On Cheating in Sealed-Bid Auctions

Ryan Porter
rwporter@stanford.edu

Yoav Shoham
shoham@stanford.edu

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 spies on the bids of a second-price auction and then inserts a fake 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 deciding on 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.

Categories and Subject Descriptors

J.4 [**Computer Applications**]: Social and Behavioral Sciences—*Economics*

General Terms

Economics, Security

Keywords

Game Theory, Sealed-Bid Auctions, Cheating

1. INTRODUCTION

Among the types of auctions commonly used in practice, sealed-bid auctions are a good practical choice because they require little communication and can be completed almost instantly. Each bidder simply submits a bid, and the winner is immediately determined. However, sealed-bid auctions do 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 that another bidder gave directly to the auctioneer. However, in an online auction the auctioneer is often little more than a server with questionable security; and, since all participants 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.

We investigate this type of cheating along two dimensions: whether it is the auctioneer or a bidder who cheats, and which variant (either first or second-price) of the sealed-bid auction is used. Note that two of these cases are trivial. In our setting, 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 truthfully. This leaves us with two cases that we examine in detail.

A seller can profitably cheat in a second-price auction by looking at the bids before the auction clears and submitting an extra bid. This possibility was pointed out as early as the seminal paper [12] 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. Possible motivation for this knowledge could be a recently published exposé on seller cheating in eBay auctions. 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 a common 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

[1]This work was supported in part by DARPA grant F30602-00-2-0598.

EC'03, June 9–12, 2003, San Diego, California, USA.

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 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. We conclude this section by exploring other distributions, including some in which the presence of cheating bidders actually induces an honest bidder to lower its bid.

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 cheating in a second price auction. Section 3 covers the case of bidders cheating in a first-price auction. In Section 4, we quantify the effects that the possibility of cheating has on 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 are found 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 [2] and [6], respectively.

2.1 Formulation

The setting consists of N bidders, or agents, (indexed by $i = 1, \cdots, n$) and a seller. Each agent has a type $\theta_i \in [0,1]$, drawn from a continuous range, which represents the agent's 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 only the distribution over the possible types of the other agents. A bidding strategy for an agent $b_i : [0,1] \rightarrow [0,1]$ 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, but ties complicate the formulation. Fortunately, we can ignore the case of ties in this paper because our continuity assumption will make them a zero probability event in equilibrium. 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 will pay 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.[5] We can then write the payment of the winning agent as follows.

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

Let $\mu(\cdot)$ be an indicator function that takes an inequality as an argument and returns is 1 if it holds, 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 price if it does.

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

We will be concerned with the expected utility of an agent, with the expectation taken over the types of the other agents and over whether or not the seller cheats. By pushing the expectation inward so that it is only over the price (conditioned on the agent winning the auction), we can write the expected utility as:

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

We assume that all agents are rational, expected utility maximizers. 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:

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

2.2 Equilibrium

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

[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, but, 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.

[5]Note that common knowledge is not necessary for the existence of an equilibrium.

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 $P^c = 1$.

2.3 Continuum of Auctions

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 price paid by the winning agent is determined by a weighted coin flip: with probability P^c it 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. Furthermore, the fact that this probabilistic auction satisfies the properties required for the Revenue Equivalence Theorem (see, e.g., [2]) provides a way to verify that the bidding strategy in Equation 5 is the symmetric equilibrium of this auction (see the alternative proof of Theorem 1 in the appendix).

