

The Influence of Privacy Cost on Threshold Strategies in Sealed-Bid First and Second-Price Auctions

Sumit Joshi, Dept. of Economics
Yu-An Sun and Poorvi L. Vora, Dept. of Computer Science
The George Washington University
{ysun,poorvi}@gwu.edu

Abstract

This paper follows up on a model described elsewhere by the authors. The model incorporates a privacy cost component to expected payoff in auctions. This is the cost, to the bidder, of the information revealed, by the bid, about the bidder's valuation. It applies even if the sale is lost, and hence, negative pay-offs, not possible in regular auctions, are possible in auctions with privacy costs. An important concern is whether such a cost affects seller revenue. Our initial analysis of eBay data on sales of the Mustang indicates that those bidders who avail of privacy technology tend to pay a larger sale price, on average, than those who do not. This indicates that privacy cost does effect seller revenue in at least some cases.

In this paper we provide a mathematical analysis of *threshold strategies* - where bidders have monotonic increasing strategies for valuations $x \geq x_t \geq 0$ and do not bid for valuations $x < x_t$ - for sealed-bid first and second-price auctions with privacy cost. We observe that bidders pass on at least some, if not all, privacy costs to the seller, and that expected revenues are lower in auctions with privacy costs than otherwise. We also find that a *zero-plus* bid - one arbitrarily close to zero - is a valid equilibrium bid, and, in equilibrium, the probability of it winning the auction is non-negligible. Most importantly, we find that when privacy costs are large enough, revenue equivalence does not hold. In such cases, first-price auctions provide higher expected payoffs, but which auction provides a higher revenue depends on the privacy cost function and the number of bidders.

Because bidders can choose between the two auctions and revenues and payoffs are not equivalent, and because the bidder passes on at least some of the privacy cost to the seller, it is in the interest of the seller to reduce privacy cost. We conclude that it could often be valuable for the seller to provide privacy-protecting technology while executing an auction, so as to increase revenue by reducing the bidder's privacy cost.

1 Introduction

The sealed-bid second price auction, also known as the Vickrey auction [21], is one of the most fundamental mechanisms in game theory. Its importance arises from the existence of a dominant strategy of truth revelation. However, the Vickrey model does not include the privacy cost of making known a truth-revealing bid in dynamic games where the seller is likely to be re-encountered as a competitor or negotiator. An enhancement of the Vickrey model to include privacy costs would be very applicable to today's markets, where tracking, automated agents, and financial profiling

are very common, and cryptographic privacy protection rare. The model would prove useful in answering questions such as the following. Does the Vickrey auction have a dominant strategy in the presence of privacy costs? Is it truth revealing? Does privacy cost decrease seller revenue? Does revenue equivalence hold?

A paper by the authors (in review [9]) takes the view that a privacy cost could affect optimal strategies in games, particularly dynamic games. It models the cost as a monotonic increasing function of valuation that applies to the bidder's payoff whether the sale is won or not. Because negative payoffs are possible when the sale is lost, auctions with privacy costs are not identical to regular auctions with lower valuations. The paper addresses the specific question of whether the cost affects seller revenue when equilibrium strategies are restricted to being monotonic increasing in first and second-price sealed-bid auctions. (The restriction of a monotonic increasing strategy provides upper bounds on the privacy cost and its derivative). The paper finds that, in such cases, the second-price auction no longer possesses a dominant strategy. On examination of the symmetric Nash equilibrium for both auctions, it finds that privacy costs reduce the expected revenue of the seller, but do not affect the expected payoff of the bidder; i.e. the bidder passes on the entire privacy cost to the seller. It also finds that revenue equivalence continues to hold. While these results are interesting, one is immediately led to ask how a relaxation of the requirement of a monotonic increasing strategy would affect the results. This is especially relevant because all indications are that a wide variety of strategies - including decreasing ones, and those that require zero bids for non-zero valuations - are possible.

