta_i) = \theta_i - \frac{\int_0^{\theta_i} F^{(\frac{N-1}{P^c})}(x)dx}{F^{(\frac{N-1}{P^c})}(\theta_i)} \quad (5)$$

It is useful to consider the extreme points of P^c. Setting $P^c = 1$ yields the correct result for a first-price auction (see, e.g., [10]). In the case of $P^c = 0$, this solution is not defined. However, in the limit, $b^*_i(\theta_i)$ approaches θ_i as P^c approaches 0, which is what we expect as the auction approaches a standard second-price auction.

The position of P^c is perhaps surprising. For example, the linear combination $b_i(\theta_i) = \theta_i - P^c \cdot \frac{\int_0^{\theta_i} F^{(N-1)}(x)dx}{F^{(N-1)}(\theta_i)}$ of the equilibrium bidding strategies of first and second-price auctions would have also given us the correct bidding strategies for the cases of $P^c = 0$ and $P^c = 1$. However, it is not a symmetric equilibrium bidding strategy for intermediate values of P^c.

An alternative perspective on the setting is as a continuum between first and second-price auctions. Consider a probabilistic sealed-bid auction in which the seller is honest, but the payment of the winning agent is determined by a weighted coin flip: with probability P^c the payment is his bid, and with probability $1 - P^c$ it is the second-highest bid. By adjusting P^c, we can smoothly move between a first and second-price auction.

2.3 Special Case: Uniform Distribution

Another way to try to gain insight into Equation 5 is by instantiating the distribution of types. We now consider the often-studied uniform distribution: $F(\theta_i) = \theta_i$.

Corollary 2 *In a second-price auction in which the seller cheats with probability P^c, and $F(\theta_i) = \theta_i$, it is a Bayes-Nash equilibrium for each agent to bid according to the following strategy:*

$$b_i^*(\theta_i) = \frac{N-1}{N-1+P^c}\theta_i$$

This equilibrium bidding strategy, parameterized by P^c, can be viewed as an interpolation between two well-known results. When $P^c = 0$ the bidding strategy is now well-defined (each agent bids its true type), while when $P^c = 1$ we get the correct result for a first-price auction: each agent bids according to the strategy $b_i^*(\theta_i) = \frac{N-1}{N}\theta_i$.

3 First-Price Auction, Cheating Agents

We now consider the case in which the seller is honest, but there is a chance that agents will cheat and examine the other bids before submitting their own (or, alternatively, they will revise their bid before the auction clears). Since this type of cheating serves no purpose in a second-price auction, we only analyze the case of a first-price auction. After revising the formulation from the previous section, we present a fixed point equation for the equilibrium strategy for an arbitrary distribution $F(\cdot)$. This equation will be useful for the analysis of the uniform distribution, in which we show that the possibility of cheating agents does not change the equilibrium strategy of honest agents. This result then has implications for the robustness of the symmetric equilibrium to overbidding in a standard first-price auction. Furthermore, we find that for other distributions overbidding actually induces a competing agent to shave more off of its bid.

3.1 Formulation

It is clear that if a single agent is cheating, he will bid (up to his valuation) the minimum amount necessary to win the auction. It is less obvious, though, what will happen if multiple agents cheat. One could imagine a scenario similar to an English auction, in which all cheating agents keep revising their bids until all but one cheater wants the good at the current winning bid. However, we are only concerned with how an honest agent should bid given that it is aware of the possibility of cheating. Thus, it suffices for an honest agent to know that it will win the auction if and only if its bid exceeds every other honest agent's bid and every cheating agent's *type*.

This intuition can be formalized as the following discriminatory auction. In the first stage, each agent's payment rule is determined. With probability P^a, the agent will pay the second highest bid if it wins the auction (essentially, he is a cheater), and otherwise it will have to pay its bid. These selections are made independently and are recorded by a vector of indicator variables $\mu^a = (\mu^{a_1}, \ldots, \mu^{a_n})$, where $\mu^{a_i} = 1$ denotes that agent i pays the second highest bid. Each agent knows the probability P^a, but does not know the payment rule for all other agents. After agents privately observe their respective payment rules, the auction proceeds as a standard, sealed-bid auction. It is thus a dominant strategy for a cheater to bid its true type, making this formulation strategically equivalent to the setting outlined in the previous paragraph. The expression for the utility of an honest agent in this discriminatory auction is as follows, where the first term is the profit conditional on winning, and each term in the succeeding product is the probability that agent i has a higher bid than agent i.

$$u_i(b(\theta), \mu^a, \theta_i) = \Big(\theta_i - b_i(\theta)\Big) \cdot \prod_{j \neq i} \Big[\mu^{a_j} \cdot \mu\Big(b_i(\theta_i) > \theta_j\Big) + (1 - \mu^{a_j}) \cdot \mu\Big(b_i(\theta_i) > b_j(\theta_j)\Big)\Big] \quad (6)$$

3.2 Equilibrium

Given that in equilibrium the cheating agents will bid according to their dominant strategy, our goal is to find the symmetric bidding strategy for the honest agents that completes the equilibrium. For the general case of an arbitrary $F(\cdot)$, we were not able to derive a closed form solution for the honest agent's bidding strategy, and instead give a fixed point equation for it.

Theorem 3 *In a first-price auction in which each agent cheats with probability P^a, it is a Bayes-Nash equilibrium for each non-cheating agent i to bid according to the strategy that is a fixed point of the following equation:*

$$b_i^*(\theta_i) = \theta_i - \frac{\int_0^{\theta_i} \Big(P^a \cdot F(b_i^*(x)) + (1-P^a) \cdot F(x)\Big)^{(N-1)} dx}{\Big(P^a \cdot F(b_i^*(\theta_i)) + (1-P^a) \cdot F(\theta_i)\Big)^{(N-1)}} \qquad (7)$$

3.3 Special Case: Uniform Distribution

In order to solve Equation 7, we return to the uniform distribution: $F(\theta_i) = \theta_i$. Recall the logic behind the symmetric equilibrium strategy in a first-price auction without cheating: $b_i^*(\theta_i) = \frac{N-1}{N}\theta_i$ is the optimal tradeoff between increasing the probability of winning and decreasing the price paid upon winning, given that the other agents are bidding according to this strategy. Since in the current setting the cheating agents do not shave their bids at all and thus decrease an honest agent's probability of winning (while obviously not affecting the price that an honest agent pays if he wins), it is natural to expect that an honest (risk-neutral) agent should compensate by increasing his bid. The idea is that sacrificing some potential profit in order to regain some of the lost probability of winning would bring the two sides of the tradeoff back into balance. However, it turns out that the equilibrium bidding strategy is unchanged.