2.4 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:*

$$b_i(\theta_i) = \frac{N-1}{N-1+P^c}\theta_i \quad (6)$$

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 is pointless 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 the uniform distribution, in which we show that the possibility of cheating agents does not change the equilibrium strategy of honest agents. This result 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 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. Otherwise, this auction is 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.

$$u_i(b(\theta), \mu^a, \theta_i) = \Big(\theta_i - b_i(\theta)\Big) \cdot \prod_{j \neq i} \Big[\mu^{a_j} \cdot \mu\big(b_i(\theta_i) > \theta_j\big) + (1 - \mu^{a_j}) \cdot \mu\big(b_i(\theta_i) > b_j(\theta_j)\big)\Big] \quad (7)$$

3.2 Equilibrium

Our goal is to find the equilibrium in which all cheating agents use their dominant strategy of bidding truthfully and honest agents bid according to a symmetric bidding strategy. Since we have left $F(\cdot)$ unspecified, we cannot present 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:*

$$b_i(\theta_i) = \theta_i - \frac{\int_0^{\theta_i} \Big(P^a \cdot F(b_i(x)) + (1 - P^a) \cdot F(x)\Big)^{(N-1)} dx}{\Big(P^a \cdot F(b_i(\theta_i)) + (1 - P^a) \cdot F(\theta_i)\Big)^{(N-1)}} \quad (8)$$

3.3 Special Case: Uniform Distribution

Since we could not solve Equation 8 in the general case, we can only see how the possibility of cheating affects the equilibrium bidding strategy for particular instances of $F(\cdot)$. A natural place to start is 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 the same strategy. Since in the current setting the cheating agents do not shave their bid 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 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 robust 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 all other agents shave their 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. Note that this result holds even if other agents are shaving their bid by a negative fraction, and are thus irrationally bidding above their type.

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 $b_i(\theta_i) = \frac{N-1}{N}\theta_i$.*

Obviously, these strategy profiles are not equilibria (unless each $\alpha_j = 0$), because each agent j has an incentive to set $\alpha_j = 0$. The point of this theorem is that a wide range of possible beliefs that an agent can hold about the strategies of the other agents will all lead him to play the equilibrium strategy. This is important because 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, if the equilibrium strategy is a best response to a large set of strategy profiles that are out of equilibrium, then it seems much more plausible that the agents will indeed converge on this equilibrium.

It is important to note, though, that while this equilibrium is robust against arbitrary deviations to strategies that shave less, it is not robust to even a single agent shaving more off of its bid. In fact, if we take any strategy profile consistent with the conditions of Theorem 5 and change a single agent j's strategy so that its corresponding α_j is negative, then agent i's best response is to shave more than $\frac{1}{N}$ off of its bid.

3.4 Effects of Overbidding for Other Distributions

A natural question is whether the best response bidding strategy is similarly robust to overbidding by competing agents for other distributions. It turns out that Theorem 5 holds for all distributions of the form $F(\theta_i) = (\theta_i)^k$, where k is some positive integer. However, 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 an agent should actually shave its bid more when other agents shave their bids less. In the example we present for this distribution (with the details in the appendix), there are only two players and the deviation by one agent is to bid his type. However, it can be generalized to a higher number of agents and to other deviations.

EXAMPLE 1. *In a first-price auction where $F(\theta_i) = \frac{\theta_i^2+\theta_i}{2}$ and $N = 2$, if agent 2 always bids its type ($b_2(\theta_2) = \theta_2$), then, for all $\theta_1 > 0$, agent 1's best response bidding strategy is strictly less than the bidding strategy of the symmetric equilibrium.*

We also note that the same result holds for the normalized exponential distribution ($F(\theta_i) = \frac{e^{\theta_i}-1}{e-1}$).

It is certainly the case that distributions can be found that support the intuition given above that agents should shave their bid less when other agents are doing likewise. 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 $F(\theta_i) = \frac{e-e^{(1-\theta_i)}}{e-1}$.

It would be useful to relate the direction of the change in the best response bidding strategy to a general condition on $F(\cdot)$. Unfortunately, we were not able to find such a condition, in part because the integral in the symmetric bidding strategy of a first-price auction cannot be solved without knowing $F(\cdot)$ (or at least some restrictions on it). We do note, however, that the sign of the second derivative of $\big(F(\theta_i)/f(\theta_i)\big)$ is an accurate predictor for all of the distributions that we considered.