Building on [9], this paper examines the first and second-price sealed-bid auctions while restricting the strategies to be *threshold strategies*, i.e. strategies that are zero for valuation $x < x_t$ and monotonic increasing for $x \geq x_t$ for some $x_t \geq 0$. Monotonic increasing strategies are a special case of threshold strategies, with $x_t = 0$. The requirement of a threshold strategy restricts the auction to being efficient, and bounds the privacy cost and its derivative from above, but allows higher values than the monotonic increasing requirement of [9]. The results are very interesting. Revenue equivalence does not always hold, the bidder does not always pass on all privacy cost to the seller, first-price auctions always provide a larger expected payoff to the bidder, and bidding strategies with infinitesimally small bids have a finite probability of winning the sale. Initial analysis of auction data from eBay indicates that the decrease in seller revenue due to privacy cost could be as large as 5-10% of sale price.

1.1 A Privacy Cost to Bidding

Knowledge of an individual's valuation of an item enables price discrimination, which could be useful if it allowed a fairer distribution of items and enabled the estimation of demand. In a stable

market equilibrium, it can be argued that vendors would be motivated to reduce consumer privacy in order to improve the accuracy of price discrimination, and customers who valued their privacy would avoid situations that enabled the estimation of their valuations. This would result in a balance that reflected the value placed on personal information by both customers and vendors. In what direction is the market for personal information progressing? What kind of financial value do individuals and data collectors place on personal information? We argue that the answers to these questions are not as simple as might appear at face value.

Experimental evidence cited in [1] describes how even those customers who profess to value their privacy do not assert its value in interactions where doing so would inconvenience them or cost money. Further, many consumers are willing to trade their privacy in exchange for something in return [20] - for example, many trade their grocery shopping profile for a small discount. On the other hand, some consumers are willing to pay for privacy - those who pay to be kept off telephone directories are examples. It appears that when the decisions are simple and not to be made on a continual basis, customers might choose to assert, economically, their value for privacy - whether it is to pay for privacy or to give it up for a convenience. However, given the extensive and continual nature of data collection today, it is currently not possible for even the most privacy-conscious individual to make the innumerable decisions required, manually and on a case-by-case basis, about whom to reveal information to and how much to reveal.

It is quite possible that customers would behave differently if the cost of asserting their privacy were not as high as it currently is - for example if automated tools, such as privacy agents, existed that could participate in transactions on their behalf. One reason such tools do not exist is that optimal strategies in games do not typically take into account the privacy cost of participating in the game, and even a theoretical examination of the effect of privacy costs on the outcomes of simple games does not exist. Cryptographic auction protocols, with or without a trusted third party, have been proposed to provide privacy in auctions [15]; however, most cryptographic solutions are computationally expensive and not practical. Because of this, and because it is thought that privacy is not valued very highly, cryptographic auction design is not widely used. For example, the FCC has deployed a new automated auction system [12] with neither encryption (except a Secure ID card for authentication purposes), nor anonymous protocols. It is possible that a combination of negotiation and cryptographic tools - where negotiating tools enable customers to assert the financial value they place on privacy, and cryptographic tools enable vendors to provide the privacy required - would lead us to a situation more conducive to elucidating the real value customers place on privacy.

1.2 Our Approach and Results

As in [9], we approach sealed-bid first and second price auctions by including the bidder's privacy costs, and derive equilibrium bid strategies which we then compare with the classical ones [13]. The equilibrium bid strategy is obtained by maximizing the payoff of the bidder; as [9] shows that there is no dominant strategy in either auction with privacy costs, we examine the symmetric Bayesian Nash equilibria. Its necessary requirements indicate that strategies are not restricted much. For example, the bid may be positive at one valuation, zero at a higher one, and positive again at a still higher one. Some bidders with non-zero valuations may not bid. Clearly, the bid strategy is *not necessarily injective* with respect to valuation - many valuations correspond to zero bids, for example. Further, the general bid strategy is also *not efficient*, because the highest valuation does not always win the sale - in general, larger valuations do not always correspond to larger bids. In particular, some bidders with non-zero valuations may not bid at all, and bidders with lower valuations could win.