Corollary 4 *In a first-price auction in which each agent cheats with probability P^a, and $F(\theta_i) = \theta_i$, it is a Bayes-Nash equilibrium for each non-cheating agent to bid according to the strategy $b_i^*(\theta_i) = \frac{N-1}{N}\theta_i$.*

This result suggests that the equilibrium of a first-price auction is particularly stable when types are drawn from the uniform distribution, since the best response is unaffected by deviations of the other agents to the strategy of always bidding their type. In fact, as long as each other agent shaves its bid by a fraction (which can differ across the agents) no greater than $\frac{1}{N}$, it is still a best response for the remaining agent to bid according to the equilibrium strategy.

Theorem 5 *In a first-price auction where $F(\theta_i) = \theta_i$, if each agent $j \neq i$ bids according a strategy $b_j(\theta_j) = \frac{N-1+\alpha_j}{N}\theta_j$, where $\alpha_j \geq 0$, then it is a best response for the remaining agent i to bid according to the strategy $b_i^R(\theta_i) = \frac{N-1}{N}\theta_i$.*

Note that this theorem applies to the case in which other agents are shaving their bid by a negative fraction, and thus are irrationally bidding above their type. Also, the condition for the theorem is tight, in the sense that if $\alpha_j < 0$ for a single agent j, then agent i's best response is to shave its bid by a fraction greater than $\frac{1}{N}$.

While the strategy profiles of this theorem are not equilibria (except for the case in which $\alpha_j = 0$), the result supports the predictive power of the symmetric equilibrium. A common (and valid) criticism of equilibrium concepts such as Nash and Bayes-Nash is that they are silent on how the agents converge on a strategy profile from which no one wants to deviate. However, Theorem 5 tells us that the symmetric equilibrium strategy is a best response to a large set of strategy profiles that are out of equilibrium, which makes it more plausible that each individual agent will choose to play according to the equilibrium strategy.

3.4 Effects of Overbidding for Other Distributions

We now consider the effect of overbidding for distributions other than uniform. It is certainly the case that distributions can be found that support the intuition given above that an agent should shave its bid less when other agents are bidding their type. Examples include $F(\theta_i) = -\frac{1}{2}\theta_i^2 + \frac{3}{2}\theta_i$ (the solution to the system of equations: $F''(\theta_i) = -1$, $F(0) = 0$, and $F(1) = 1$), and $F(\theta_i) = \frac{e - e^{(1-\theta_i)}}{e-1}$.

However, the robustness result of Theorem 5 is not unique to the uniform distribution– it holds for all distributions of the form $F(\theta_i) = (\theta_i)^k$, for $k > 0$. More surprisingly, taking a simple linear combination of two such distributions to produce $F(\theta_i) = \frac{\theta_i^2 + \theta_i}{2}$ yields a distribution in which the best response for the remaining agent is to shave its bid more than it does in equilibrium. The same result also holds for the normalized exponential distribution $F(\theta_i) = \frac{e^{\theta_i} - 1}{e-1}$.

While we were not able to derive a general condition for the predicting the direction of change, we were able to find sufficient conditions, based on the function $g(\theta_i) = \big(F(\theta_i)/f(\theta_i)\big)$. If $g(\cdot)$ is strictly increasing, and if its second derivative is monotonic, then the sign of the second derivative is the direction of change.

Theorem 6 *In a first-price auction, if* $g'(\theta_i) > 0$ *and* $g''(\theta_i) < 0$ *(resp.* $g''(\theta_i) = 0$, $g''(\theta_i) > 0$*) hold for all* θ_i*, where* $g(\theta_i) = \frac{F(\theta_i)}{f(\theta_i)}$*, and if each agent* $j \neq i$ *bids according the strategy* $b_j(\theta_j) = \theta_j$*, then, for all* $\theta_i > 0$*, the best response bidding strategy* $b_i^R(\theta_i)$ *for agent i is strictly less than (resp. equal to, strictly greater than) the symmetric equilibrium bidding strategy* $b_i^*(\theta_i)$*.*

Note that this theorem applies to all of the distributions mentioned above.

4 Revenue Loss for an Honest Seller

In both of the settings we covered, an honest seller suffers a loss in expected revenue due to the possibility of cheating. The equilibrium bidding strategies that we derived allow us to quantify this loss. Although this is as far as we will take the analysis, we note that it could be applied to more general settings, in which the seller could, for example, choose the market in which he sells his good or pay a trusted third party to oversee the auction.

In a second-price auction in which the seller may cheat, an honest seller suffers due to the fact that the agents will shave their bids out of fear that he will cheat. For the case in which agent types are drawn from the uniform distribution, every agent will shave its bid by $\frac{P^c}{N-1+P^c}$, which is thus also the fraction by which an honest seller's revenue decreases due to the possibility of cheating.

Analysis of the case of a first-price auction in which agents may cheat is not as straightforward. If $P^a = 1$ (each agent cheats with certainty), then the setting is equivalent to a standard second-price auction, and the seller's expected revenue will be unchanged. Again considering the uniform distribution for agent types, it is not surprising that $P^a = \frac{1}{2}$ causes the seller to lose the most revenue. However, even in this worst case, the percentage of expected revenue lost is significantly less than it is for the second-price auction in which $P^c = \frac{1}{2}$, as shown in Table 1.[5] It turns out that setting $P^c = 0.2$ would make the expected loss of these two settings approximately equal. While this comparison between the settings is unlikely to be useful for a seller, it is interesting to note that, at least for the case of the uniform distribution, agent suspicions of possible cheating by the seller are in some sense worse than agents actually cheating themselves.

[5] Note that we have not considered the costs of the seller. Thus, the expected loss in profit could be much greater than the numbers that appear here.

	Percentage of Revenue lost for an Honest Seller	
Agents	Second-Price Auction ($P^c = 0.5$)	First-Price Auction ($P^a = 0.5$)
2	33	12
5	11	4.0
10	5.3	1.8
15	4.0	1.5
25	2.2	0.83
50	1.1	0.38
100	0.50	0.17

Table 1
The percentage of expected revenue lost by an honest seller due to the possibility of cheating in the two settings considered in this paper. Agent valuations are drawn from the uniform distribution.