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, 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 the fact that the agents will shave their bids. 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 so straightforward. If $P^a = 1$ (each agent cheats with certainty), then we simply have a 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.[6] It turns out that setting $P^c = 0.2$ would make the expected loss of these two settings comparable. While this comparison between the settings is unlikely to be useful for a seller, it is interesting to note that agent suspicions of possible cheating by the seller are in some sense worse than agents actually cheating themselves.

	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. RELATED WORK

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, a bidder's only source of information about N could be a counter that the seller places on the auction webpage, or a statement by the seller about the number of potential bidders who have indicated that they will participate. In these cases, the seller could arbitrarily manipulate the perceived N. In a first-price auction, the seller obviously has an incentive to increase the perceived value of N in order to induce agents to bid closer to their true valuation. 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 make the natural assumption of a common prior, then the setting reduces to the one tackled by [5], which derived the equilibrium bidding strategies of a first-price auction when the number of bidders is drawn from a known distribution but not revealed to any of the bidders. Of course, 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 unreasonable for bidders to be able to alter the perceived N arbitrarily, collusion provides an opportunity to decrease the perceived N by having only one of a group of colluding agents participate in the auction. 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 the collusion did not take place. This issue is tackled in [7].

Other types of collusion are of course related to the general topic of cheating in auctions. Results on collusion in first and second-price auctions can be found in [8] and [3], respectively.

The work most closely related to our first setting is [11], 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 who decides between running a first-price auction (where profitable cheating is never possible) or 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.

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 an 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 [1, 4, 9]).

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 auction allow us to relate the two auctions as endpoints along a continuum. The case of agents cheating in a first-price auction showed 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 showed that even for relatively simple distributions it can be positive, negative, or neutral. 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.

7. REFERENCES

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

[2] D. Fudenberg and J. Tirole. *Game Theory.* MIT Press, 1991.

[6]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.

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

[4] M. Harkavy, J. D. Tygar, and H. Kikuchi. Electronic auctions with private bids. In *Proceedings of the 3rd USENIX Workshop on Electronic Commerce*, 1998.

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

[6] P. Klemperer. Auction theory: A guide to the literature. *Journal of Economic Surveys*, 13(3):227–286, 1999.

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

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

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

[10] J. Riley and W. Samuelson. Optimal auctions. *American Economic Review*, 71(3):381–392, 1981.

[11] M. Rothkopf and R. Harstad. Two models of bid-taker cheating in vickrey auctions. *The Journal of Business*, 68(2):257–267, 1995.

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

APPENDIX

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:*

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

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.

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 that of agent i is equal to the *cdf* at the type that would induce a bid equal to that of agent i, which is formally written 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:

$$E_{\theta_{-i},\mu^c}\big[u_i(b(\theta), \mu^c, \theta_i)\big] = F^{N-1}(\Phi_i(b_i(\theta_i)))\cdot \\ \Big(\theta_i - E_{\theta_{-i},\mu^c}\Big[p_i(b(\theta), \mu^c) \mid \Big(b_i(\theta_i) > b_{[1]}(\theta_{-i})\Big)\Big]\Big) \tag{9}$$

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 9, and then simplifying the expectation yields:

$$E_{\theta_{-i},\mu^c}\Big[p_i(b(\theta), \mu^c) \mid \Big(b_i(\theta_i) > b_{[1]}(\theta_{-i})\Big)\Big] \\ = E_{\theta_{-i},\mu^c}\Big[\Big(\mu^c \cdot b_i(\theta_i) + (1-\mu^c)\cdot b_{[1]}(\theta_{-i})\Big) \mid \\ \Big(b_i(\theta_i) > b_{[1]}(\theta_{-i})\Big)\Big] \\ = P^c \cdot b_i(\theta_i) + (1-P^c)\cdot E_{\theta_{-i}}\Big[\Big(b_{[1]}(\theta_{-i})\Big) \mid \\ \Big(b_i(\theta_i) > b_{[1]}(\theta_{-i})\Big)\Big] \\ = P^c \cdot b_i(\theta_i) + (1-P^c)\cdot \Big[\int_0^{b_i(\theta_i)} b_{[1]}(\theta_{-i})\cdot \\ pdf\Big(b_{[1]}(\theta_{-i}) \mid \big(b_i(\theta_i) > b_{[1]}(\theta_{-i})\big)\Big) db_{[1]}(\theta_{-i})\Big] \tag{10}$$