Because the strategies are not necessarily injective or efficient, the standard approaches for deriving symmetric equilibria do not hold, and the most general case of bidding with privacy costs in first and second-price sealed bid auctions is beyond the scope of this paper. We limit ourselves to threshold strategies. We observe that, when bidders are risk neutral, valuations independent and identically distributed, privacy costs monotonic increasing functions of estimated valuation, and equilibrium bid strategies threshold (i.e. there is a valuation $x_t \geq 0$ above which bid strategies are monotonic increasing, and below which bids are zero or infinitesimally small):

1. Privacy costs always result in lower bids, lower expected sale price, and revenue loss.
2. When $x_t \neq 0$:
 - (a) Infinitesimally small bids are possible, have a non-infinitesimal probability of winning the auction, and correspond to a non-infinitesimal expected payoff.
 - (b) First-price auctions are better for bidders. The first-price bidder passes on all the privacy cost to the seller; her expected payoff is identical to that when there is no privacy cost. The second-price bidder takes on some of the privacy cost; her expected payoff is strictly smaller than that when there is no privacy cost.
 - (c) Revenue equivalence does not hold.
 - (d) Expected revenues are strictly smaller than in auctions without privacy costs, and strictly larger than those for $x_t = 0$.

Initial experimental results also indicate that privacy cost reduces seller revenue

1.3 Organization of the paper

The paper is organized as follows: related work is discussed in section 2. Section 3 deals with preliminaries such as notation and assumptions, summarizes the model and results of [9] and sketches the approach and techniques used. Results are presented in section 4. Section 5 provides conclusions and future work. The appendix contains a table of notation and proof details.

2 Related Work

Cryptographic protocols have been developed to provide many strong provable results regarding privacy protection in auctions. They tend to address the following security issues: secrecy of the bids [5, 18], bidder anonymity [6], non-repudiation [5] and correctness of the auction result [3, 5, 6, 14, 18]. Thus cryptography may be used to reduce the bidder's privacy cost in auctions; however the cryptographic techniques are computationally expensive or required a trusted third party, and are hence not yet practical for widespread use. Further, all cryptographic techniques depend on the difficulty of particular problems in computer science, and are vulnerable to attacks by computationally unbounded adversaries.

There has been much interest on where markets are headed with respect to the value of personal information and the ease with which it is obtainable [1, 2, 19, 20]. For example, [19] argues that because commercial interests are likely to use personal information for the purposes of price discrimination, they have little incentive to provide privacy protection. This is a strong argument for financial value being associated with personal information. [20] makes an argument for the consideration of privacy markets while determining privacy regulation - for example, he argues that unnecessarily strict regulation may prevent an individual from trading her privacy for convenience. [2] shows that price discrimination is profitable in the presence of marketplace competition.

[11, 16, 17] provide excellent auction theory surveys. The Vickrey is known to be a truth-revealing auction, with a weakly dominant strategy [21]. Though the first-price auction does not have a dominant strategy, its Nash equilibrium strategy is an invertible function of the valuation, and valuations can be accurately estimated from knowledge of bids. Repeated auctions are common in the current commerce environment and the revelation of the winning bid provides valuable information when there are sequential, repeated auctions for similar items [8]. It is possible to classify competitors (high or low valuation, for example) after the auction is over and the winning bid is publicly known. With knowledge of a competitor's classification, other players (bidders, negotiators, vendors) are able to adjust their strategies. Knowledge of unsuccessful auction bids may also be used to the bidder's benefit, for example, [7] provides a means of using it to estimate

latent demand.