5 Discussion and Related Work

The work most closely related to our first setting is [20], which also presents a model in which the seller may cheat in a second-price auction. In their setting, the seller is a participant in the Bayesian game, and he decides between running a first-price auction (where profitable cheating is never possible) or a second-price auction. The seller makes this choice after observing his type, which is his probability of having the opportunity and willingness to cheat in a second-price auction. The bidders, who know the distribution from which the seller's type is drawn, then place their bid. It is shown that, in equilibrium, only a seller with the maximum probability of cheating would ever choose to run a second-price auction. Our work differs in that we focus on the agents' strategies in a second-price auction for a given probability of cheating by the seller. An explicit derivation of the equilibrium strategies then allows us relate first and second-price auctions.

In the same paper, another model of seller cheating is studied. The setting consists of repeated second-price auctions in which the seller decides, in each auction, whether to cheat. The bidders assume that the seller is honest until he is caught cheating, after which they assume that the seller always cheats. Another paper by Rothkopf [19] discusses a form of agent cheating in sealed-bid auctions that also gives rise to auctions that lie between first and second-price auctions. Specifically, agents may submit multiple bids, and then withdraw bids after all of the bids are opened.

Also, a special case of our robustness result for a first-price auction was independently derived by [21], using a different analysis. The focus of that paper was on the maxmin strategy, the strategy that guarantees the agent the highest expected payoff over the space of all possible strategies of the remaining agents. One of the settings considered was a first-price auction in which the agent types are drawn from the uniform distribution. When it is assumed that the competing agents will not play according to a dominated strategy, the payoff of the remaining agent is minimized when each competing agent bids its valuation. Since this is equivalent to the case in our setting in which all other agents are cheating, the maxmin strategy for an agent is to bid according to the symmetric equilibrium bidding strategy.

A different model of bidder cheating was considered in [2] for a multi-unit reverse auction setting. In their model, each bidder privately knows both his cost function and his "trustworthiness", which can be interpreted as his probability of delivery. They present an auction under which it is dominant strategy for each bidder to bid truthfully along both dimensions. The key to this auction is the fact that the payment to the winning bidder is based on the declared trustworthiness of the second-highest bid. A similar incentive compatible mechanism was presented in [16] for a task allocation setting in which each bidder has a privately-known cost of attempting and probability of successfully completing each task.

Work more generally related to ours includes analyses of various forms of collusion. Results on collusion in first and second-price auctions can be found in [12] and [8], respectively. Further opportunities for collusion exist when multiple goods are auctioned. For example, in a simultaneous ascending auction, bidders can encode signals in their bids in order to implicitly collude with other bidders not to compete for the same sets of goods, which occurred in an FCC spectrum auction [5]. Also, in a combinatorial auction, a bidder can profit from bidding under multiple identities, whereby he essentially colludes with himself [23].

An area of related work that can be seen as complementary to ours is that of secure auctions, which takes the point of view of the auction designer. The goals often extend well beyond simply preventing cheating, including properties such as anonymity of the bidders and nonrepudiation of bids. Cryptographic methods are the standard weapon of choice here (see, e.g., [3,6,14,15]). An alternative use of cryptography, proposed in [1], is to bind the online identity of each participant to a single digital certificate. Then, a bidder or seller who is caught cheating can either be punished or made to suffer a permanent loss of reputation.

Existing work covers another dimension along which we could analyze cheating: altering the perceived value of N. In this paper, we have assumed that N is known by all of the bidders. However, in an online setting this assumption is rather tenuous. For example, the seller could arbitrarily increase the perceived N by registering for the auction under different identities. In a first-price auction, the seller has an incentive to do so in order to induce agents to bid closer to their true valuation [6]. However, if agents are aware that the seller has this power, then any communication about N to the agents is "cheap talk", and furthermore is not credible. Thus, in equilibrium the agents would ignore the declared value of N, and bid according to their own prior beliefs about the number of agents. If we assume a common prior, then the setting reduces to the one tackled by [9], which derived the equilibrium bidding strategy of a first-price auction when the number of bidders is drawn from a known distribution but not revealed to any of the bidders. Instead of assuming that the seller can always exploit this power, we could assume that it can only do so with some probability that is known by the agents. The analysis would then proceed in a similar manner as that of our cheating seller model.

The other interesting case of this form of cheating is by bidders in a first-price auction. Bidders would obviously want to decrease the perceived number of agents in order to induce their competition to lower their bids. While it is unlikely that bidders will be able to alter the perceived N arbitrarily, collusion among a group agents in which only one member of the group participates in the auction provides an opportunity to decrease the perceived N. While the non-colluding agents would account for this possibility, as long as they are not certain of the collusion they will still be induced to shave more off of their bids than they would if collusion were not possible. This issue is tackled in [11].

Modifying other assumptions of our setting provides an opportunity for future work. For example, instead of risk-neutral bidders, one can consider risk-averse or risk-seeking bidders. Also, the assumption of independent private values is an important one, and the cases of affiliated and common-value (see, e.g., [13] and [4], respectively) should also be considered. Finally, the setting could be extended so that the auctioneer also may cheat, possibly forming collusive agreements with bidders or the seller.

6 Conclusion

In this paper we presented the equilibria of sealed-bid auctions in which cheating is possible. In addition to providing strategy profiles that are stable against deviations, these results give us with insights into both first and second-price auctions. The results for the case of a cheating seller in a second-price auc-

[6] This assumes that a larger perceived N, and a corresponding lower perceived probability of winning the auction, does not affect an agent's decision of whether or not to bid (due to, for example, the cost of determining its valuation)

tion allow us to relate the two auctions as endpoints along a continuum. The case of agents cheating in a first-price auction shows the robustness of the first-price auction equilibrium when agent types are drawn from the uniform distribution. We also explored the effect of overbidding on the best response bidding strategy for other distributions, and provided sufficient conditions for predicting the direction of this effect. Finally, results from both of our settings allowed us to quantify the expected loss in revenue for a seller due to the possibility of cheating.