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 $b_{[1]}(\theta_{-i})$:

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

Substituting the *pdf* into Equation 10 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 (9), 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 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:

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

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:

$$b_i(\theta_i) = \theta_i - \frac{P^c\cdot F(\theta_i)\cdot b_i'(\theta_i)}{(N-1)\cdot f(\theta_i)} \tag{11}$$

Attempts at a derivation of the solution from this point proved fruitless, but we are at a point now where a guessed solution can be quickly verified. We used the solution for the first-price auction (see, e.g., [10]) as our starting point to find the answer:

$$b_i(\theta_i) = \theta_i - \frac{\int_0^{\theta_i} F^{(\frac{N-1}{P^c})}(x)dx}{F^{(\frac{N-1}{P^c})}(\theta_i)} \tag{12}$$

To verify the solution, we first take its derivative:

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

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 (11):

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

Cancelling terms yields Equation 12, verifying that our guessed solution is correct. □

Alternative Proof of Theorem 1:

The following proof uses the Revenue Equivalence Theorem (RET) and the probabilistic auction given as an interpretation of our cheating seller setting.

In a first-price auction without the possibility of cheating, the expected payment for an agent with type θ_i is simply the product of its bid and the probability that this bid is the highest. For the symmetric equilibrium, this is equal to:

$$F^{(N-1)}(\theta_i)\cdot\Big[\theta_i - \frac{\int_0^{\theta_i} F^{(N-1)}(x)dx}{F^{(N-1)}(\theta_i)}\Big]$$

For our probabilistic auction, the expected payment of the winning agent is a weighted average of its bid and the second highest bid. For the $b_i(\cdot)$ we found in the original interpretation of the setting, it can be written as follows.

$$F^{(N-1)}(\theta_i)\cdot\Big[P^c\Big(\theta_i - \frac{\int_0^{\theta_i} F^{(\frac{N-1}{P^c})}(x)dx}{F^{(\frac{N-1}{P^c})}(\theta_i)}\Big) + (1-P^c)\Big(\frac{1}{F^{(N-1)}(\theta_i)}\cdot\int_0^{\theta_i}\Big(x - \frac{\int_0^x F^{(\frac{N-1}{P^c})}(y)dy}{F^{(\frac{N-1}{P^c})}(x)}\Big)\cdot(N-1)\cdot F^{(N-2)}(x)\cdot f(x)dx\Big)\Big]$$

By the RET, the expected payments will be the same in the two auctions. Thus, we can verify our equilibrium bidding strategy by showing that the expected payment in the two auctions is equal. Since the expected payment is zero at $\theta_i = 0$ for both functions, it suffices to verify that the derivatives of the expected payment functions with respect to θ_i are equal, for an arbitrary value θ_i. Thus, we need to verify the following equation:

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

This simplifies to:

$$0 = P^c \cdot \Big[\frac{(\frac{N-1}{P^c}) \cdot F^{(N-2)}(\theta_i) \cdot f(\theta_i) \cdot \int_0^{\theta_i} F^{(\frac{N-1}{P^c})}(x)dx}{F^{(\frac{N-1}{P^c})}(\theta_i)} +$$

$$(N-1) \cdot F^{(N-2)}(\theta_i) \cdot f(\theta_i) \cdot \Big(-\frac{\int_0^{\theta_i} F^{(\frac{N-1}{P^c})}(x)dx}{F^{(\frac{N-1}{P^c})}(\theta_i)} \Big) \Big] +$$

$$(1-P^c) \Big[\Big(-\frac{\int_0^{\theta_i} F^{(\frac{N-1}{P^c})}(y)dy}{F^{(\frac{N-1}{P^c})}(\theta_i)} \Big) \cdot (N-1) \cdot F^{(N-2)}(\theta_i) \cdot f(\theta_i) \Big]$$