In the existing scenario, it appears unreasonable to ignore privacy costs altogether and to assume that privacy costs do not affect a bidder’s strategy at all. We assume, as in [9], that privacy costs depend on the bidder’s valuation. There has been some literature on costs in auctions, such as entry cost in fixed interval [10] and identical transaction cost [4], neither of which is valuation dependent.

3 Preliminaries

This section contains a review of the model and results of [9], notation, and a summary of the techniques used in this paper.

3.1 The Model and Results of [9]

We now summarize the model and results of [9] to provide the context for the work described in this paper. Equilibrium bid strategies in the classical first and second-price sealed-bid auctions are well-known to be $E[x_2|x = x_1]$ and x respectively for valuation x , x_1 the highest valuation, and x_2 the second highest one [13]. As the strategies for both auctions are invertible functions of the valuations, the first-price auction reveals as much information about the valuation as does the second-price auction. In both cases, the valuation can be determined (exactly) from the bid, and there is a privacy cost to bidding. In this paper as in [9], this cost is assumed to be a function of the estimated valuation and denoted $\psi(x)$; further, it is assumed that higher valuations would be more worth protecting and would hence correspond to higher privacy costs.

In auctions without privacy costs, the payoff is the difference between x and the sale price, if the bid is won, and zero if it is lost. Thus the payoff from a won sale is:

$$Payoff^{regular}(b) = x - saleprice$$

and is always non-negative for symmetric equilibria. Further, non-zero valuations always result in non-zero bids. In auctions with privacy costs there is an additional (negative) component of the pay-off - the privacy cost. It exists whether the bidder wins or loses, and is zero only if she does not bid. This paper shows that this might result in zero bids for non-zero valuations.

The payoff resulting from a sale in an auction with privacy cost $\psi(x)$ is

$$Payoff^{f^\psi, win}(b) = Payoff^{regular}(b) - \psi(x)$$

Not winning results not in a zero payoff, but in a payoff of value $-\psi(x)$:

$$Payoff^{f^\psi, loss}(b) = -\psi(x)$$

The different payoff expression poses a distinct optimization problem for the bidder. In particular, should the privacy cost contribute to a reduction in the bid, and hence also to the chances of winning, and to seller revenue? [9] reports that, when privacy cost and its rate of increase is small enough to result in a monotonic increasing strategy, the bid is reduced enough to pass on the entire privacy cost to the seller and to maintain the bidder's expected payoff. In this paper we show that, when the strategy is restricted to be a threshold strategy - of which monotonic increasing is a special case - some of the privacy cost is absorbed by the bidder and some of it by the seller.

[9] shows that there is no dominant strategy possible in either auction, and examines the symmetric Bayesian Nash equilibrium for both auctions. Its analysis of the first-price auction is similar to that of first-price auctions without privacy costs, which do not possess a dominant strategy. However, its analysis of the second-price auction is quite distinct, as second-price auctions without privacy costs have a weakly dominant strategy. In the second-price auction, the possible payoffs can be $-\psi(x)$ if the sale is lost, or $x - b_2 - \psi(x)$ if the sale is won, where b_2 is the second-highest bid. Thus, depending on the value of $b_2 \leq b$, any value in $[x - b - \psi(x), x - \psi(x)]$ is possible if the sale is won. Contrast this with the first-price auction where there are only two possible payoffs for bid $b \neq 0$: $x - b - \psi(x)$ and $-\psi(x)$.

The expected payoff is the quantity that a bidder seeks to maximize for the Nash equilibrium strategy. In the first-price auction, the expected value is taken over the probabilities of winning and losing the bid. Its value for bid b and equilibrium strategy β is:

$$E[\Pi_I(x)] = \begin{cases} 0 & b = 0 \\ (x - b)Pr[b = b_1] - \psi(\beta^{-1}(b)) & else \end{cases} \quad (1)$$

In the second-price auction, the expected payoff, in addition to being the expected value over winning and losing as in the first-price auction, is also the expected value over all possible second-highest bids if the sale is won. Its value is hence:

$$E[\Pi_{II}(x)] = \begin{cases} 0 & b = 0 \\ (x - E[b_2|b = b_1])Pr[b = b_1] - \psi(\beta^{-1}(b)) & else \end{cases} \quad (2)$$

The equilibrium strategy in both cases is determined by (i) differentiating the expected payoff wrt b , setting to zero and solving for b (ii) determining that another strategy would not provide a higher expected payoff to an individual bidder when all others are using the strategy obtained in (i).