A Proofs

Theorem 1 *In a second-price auction in which the seller cheats with probability P^c, it is a Bayes-Nash equilibrium for each agent to bid according to the following strategy:*

$$b_i^*(\theta_i) = \theta_i - \frac{\int_0^{\theta_i} F^{(\frac{N-1}{P^c})}(x)dx}{F^{(\frac{N-1}{P^c})}(\theta_i)} \tag{A.1}$$

Proof:

To find an equilibrium, we start by guessing that there exists an equilibrium in which all agents bid according to the same function $b_i(\theta_i)$, because the game is symmetric. Further, we guess that $b_i(\theta_i)$ is strictly increasing and differentiable over the range $[0, 1]$. We can also assume that $b_i(0) = 0$, because negative bids are not allowed and a positive bid is not rational when the agent's valuation is 0. Note that these are not assumptions on the setting– they are merely limitations that we impose on our search. The fact that we find an equilibrium that satisfies them shows that are indeed valid.

Let $\Phi_i : [0, b_i(1)] \to [0, 1]$ be the inverse function of $b_i(\theta_i)$. That is, it takes a bid for agent i as input and returns the type θ_i that induced this bid. Recall Equation 3:

$$E_{\theta_{-i},\mu^c}\Big[u_i(b(\theta), \mu^c, \theta_i)\Big] = Prob\Big(b_i(\theta_i) > b_{[1]}(\theta_{-i})\Big) \cdot \\ \Big(\theta_i - E_{\theta_{-i},\mu^c}\Big[p_i(b(\theta), \mu^c) \mid \Big(b_i(\theta_i) > b_{[1]}(\theta_{-i})\Big)\Big]\Big)$$

The probability that a single other bid is below $b_i(\theta_i)$ is equal to the *cdf* at the type that would induce a bid equal to $b_i(\theta_i)$, which is formally expressed as $F(\Phi_i(b_i(\theta_i)))$. Since all agents are independent, the probability that all other bids are below agent i's is simply this term raised the $(N-1)$-th power. Thus, we can re-write the expected utility as:

$$E_{\theta_{-i},\mu^c}\Big[u_i(b(\theta),\mu^c,\theta_i)\Big] = F^{N-1}(\Phi_i(b_i(\theta_i)))\cdot \\ \Big(\theta_i - E_{\theta_{-i},\mu^c}\Big[p_i(b(\theta),\mu^c) \mid \Big(b_i(\theta_i) > b_{[1]}(\theta_{-i})\Big)\Big]\Big) \quad \text{(A.2)}$$

We now solve for the expected payment. Plugging Equation 1 (which gives the price for the winning agent) into the term for the expected price in Equation A.2, and then simplifying the expectation yields:

$$\begin{aligned} &E_{\theta_{-i},\mu^c}\Big[p_i(b(\theta),\mu^c) \mid \Big(b_i(\theta_i) > b_{[1]}(\theta_{-i})\Big)\Big] \\ &= E_{\theta_{-i},\mu^c}\Big[\Big(\mu^c\cdot b_i(\theta_i) + (1-\mu^c)\cdot b_{[1]}(\theta_{-i})\Big) \mid \Big(b_i(\theta_i) > b_{[1]}(\theta_{-i})\Big)\Big] \\ &= P^c\cdot b_i(\theta_i) + (1-P^c)\cdot E_{\theta_{-i}}\Big[\Big(b_{[1]}(\theta_{-i})\Big) \mid \Big(b_i(\theta_i) > b_{[1]}(\theta_{-i})\Big)\Big] \\ &= P^c\cdot b_i(\theta_i) + (1-P^c)\cdot\Big[\int_0^{b_i(\theta_i)} b_{[1]}(\theta_{-i})\cdot \\ &\qquad pdf\Big(b_{[1]}(\theta_{-i}) \mid \Big(b_i(\theta_i) > b_{[1]}(\theta_{-i})\Big)\Big) db_{[1]}(\theta_{-i})\Big] \quad \text{(A.3)} \end{aligned}$$

Note that the integral on the last line is taken up to $b_i(\theta_i)$ because we are conditioning on the fact that $b_i(\theta_i) > b_{[1]}(\theta_{-i})$. To derive the *pdf* of $b_{[1]}(\theta_{-i})$ given this condition, we start with the *cdf*. For a given value $b_{[1]}(\theta_{-i})$, the probability that any one agent's bid is less than this value is equal to $F(\Phi_i(b_{[1]}(\theta_{-i})))$. We then condition on the agent's bid being below $b_i(\theta_i)$ by dividing by $F(\Phi_i(b_i(\theta_i)))$. The *cdf* for the $N-1$ agents is then this value raised to the $(N-1)$-th power.

$$cdf\Big[b_{[1]}(\theta_{-i}) \mid \Big(b_i(\theta_i) > b_{[1]}(\theta_{-i})\Big)\Big] = \frac{F^{N-1}(\Phi_i(b_{[1]}(\theta_{-i})))}{F^{N-1}(\Phi_i(b_i(\theta_i)))}$$

The *pdf* is then the derivative of the *cdf* with respect to $b_{[1]}(\theta_{-i})$:

$$\begin{aligned} &pdf\Big[b_{[1]}(\theta_{-i}) \mid \Big(b_i(\theta_i) > b_{[1]}(\theta_{-i})\Big)\Big] \\ &= \frac{N-1}{F^{N-1}(\Phi_i(b_i(\theta_i)))}\cdot F^{N-2}(\Phi_i(b_{[1]}(\theta_{-i})))\cdot f(\Phi_i(b_{[1]}(\theta_{-i})))\cdot \Phi_i'(b_{[1]}(\theta_{-i})) \end{aligned}$$

Substituting the *pdf* into Equation A.3 and pulling terms out of the integral that do not depend on $b_{[1]}(\theta_{-i})$ yields:

$$E_{\theta_{-i},\mu^c}\Big[p_i(b(\theta),\mu^c) \mid \Big(b_i(\theta_i) > b_{[1]}(\theta_{-i})\Big)\Big] = P^c \cdot b_i(\theta_i) + \frac{(1-P^c)\cdot(N-1)}{F^{N-1}(\Phi_i(b_i(\theta_i)))} \cdot \int_0^{b_i(\theta_i)} \Big(b_{[1]}(\theta_{-i}) \cdot F^{N-2}(\Phi_i(b_{[1]}(\theta_{-i}))) \cdot f(\Phi_i(b_{[1]}(\theta_{-i}))) \cdot \Phi_i'(b_{[1]}(\theta_{-i}))\Big)\, db_{[1]}(\theta_{-i})$$