After distributing P^c, the right-hand side of this equation cancels out, and we have verified our equilibrium bidding strategy. □

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:*

$$b_i(\theta_i) = \frac{N-1}{N-1+P^c}\theta_i \tag{6}$$

PROOF. Plugging $F(\theta_i) = \theta_i$ into Equation 5 (repeated as 12), we get:

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

□

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:*

$$b_i(\theta_i) = \theta_i - \frac{\int_0^{\theta_i} \Big(P^a \cdot F(b_i(x)) + (1-P^a) \cdot F(x) \Big)^{(N-1)} dx}{\Big(P^a \cdot F(b_i(\theta_i)) + (1-P^a) \cdot F(\theta_i) \Big)^{(N-1)}} \tag{8}$$

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 (7), 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 and set it to zero:

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

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:

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

Rearranging terms yields:

$$b_i(\theta_i) = \theta_i - \frac{\Big(P^a \cdot F(b_i(\theta_i)) + (1-P^a) \cdot F(\theta_i) \Big) \cdot b_i'(\theta_i)}{(N-1) \cdot \Big(P^a \cdot f(b_i(\theta_i)) \cdot b_i'(\theta_i) + (1-P^a) \cdot f(\theta_i) \Big)} \tag{13}$$

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 $b_i'(\theta_i)$.

$$b_i(\theta_i) = \theta_i - \frac{\int_0^{\theta_i} \Big(P^a \cdot F(b_i(x)) + (1-P^a) \cdot F(x) \Big)^{(N-1)} dx}{\Big(P^a \cdot F(b_i(\theta_i)) + (1-P^a) \cdot F(\theta_i) \Big)^{(N-1)}} \tag{14}$$

To verify Equation 14, first take its derivative:

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

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 13 yields Equation 14, 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 (8, and repeated as 14) 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 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.

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

$$= \theta_i - \frac{\int_0^{\theta_i} x^{(N-1)}\, dx}{\theta_i^{(N-1)}} \quad = \theta_i - \frac{\frac{1}{N}\theta_i^N}{\theta_i^{(N-1)}} \quad = \frac{N-1}{N}\theta_i$$

□

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 $b_i(\theta_i) = \frac{N-1}{N}\theta_i$.*

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 $F(\Phi_j(b_i(\theta_i))) = \frac{N}{N-1+\alpha_j}b_i(\theta_i)$. Note, however, that since $\Phi_j(\cdot)$ is only defined over the range $[0, b_j(1)]$, it must be the case that $b_i(1) \leq b_j(1)$, which is why $\alpha_j \geq 0$ is necessary, in addition to being sufficient. Assuming that $b_i(\theta_i) = \frac{N-1}{N}\theta_i$, then indeed $\Phi_j(b_i(\theta_i))$ is always well-defined. We will now show that this assumption is correct. The expected utility for agent i can then be written as:

$$E_{\theta_{-i}}\big[u_i(b(\theta), \theta_i)\big] = \Big[\Pi_{j\neq i}\frac{N}{N-1+\alpha_j}b_i(\theta_i)\Big] \cdot \Big[\theta_i - b_i(\theta)\Big]$$
$$= \Big[\Pi_{j\neq i}\frac{N}{N-1+\alpha_j}\Big] \cdot \Big[\theta_i \cdot \Big(b_i(\theta_i)\Big)^{(N-1)} - \Big(b_i(\theta_i)\Big)^N\Big]$$

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

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

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

Full Version of Example 1:

In a first-price auction where $F(\theta_i) = \frac{\theta_i^2+\theta_i}{2}$ and $N = 2$, if agent 2 always bids its type ($b_2(\theta_2) = \theta_2$), then, for all $\theta_1 > 0$, agent 1's best response bidding strategy is strictly less than the bidding strategy of the symmetric equilibrium.