[9] shows that, when bidders are risk neutral, valuations independent and identically distributed, privacy costs monotonic increasing functions of estimated valuation, and equilibrium bid strategies monotonic increasing:

1. There is no dominant strategy in either auction

2. Privacy costs always influence symmetric Nash equilibrium strategies for both auctions, resulting in lower bids and lower expected revenue
3. The bidder passes on all the privacy cost to the seller, and her expected payoff remains as without privacy cost
4. The seller's loss of revenue is equal to N times the expected value of an individual bidder's privacy cost, where N is the number of bidders
5. In expected value of payoff and expected revenue, both first and second-price auctions with the same privacy costs are the same, i.e. revenue equivalence holds

In this paper - when strategies are not required to be monotonic increasing, but are required simple to be threshold (of which monotonic increasing is a special case) - we show that revenue equivalence need not hold, and that the bidder does not always pass on all the privacy cost to the seller.

3.2 Notation, Definitions and Assumptions

As in [9], as far as possible, we follow the notation of Krishna [13]. We denote the valuation by x , the bid by b , the optimal bidding function by β , the payoff by Π , the revenue by R , the expectation operator by $E[\cdot]$, and the number of bidders by N . A subscript on the valuation or bid indicates its order in a non-increasing sequence; so x_1 denotes the highest valuation, b_1 the highest bid, and so on. A subscript of I or II on functions, such as Π , R , β etc. denotes a general expression for the first and second-price auction respectively. A superscript denotes Nash equilibrium values. Thus, the Nash equilibrium strategy for first-price auctions is denoted β^I , and for second-price by β^{II} . On the other hand, $\beta_I(x)$ and $\beta_{II}(x)$ denote an arbitrary bidding strategy for first and second price auctions respectively. We introduce notation as needed, a table of all our notation may be found in the appendix.

Definition 1: A *zero-plus bid* is a bid arbitrarily close to zero but not zero.

Definition 2: A *threshold strategy* is a strategy $\beta(x)$ such that \exists threshold x_t , $1 > x_t \geq 0$, and $\beta(x)$ is monotonic increasing for $x \geq x_t$, and zero or zero-plus for $x \leq x_t$.

Assumption 1: The Strategy We assume that the strategies are threshold, symmetric and known. Thus the valuation x corresponding to a non-zero, non-zero-plus bid b can always be determined: $\hat{x} = \beta^{-1}(b) = x$.

Assumption 2: The Bidders We assume that bidders are risk-neutral and symmetric; their valuations are independent and identically distributed over $[0, \omega]$. We denote the cumulative distribution function on each valuation by F and the corresponding probability distribution function

by f . We denote by $G(x)$ the probability that a given valuation x is the highest in a set of $N - 1$ bidders, $G(x) = [F(x)]^{N-1}$; this is the probability distribution of the highest valuation among $N - 1$ bidders.

Assumption 3: Privacy Cost We assume that an estimate $\hat{x} = f(b)$ may be made of the valuation, from the bid. The privacy cost is a function of the estimated valuation, and denoted $\psi(\hat{x})$. We also assume that revealing a larger valuation has a larger privacy cost, i.e. $\psi(x)$ is monotonic increasing as a function of x :

$$\frac{d\psi}{dx} \geq 0$$

and that the privacy cost of a zero valuation being known with complete accuracy is zero, i.e.

$$\psi(0) = 0$$

We use an additional superscript to indicate an auction with non-zero privacy cost $\psi(x)$ - for example, $\beta^{I,\psi}$ is the equilibrium strategy in a first price auction with privacy cost $\psi(x)$.