Plugging the expected price back into the expected utility equation (A.2), and distributing $F^{N-1}(\Phi_i(b_i(\theta_i)))$, yields:

$$E_{\theta_{-i},\mu^c}\Big[u_i(b(\theta),\mu^c,\theta_i)\Big] = F^{N-1}(\Phi_i(b_i(\theta_i))) \cdot \theta_i - F^{N-1}(\Phi_i(b_i(\theta_i))) \cdot P^c \cdot b_i(\theta_i) - (1-P^c)\cdot(N-1)\cdot\Big[\int_0^{b_i(\theta_i)} \Big(b_{[1]}(\theta_{-i}) \cdot F^{N-2}(\Phi_i(b_{[1]}(\theta_{-i}))) \cdot f(\Phi_i(b_{[1]}(\theta_{-i}))) \cdot \Phi_i'(b_{[1]}(\theta_{-i}))\Big)\, db_{[1]}(\theta_{-i})\Big]$$

We are now ready to optimize the expected utility by taking the derivative with respect to $b_i(\theta_i)$ and setting it to 0. Note that θ_i is not variable at this point in the proof– we are finding the bid $b_i(\theta_i)$ that maximizes expected utility for any given θ_i. Note also that we do not need to solve the integral, because it will disappear when the derivative is taken (by application of the Fundamental Theorem of Calculus).

$$0 = (N-1)\cdot F^{N-2}(\Phi_i(b_i(\theta_i))) \cdot f(\Phi_i(b_i(\theta_i))) \cdot \Phi_i'(b_i(\theta_i)) \cdot \theta_i - F^{N-1}(\Phi_i(b_i(\theta_i))) \cdot P^c - P^c \cdot (N-1) \cdot F^{N-2}(\Phi_i(b_i(\theta_i))) \cdot f(\Phi_i(b_i(\theta_i))) \cdot \Phi_i'(b_i(\theta_i)) \cdot b_i(\theta_i) - (1-P^c)\cdot(N-1)\cdot\Big[b_i(\theta_i) \cdot F^{N-2}(\Phi_i(b_i(\theta_i))) \cdot f(\Phi_i(b_i(\theta_i))) \cdot \Phi_i'(b_i(\theta_i))\Big]$$

Dividing through by $F^{N-2}(\Phi_i(b_i(\theta_i)))$ and combining like terms yields:

$$0 = \Big[\Big(\theta_i - P^c \cdot b_i(\theta_i) - (1 - P^c) \cdot b_i(\theta_i)\Big) \cdot (N-1) \cdot f(\Phi_i(b_i(\theta_i))) \cdot \Phi_i'(b_i(\theta_i))\Big] - P^c \cdot F(\Phi_i(b_i(\theta_i)))$$

Simplifying the expression and rearranging terms produces:

$$b_i(\theta_i) = \theta_i - \frac{P^c \cdot F(\Phi_i(b_i(\theta_i)))}{(N-1) \cdot f(\Phi_i(b_i(\theta_i))) \cdot \Phi_i'(b_i(\theta_i))}$$

Note that from this point on, we treat θ_i as a variable. Since the above steps for optimizing $b_i(\theta_i)$ apply regardless of the actual value of θ_i, we are implicitly performing this same optimization for all possible values of θ_i and then aggregating the results into a single function $b_i(\theta_i)$.

To further simplify, we use the formula $f'(x) = \frac{1}{g'(f(x))}$, where $g(x)$ is the inverse function of $f(x)$. Plugging in function from our setting gives us: $\Phi_i'(b_i(\theta_i)) = \frac{1}{b_i'(\theta_i)}$. Applying both this equation and $\Phi_i(b_i(\theta_i)) = \theta_i$ yields:

$$b_i(\theta_i) = \theta_i - \frac{P^c \cdot F(\theta_i) \cdot b_i'(\theta_i)}{(N-1) \cdot f(\theta_i)} \tag{A.4}$$

While we are not able to derive a solution directly, this equation allows us to quickly verify a guessed solution that is based on the solution for the first-price auction (see, e.g., [18]).

$$b_i(\theta_i) = \theta_i - \frac{\int_0^{\theta_i} F^{(\frac{N-1}{P^c})}(x)dx}{F^{(\frac{N-1}{P^c})}(\theta_i)} \tag{A.5}$$

To verify that this is indeed the solution, we first take its derivative:

$$b_i'(\theta_i) \quad = \quad 1 \; - \; \frac{F^{(2 \cdot \frac{N-1}{P^c})}(\theta_i) - \frac{N-1}{P^c} \cdot F^{(\frac{N-1}{P^c} - 1)}(\theta_i) \cdot f(\theta_i) \cdot \int_0^{\theta_i} F^{(\frac{N-1}{P^c})}(x)dx}{F^{(2 \cdot \frac{N-1}{P^c})}(\theta_i)}$$

This simplifies to:

$$b_i'(\theta_i) = \frac{\frac{N-1}{P^c} \cdot f(\theta_i) \cdot \int_0^{\theta_i} F^{(\frac{N-1}{P^c})}(x)dx}{F^{(\frac{N-1}{P^c}+1)}(\theta_i)}$$

We then plug this derivative into the equation we derived (A.4):

$$b_i(\theta_i) = \theta_i - \frac{P^c \cdot F(\theta_i) \cdot \frac{N-1}{P^c} \cdot f(\theta_i) \cdot \int_0^{\theta_i} F^{(\frac{N-1}{P^c})}(x)dx}{(N-1) \cdot f(\theta_i) \cdot F^{(\frac{N-1}{P^c}+1)}(\theta_i)}$$

Cancelling terms yields Equation A.5, verifying that our guessed solution is correct. □

Corollary 2 *In a second-price auction in which the seller cheats with probability P^c, and $F(\theta_i) = \theta_i$, it is a Bayes-Nash equilibrium for each agent to bid according to the following strategy:*

$$b_i^*(\theta_i) = \frac{N-1}{N-1+P^c}\theta_i$$

Proof:

Plugging $F(\theta_i) = \theta_i$ into Equation A.1 (repeated as A.5), yields:

$$\begin{aligned} b_i(\theta_i) = \theta_i - \frac{\int_0^{\theta_i} x^{(\frac{N-1}{P^c})}dx}{\theta_i^{(\frac{N-1}{P^c})}} &= \theta_i - \frac{\frac{P^c}{N-1+P^c}\theta_i^{(\frac{N-1+P^c}{P^c})}}{\theta_i^{(\frac{N-1}{P^c})}} \\ = \theta_i - \frac{P^c}{N-1+P^c} \cdot \theta_i &= \frac{N-1}{N-1+P^c} \cdot \theta_i \end{aligned}$$