After calculating the symmetric equilibrium in which both agents shave their bid by the same amount, we find the best response to an agent who instead does not shave its bid. We then show that this best response is strictly less than the equilibrium strategy. To find the symmetric equilibrium bidding strategy, we instantiate $N = 2$ in the general formula the equation found in [10], plug in $F(\theta_i) = \frac{\theta_i^2+\theta_i}{2}$, and simplify:

$$b_i(\theta_i) = \theta_i - \frac{\int_0^{\theta_i} F(x)dx}{F(\theta_i)}$$
$$= \theta_i - \frac{\frac{1}{2} \cdot \int_0^{\theta_i}(x^2 + x)dx}{\frac{1}{2} \cdot (\theta_i^2 + \theta_i)} \quad = \theta_i - \frac{\frac{1}{3}\theta_i^3 + \frac{1}{2}\theta_i^2}{\theta_i^2 + \theta_i} \quad = \frac{\frac{2}{3}\theta_i^2 + \frac{1}{2}\theta_i}{\theta_i + 1}$$

We now derive the best response for agent 1 to the strategy $b_2(\theta_2) = \theta_2$, denoting the best response strategy $b_1^*(\theta_1)$ to distinguish it from the symmetric case. The probability of agent 1 winning is $F(b_1^*(\theta_1))$, which is the probability that agent 2's type is less than agent 1's bid. Thus, agent 1's expected utility is:

$$E_{\theta_2}\big[u_i((b_1^*(\theta_1), b_2(\theta_2)), \theta_1)\big] = F(b_1^*(\theta_1)) \cdot \Big(\theta_1 - b_1^*(\theta_1)\Big)$$
$$= \frac{(b_1^*(\theta_1))^2 + b_1^*(\theta_1)}{2} \cdot \Big(\theta_1 - b_1^*(\theta_1)\Big)$$

Taking the derivative with respect to $b_1^*(\theta_1)$, setting it to zero, and then rearranging terms gives us:

$$0 = \frac{1}{2} \cdot \Big(2 \cdot b_1^*(\theta_1) \cdot \theta_1 - 3 \cdot (b_1^*(\theta_1))^2 + \theta_1 - 2 \cdot b_1^*(\theta_1)\Big)$$
$$0 = 3 \cdot (b_1^*(\theta_1))^2 + (2 - 2 \cdot \theta_1) \cdot b_1^*(\theta_1) - \theta_1$$

Of the two solutions of this equation, one always produces a negative bid. The other is:

$$b_1^*(\theta_1) = \frac{\theta_1 - 1 + \sqrt{\theta_1^2 + \theta_1 + 1}}{3}$$

We now need to show that $b_1(\theta_1) > b_1^*(\theta_1)$ holds for all $\theta_i > 0$. Substituting in for both terms, and then simplifying the inequality gives us:

$$\frac{\frac{2}{3}\theta_1^2 + \frac{1}{2}\theta_1}{\theta_1 + 1} > \frac{\theta_1 - 1 + \sqrt{\theta_1^2 + \theta_1 + 1}}{3}$$
$$\theta_1^2 + \frac{3}{2}\theta_1 + 1 > (\theta_1 + 1)\sqrt{\theta_1^2 + \theta_1 + 1}$$

Since $\theta_1 \geq 0$, we can square both sides of the inequality, which then allows us to verify the inequality for all $\theta_1 > 0$.

$$\theta_1^4 + 3\theta_1^3 + \frac{17}{4}\theta_1^2 + 3\theta_1 + 1 > \theta_1^4 + 3\theta_1^3 + 4\theta_1^2 + 3\theta_1 + 1$$
$$\frac{1}{4}\theta_1^2 > 0$$