3.3 Review of Results for Classical Auctions Without Privacy Costs

The first-price auction without privacy costs does not possess a dominant strategy; its symmetric Nash equilibrium bidding strategy is well known to be the expected value of the second-highest valuation, conditional to x being the highest one:

$$\beta^I(x) = E[x_2|x = x_1] = \frac{\int_0^x yG'(y)dy}{G(x)} \quad (3)$$

and is the one that maximizes:

$$E[\Pi_I(x)] = (x - b)Pr[b = b_1] = (x - b)G(x) \quad (4)$$

The second-price auction without privacy costs possesses a dominant strategy, it is:

$$\beta^{II}(x) = x \quad (5)$$

and maximizes:

$$\Pi_{II}(x) = \begin{cases} 0 & b \neq b_1 \\ (x - b_2) & else \end{cases} \quad (6)$$

The expected payoff is identical for both auctions (where the expectation is over the probability of a win for a given valuation x in both auctions, and, additionally, also over the value of the second-highest bid in the second-price auction):

$$E[\Pi^{I,II}(x)] = \int_0^x G(y)dy \quad (7)$$

4 Our Main Results

In this section we state formally our main results and explain them. Detailed proofs may be found in the appendix.

4.1 Symmetric Nash Equilibrium: First-Price Auctions with Non-zero Privacy Cost

In a first-price auction with privacy cost $\psi(x)$, the expected payoff for a bid b and valuation x when all others are bidding $\beta(x)$ is (1), repeated here for convenience:

$$E[\Pi_I(x)] = \begin{cases} 0 & b = 0 \\ (x - b)Pr[b = b_1] - \psi(\beta^{-1}(b)) & else \end{cases}$$

where $Pr[b = b_1] = G(\beta^{-1}(b))$. On differentiating (1) wrt the bid, we obtain the equilibrium strategy, and expected payoff (see Appendix for proof). Simply put, the equilibrium strategy is as follows:

1. $\beta_{I,\psi}(x)$ denotes the strategy that sets the derivative of the expected payoff to zero:

$$\beta_{I,\psi}(x) = \frac{\int_0^x yG'(y)dy}{G(x)} - \frac{\psi(x)}{G(x)} = \beta^I(x) - \frac{\psi(x)}{G(x)} \quad (8)$$

If $\beta_{I,\psi}(x)$ is positive, the bidder bids $\beta_{I,\psi}(x)$.

2. We do not consider those cases for which $\beta_{I,\psi}(x)$ is always positive but not one-to-one, such a strategy is not a threshold strategy. See Case a, Example 1, strategy $\beta^a(x)$ and Figure 1 for an example of such a strategy.
3. If $\beta_{I,\psi}(x) < 0$, a zero bid, and the corresponding zero payoff, is not the only possible bid.

It may be possible for the bidder to obtain an expected payoff that is non-zero, though lower than $E[\Pi_{I,\psi}(x)]$ (which is the payoff that requires the negative bid of $\beta_{I,\psi}(x)$). Suppose the bidder were to chose another (positive) bid, corresponding to, say, a valuation $z \neq x$. It can be shown that the expected payoff would be smaller than $E[\Pi_{I,\psi}(x)]$, but greater than 0 if (see Appendix):

$$x > \beta^I(z) \quad (9)$$

Among all such possible values of z , the bidder would choose the one offering the largest expected payoff. It can be shown that this value corresponds to a bid infinitesimally close to

zero, which we denote by $0+$. This bid is distinct from not bidding in the following way: a bidder who does not bid has no probability of winning, even if all other bidders do not bid. The bidder with a bid of $0+$ wins with a non-negligible probability, if and only if all other bids are 0 or $0+$, and pays a sale price of 0. This bid of value zero-plus has the same effect on seller revenue as does not bidding, but provides a larger expected payoff to the bidder, and can result in a sale.