□

Theorem 3 *In a first-price auction in which each agent cheats with probability P^a, it is a Bayes-Nash equilibrium for each non-cheating agent i to bid according to the strategy that is a fixed point of the following equation:*

$$b_i^*(\theta_i) = \theta_i - \frac{\int_0^{\theta_i} \Big(P^a \cdot F(b_i^*(x)) + (1-P^a) \cdot F(x)\Big)^{(N-1)} dx}{\Big(P^a \cdot F(b_i^*(\theta_i)) + (1-P^a) \cdot F(\theta_i)\Big)^{(N-1)}}$$

Proof:

We make the same guesses about the equilibrium strategy to aid our search as we did in the proof of Theorem 1. When simplifying the expectation of this setting's utility equation (6), we use the fact that the probability that agent i will have a higher bid than another honest agent is still $F(\Phi_i(b_i(\theta_i)))$, while the probability is $F(b_i(\theta_i))$ if the other agent cheats. The probability that agent i beats a single other agent is then a weighted average of these two probabilities. Thus, we can write agent i's expected utility as:

$$E_{\theta_{-i},\mu^a}\Big[u_i(b(\theta),\mu^a,\theta_i)\Big] = \Big(\theta_i - b_i(\theta_i)\Big)\cdot\Big[P^a\cdot F(b_i(\theta_i)) + (1-P^a)\cdot F(\Phi_i(b_i(\theta_i)))\Big]^{N-1}$$

As before, to find the equilibrium $b_i(\theta_i)$, we take the derivative with respect to $b_i(\theta_i)$ and set it to zero:

$$0 = \Bigg[\Big(\theta_i - b_i(\theta_i)\Big)\cdot(N-1)\cdot\Big(P^a\cdot F(b_i(\theta_i)) + (1-P^a)\cdot F(\Phi_i(b_i(\theta_i)))\Big)^{N-2}\cdot \Big(P^a\cdot f(b_i(\theta_i)) + (1-P^a)\cdot f(\Phi_i(b_i(\theta_i)))\cdot\Phi_i'(b_i(\theta_i))\Big)\Bigg] - \Big(P^a\cdot F(b_i(\theta_i)) + (1-P^a)\cdot F(\Phi_i(b_i(\theta_i)))\Big)^{N-1}$$

Aggregating this result over all possible values of θ_i (which is now a variable), applying the equations $\Phi_i'(b_i(\theta_i)) = \frac{1}{b_i'(\theta_i)}$ and $\Phi_i(b_i(\theta_i)) = \theta_i$, and dividing through, produces:

$$0 = \Bigg[\Big(\theta_i - b_i(\theta_i)\Big)\cdot(N-1)\cdot\Big(P^a\cdot f(b_i(\theta_i)) + (1-P^a)\cdot f(\theta_i)\cdot\frac{1}{b_i'(\theta_i)}\Big)\Bigg] - \Big(P^a\cdot F(b_i(\theta_i)) + (1-P^a)\cdot F(\theta_i)\Big)$$

Rearranging terms yields:

$$b_i(\theta_i) = \theta_i - \frac{\Big(P^a\cdot F(b_i(\theta_i)) + (1-P^a)\cdot F(\theta_i)\Big)\cdot b_i'(\theta_i)}{(N-1)\cdot\Big(P^a\cdot f(b_i(\theta_i))\cdot b_i'(\theta_i) + (1-P^a)\cdot f(\theta_i)\Big)} \tag{A.6}$$

In this setting, because we leave $F(\cdot)$ unspecified, we cannot present a closed form solution. However, we can simplify the expression by removing its dependence on $b_i'(\theta_i)$.

$$b_i(\theta_i) = \theta_i - \frac{\int_0^{\theta_i} \Big(P^a \cdot F(b_i(x)) + (1-P^a) \cdot F(x)\Big)^{(N-1)} dx}{\Big(P^a \cdot F(b_i(\theta_i)) + (1-P^a) \cdot F(\theta_i)\Big)^{(N-1)}} \tag{A.7}$$

To verify Equation A.7, first take its derivative with respect to θ_i:

$$b_i'(\theta_i) = 1 - \Big[1 - (N-1) \cdot \Big(P^a \cdot F(b_i(\theta_i)) + (1-P^a) \cdot F(\theta_i)\Big)^{(N-2)} \cdot \\ \Big(P^a \cdot f(b_i(\theta_i)) \cdot b_i'(\theta_i)) + (1-P^a) \cdot f(\theta_i)\Big) \cdot \\ \frac{\int_0^{\theta_i} \Big(P^a \cdot F(b_i(x)) + (1-P^a) \cdot F(x)\Big)^{(N-1)} dx}{\Big(P^a \cdot F(b_i(\theta_i)) + (1-P^a) \cdot F(\theta_i)\Big)^{2(N-1)}}\Big]$$

This equation simplifies to:

$$b_i'(\theta_i) = (N-1) \cdot \Big(P^a \cdot f(b_i(\theta_i)) \cdot b_i'(\theta_i)) + (1-P^a) \cdot f(\theta_i)\Big) \cdot \\ \frac{\int_0^{\theta_i} \Big(P^a \cdot F(b_i(x)) + (1-P^a) \cdot F(x)\Big)^{(N-1)} dx}{\Big(P^a \cdot F(b_i(\theta_i)) + (1-P^a) \cdot F(\theta_i)\Big)^{N}}$$

Plugging this equation into the $b_i(\theta_i)$ in the numerator of Equation A.6 yields Equation A.7, verifying the solution. □

Corollary 4 *In a first-price auction in which each agent cheats with probability P^a, and $F(\theta_i) = \theta_i$, it is a Bayes-Nash equilibrium for each non-cheating agent to bid according to the strategy $b_i^*(\theta_i) = \frac{N-1}{N}\theta_i$.*

Proof:

Instantiating the fixed point equation (7, repeated as A.7) with $F(\theta_i) = \theta_i$ yields:

$$b_i(\theta_i) = \theta_i - \frac{\int_0^{\theta_i} \Big(P^a \cdot b_i(x) + (1-P^a) \cdot x\Big)^{(N-1)} dx}{\Big(P^a \cdot b_i(\theta_i) + (1-P^a) \cdot \theta_i\Big)^{(N-1)}}$$

We can then plug the strategy $b_i(\theta_i) = \frac{N-1}{N}\theta_i$ into this equation in order to verify that it is a fixed point.