Because we are addressing threshold bids, the value of z which would provide the largest payoff, i.e. the closest zero of $\beta_{I,\psi}(x)$, would be the valuation at which the strategy turns from zero to positive, x_t .

4. Hence the bidder bids $\beta_{I,\psi}(x)$ for values of $x \in (x_t, 1]$, $0+$ for values of $x \in [\beta^I(x_t), x_t]$ and 0 for $x \in [0, \beta^I(x_t))$

4.1.1 Example 1

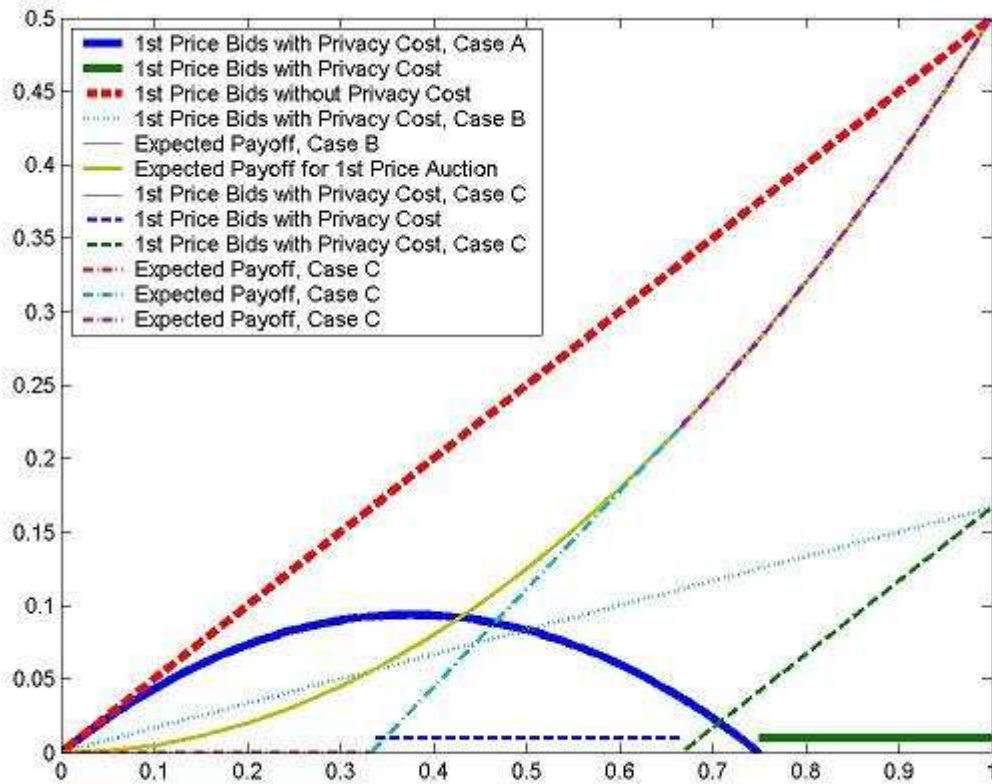


Figure 1: An illustration of strategies for first-price auctions with privacy costs

Figure 1 illustrates the results described above assuming $N = 2$, $F(x) = x$, i.e. the valuations are uniformly distributed over $[0, 1]$, and $\psi(x) = cx^k$. Hence $G(x) = x$, and, from (8), $\beta_{I,\psi}(x) =$