$$
\begin{aligned}
b_i(\theta_i) &= \theta_i - \frac{\int_0^{\theta_i}\left(P^a \cdot \frac{N-1}{N}x + (1-P^a)\cdot x\right)^{(N-1)} dx}{\left(P^a \cdot \frac{N-1}{N}\theta_i + (1-P^a)\cdot \theta_i\right)^{(N-1)}} \\
&= \theta_i - \frac{\int_0^{\theta_i} x^{(N-1)}\, dx}{\theta_i^{(N-1)}} \quad = \theta_i - \frac{\frac{1}{N}\theta_i^N}{\theta_i^{(N-1)}} \quad = \frac{N-1}{N}\theta_i
\end{aligned}
$$

□

Theorem 5 *In a first-price auction where $F(\theta_i) = \theta_i$, if each agent $j \neq i$ bids according a strategy $b_j(\theta_j) = \frac{N-1+\alpha_j}{N}\theta_j$, where $\alpha_j \geq 0$, then it is a best response for the remaining agent i to bid according to the strategy $b_i^R(\theta_i) = \frac{N-1}{N}\theta_i$.*

Proof:

We again use $\Phi_j : [0, b_j(1)] \to [0,1]$ as the inverse of $b_j(\theta_j)$. For all $j \neq i$ in this setting, $\Phi_j(x) = \frac{N}{N-1+\alpha_j}x$. The probability that agent i has a higher bid than a single agent j is $F(\Phi_j(b_i^R(\theta_i))) = \frac{N}{N-1+\alpha_j}b_i^R(\theta_i)$. Note, however, that since $\Phi_j(\cdot)$ is only defined over the range $[0, b_j(1)]$, it must be the case that $b_i^R(1) \leq b_j(1)$, which is why $\alpha_j \geq 0$ is necessary, in addition to being sufficient. Assuming that $b_i^R(\theta_i) = \frac{N-1}{N}\theta_i$, then indeed $\Phi_j(b_i^R(\theta_i))$ is always well-defined. We will now show that this assumption is correct. The expected utility for agent i can then be written as:

$$
\begin{aligned}
E_{\theta_{-i}}\Big[u_i((b_i^R(\theta_i), b_{-i}(\theta_{-i})), \theta_i)\Big] &= \Big[\Pi_{j\neq i}\frac{N}{N-1+\alpha_j}b_i^R(\theta_i)\Big] \cdot \Big[\theta_i - b_i^R(\theta)\Big] \\
&= \Big[\Pi_{j\neq i}\frac{N}{N-1+\alpha_j}\Big] \cdot \Big[\theta_i \cdot \Big(b_i^R(\theta_i)\Big)^{(N-1)} - \Big(b_i^R(\theta_i)\Big)^N\Big]
\end{aligned}
$$

Taking the derivative with respect to $b_i^R(\theta_i)$, setting it to zero, and dividing out $\Pi_{j\neq i}\frac{N}{N-1+\alpha_j}$ yields:

$$
0 = \theta_i \cdot (N-1) \cdot \Big(b_i^R(\theta_i)\Big)^{(N-2)} - N \cdot \Big(b_i^R(\theta_i)\Big)^{(N-1)}
$$

This simplifies to the solution: $b_i^R(\theta_i) = \frac{N-1}{N}\theta_i$. □

Theorem 6 *In a first-price auction, if* $g'(\theta_i) > 0$ *and* $g''(\theta_i) < 0$ *(resp.* $g''(\theta_i) = 0$ *and* $g''(\theta_i) > 0$*) hold for all* θ_i*, where* $g(\theta_i) = \frac{F(\theta_i)}{f(\theta_i)}$*, and if each agent* $j \neq i$ *bids according the strategy* $b_j(\theta_j) = \theta_j$*, then, for all* $\theta_i > 0$*, the best response bidding strategy* $b_i^R(\theta_i)$ *for agent* i *is strictly less than (resp. equal to, strictly greater than) the symmetric equilibrium bidding strategy* $b_i^*(\theta_i)$*.*

Proof: We will only show the proof of the first case (that $b_i^R(\theta_i) < b_i^*(\theta_i)$ if $g''(\theta_i) < 0$), and point out the places in which the proof changes for the other two cases.

The expected utility for agent i when each other agent always bids its type is:

$$E_{\theta_{-i}}\Big[u_i((b_i^R(\theta_i), b_{-i}(\theta_{-i})), \theta_i)\Big] = F^{(N-1)}(b_i^R(\theta_i)) \cdot \Big(\theta_i - b_i^R(\theta_i)\Big)$$

Taking the derivative with respect to $b_i^R(\theta_i)$:

$$\begin{aligned}
&\frac{d}{db_i^R(\theta_i)} E_{\theta_{-i}}\Big[u_i((b_i^R(\theta_i), b_{-j}(\theta_{-j})), \theta_i)\Big] \\
&\quad = (N-1) \cdot F^{(N-2)}(b_i^R(\theta_i)) \cdot f(b_i^R(\theta_i)) \cdot \Big(\theta_i - b_i^R(\theta_i)\Big) - F^{(N-1)}(b_i^R(\theta_i)) \\
&\qquad\quad = F^{(N-2)}(b_i^R(\theta_i)) \cdot f(b_i^R(\theta_i)) \cdot \Big[(N-1) \cdot \Big(\theta_i - b_i^R(\theta_i)\Big) - g(b_i^R(\theta_i))\Big]
\end{aligned}$$

Setting the derivative to zero, the best response bidding strategy must satisfy $b_i^R(\theta_i) = \theta_i - \frac{g(b_i^R(\theta_i))}{N-1}$ (note that $f(b_i^R(\theta_i))$ cannot be zero because of the assumption that $F(\cdot)$ is strictly increasing). Due to the condition that $g'(\theta_i) > 0$, this equation has only one solution. Thus, it suffices to check that the derivative is negative (resp. equal, positive for the other two cases) when we set the best response bidding strategy equal to the symmetric equilibrium bidding strategy. Thus, we make the substitution $b_i^R(\theta_i) = b_i^*(\theta_i)$ and check that the following inequality holds:

$$\begin{aligned}
0 &> (N-1) \cdot \Big(\theta_i - b_i^*(\theta_i)\Big) - g(b_i^*(\theta_i)) \\
\frac{g(b_i^*(\theta_i))}{N-1} &> \theta_i - b_i^*(\theta_i)
\end{aligned}$$

Plugging the symmetric equilibrium bidding strategy $b_i^*(\theta_i) = \theta_i - \frac{\int_0^{\theta_i} F^{(\frac{N-1}{P^c})}(x)dx}{F^{(\frac{N-1}{P^c})}(\theta_i)}$ into the RHS yields:

$$\frac{g(b_i^*(\theta_i))}{N-1} > \theta_i - (\theta_i - \frac{\int_0^{\theta_i} F^{(N-1)}(x)dx}{F^{(N-1)}(\theta_i)})$$
$$F^{(N-1)}(\theta_i) \cdot \frac{g(b_i^*(\theta_i))}{N-1} > \int_0^{\theta_i} F^{(N-1)}(x)dx$$

Because both sides are 0 at $\theta_i = 0$, it suffices to check that the LHS has a larger first derivative than the RHS.

$$F^{(N-1)}(\theta_i) \cdot \frac{g'(b_i^*(\theta_i))}{N-1} \cdot \frac{d}{d\theta_i}(b_i^*(\theta_i)) + F^{(N-2)}(\theta_i) \cdot f(\theta_i) \cdot g(b_i^*(\theta_i)) > F^{(N-1)}(\theta_i)$$
$$F(\theta_i) \cdot \frac{g'(b_i^*(\theta_i))}{N-1} \cdot \left[1 - \frac{F^{2(N-1)}(\theta_i) - (N-1) \cdot F^{(N-2)}(\theta_i) \cdot f(\theta_i) \cdot \int_0^{\theta_i} F^{(N-1)}(x)dx}{F^{2(N-1)}(\theta_i)}\right] +$$
$$f(\theta_i) \cdot g(b_i^*(\theta_i)) > F(\theta_i)$$
$$g'(b_i^*(\theta_i)) \cdot \frac{\int_0^{\theta_i} F^{(N-1)}(x)dx}{F^{(N-1)}(\theta_i)} + g(b_i^*(\theta_i)) > \frac{F(\theta_i)}{f(\theta_i)}$$
$$g'(b_i^*(\theta_i)) \cdot \frac{\int_0^{\theta_i} F^{(N-1)}(x)dx}{F^{(N-1)}(\theta_i)} > g(\theta_i) - g(b_i^*(\theta_i))$$

By the Mean Value Theorem, we know that there exists some value x satisfying $b_i^*(\theta_i) < x < \theta_i$) such that:

$$g'(x) \cdot (\theta_i - b_i^*(\theta_i)) = g(\theta_i) - g(b_i^*(\theta_i))$$

This equation, together with the fact that $g''(\theta_i) < 0$ (resp. $g''(\theta_i) = 0$ and $g''(\theta_i) > 0$), implies the above inequality, completing the proof.

$$g'(b_i^*(\theta_i)) \cdot (\theta_i - b_i^*(\theta_i)) > g(\theta_i) - g(b_i^*(\theta_i))$$
$$g'(b_i^*(\theta_i)) \cdot (\frac{\int_0^{\theta_i} F^{(N-1)}(x)dx}{F^{(N-1)}(\theta_i)}) > g(\theta_i) - g(b_i^*(\theta_i))$$

□

References

[1] S. Ba, A. Whinston, and H. Zhang, *Building trust in online auction markets through an economic incentive mechanism*, Decision Support Systems **35** (2003),

273–286.

[2] S. Brainov and T. Sandholm, *Auctions with untrustworthy bidders*, Proceedings of IEEE Conference on Electronic Commerce, 2003.

[3] F. Brandt, *Fully private auctions in a constant number of rounds*, Proceedings of Financial Cryptography Conference, 2003.

[4] E. Capen, , R. Clapp, and W. Campbell, *Bidding in high risk situations*, Journal of Petroleum Technology **23** (1971), 641–653.

[5] P. Cramton and J. Schwartz, *Collusive bidding in the fcc spectrum auctions*, Contributions to Economic Analysis & Policy **1** (2002).

[6] M. Franklin and M. Reiter, *The Design and Implementation of a Secure Auction Service*, Proc. IEEE Symp. on Security and Privacy, 1995.

[7] D. Fudenberg and J. Tirole, *Game theory*, MIT Press, 1991.

[8] D. Graham and R. Marshall, *Collusive bidder behavior at single-object second-price and english auctions*, Journal of Political Economy **95** (1987), 579–599.

[9] R. Harstad, J. Kagel, and D. Levin, *Equilibrium bid functions for auctions with an uncertain number of bidders*, Economic Letters **33** (1990), 35–40.

[10] P. Klemperer, *Auction theory: A guide to the literature*, Journal of Economic Surveys **13** (1999), no. 3, 227–286.

[11] K. Leyton-Brown, Y. Shoham, and M. Tennenholtz, *Bidding clubs in first-price auctions*, AAAI-02, 2002.

[12] R. McAfee and J. McMillan, *Bidding rings*, The American Economic Review **71** (1992), 579–599.

[13] P. Milgrom and R. Weber, *A theory of auctions and competitive bidding*, Econometrica **50** (1982), 1089–1122.

[14] M. Naor, B. Pinkas, and R. Sumner, *Privacy preserving auctions and mechanism design*, EC-99, 1999.

[15] H. Nurmi and A. Salomaa, *Cryptographic protocols for vickrey auctions*, Group Decisions and Negotiation **2** (1993), 363–373.

[16] R. Porter, A. Ronen, Y. Shoham, and M. Tennenholtz, *Mechanism design with execution uncertainty*, Proceedings of UAI-02, 2002.

[17] R. Porter and Y. Shoham, *On cheating in sealed-bid auctions*, EC'03, 2003.

[18] J. Riley and W. Samuelson, *Optimal auctions*, American Economic Review **71** (1981), no. 3, 381–392.

[19] M. Rothkopf, *On auctions with withdrawable winning bids*, Marketing Science **10** (1991), 40–57.

[20] M Rothkopf and R. Harstad, *Two models of bid-taker cheating in vickrey auctions*, The Journal of Business **68** (1995), no. 2, 257–267.

[21] M. Tennenholtz, *Competitive safety analysis: Robust decision-making in multi-agent systems*, Journal of Artificial Intelligence **17** (2002), 363–378.

[22] W. Vickrey, *Counterspeculations, auctions, and competitive sealed tenders*, Journal of Finance **16** (1961), 15–27.

[23] M. Yokoo, Y. Sakurai, and S. Matsubara, *Robust combinatorial auction protocol against false-name bids*, Artificial Intelligence **130** (2001), no. 2, 167–181.