$$\frac{x}{2} - cx^{k-1}.$$

$$(a) k = 3, c = \frac{2}{3}$$

In this case, $\beta_{I,\psi}(x) = \frac{1}{2}x - \frac{2}{3}x^2 > 0 \Leftrightarrow x \in (0, \frac{3}{4})$, and the strategy is not threshold. Further, it is not injective and hence does not satisfy the assumptions under which it was derived. For example, consider the value $b = \frac{1}{12}$. It corresponds to two possible valuations, $x_1 = \frac{1}{4}$ and $x_2 = \frac{1}{2}$. The probability that another bid is smaller than $\frac{1}{12}$ is at least $Pr[0 \leq x \leq \frac{1}{4}] + Pr[\frac{1}{2} \leq x \leq \frac{3}{4}]$. The expression for $Pr[b = b_1]$ in the expected payoff, (4), which was assumed to derive $\beta_{I,\psi}(x)$, accounts for only one of the above intervals. Hence, β^a is not necessarily an equilibrium strategy. Further discussions about such strategies are outside the scope of this paper.

$$(b) k = 2, c = \frac{1}{3}$$

$\beta_{I,\psi}(x) = \frac{1}{2}x - \frac{1}{3}x = \frac{1}{6}x > 0 \quad \forall x \in (0, 1]$ and $x_t = 0$. The expected payoff is:

$$E[\Pi^b(x)] = \frac{x^2}{2} \quad \forall x$$

which is identical to the expected payoff for the regular first-price auction, (7).

$$(c) k = 1, c = \frac{1}{3}$$

$\beta_{I,\psi}(x) = \frac{1}{2}x - \frac{1}{3} > 0 \Leftrightarrow x \in (\frac{2}{3}, 1]$, and hence $x_t = \frac{2}{3}$. Hence,

$$\beta^c(x) = \begin{cases} \frac{1}{2}x - \frac{1}{3} & x \in (\frac{2}{3}, 1] \\ 0 & x \in (\frac{1}{3}, \frac{2}{3}] \\ 0 & x \in (0, \frac{1}{3}] \end{cases}$$

The expected payoff is (see (14)):

$$E[\Pi^c(x)] = \begin{cases} \frac{x^2}{2} & x \in (\frac{2}{3}, 1] \\ \frac{2}{3}(x - \frac{1}{3}) & x \in (\frac{1}{3}, \frac{2}{3}] \\ 0 & x \in (0, \frac{1}{3}] \end{cases}$$

The results described above may be stated more formally as in the next section .

4.1.2 Formal Statements of Results

Theorem 1: For iid valuation x distributed according to cumulative distribution function $F(x)$ over $[0, \omega]$, N risk neutral bidders, threshold bidding strategy, and monotonic increasing privacy

cost $\psi(x)$, if

$$\beta_{I,\psi}(x) = \beta^I(x) - \frac{\psi(x)}{G(x)}$$

and x_t is the only root of $\beta_{I,\psi}(x)$, then the symmetric equilibrium strategy is

$$\beta^{I,\psi}(x) = \begin{cases} \beta_{I,\psi}(x) & x > x_t \geq 0 \\ 0+ & \beta^I(x_t) \leq x \leq x_t \text{ and } x_t \neq 0 \\ 0 & \text{else} \end{cases} \quad (10)$$

and the corresponding expected payoff is

$$E[\Pi^{I,\psi}(x)] = \begin{cases} \int_0^x G(x)dx & x > x_t \\ \frac{(x - \beta^I(x_t))G(x_t)}{N(1 - \frac{F(\beta^I(x_t))}{F(x_t)})} & \beta^I(x_t) \leq x \leq x_t \text{ and } x_t \neq 0 \\ 0 & \text{else} \end{cases} \quad (11)$$

The proof is in the appendix. Notice that a single bid value, b , could correspond to, in general, many values of x , (roots of $\frac{\int_0^x yG'(y)dy - \psi(x)}{G(x)} = b$ - see (10)) and the assumption that $\beta(x)$ is injective, section 3.2, may not, in general, hold for all $\psi(x)$. We do not consider the most general case of a non-injective equilibrium strategy. Our assumption of a threshold strategy requires $\beta_{I,\psi}$ is monotonic increasing when it is greater than 0, i.e. that

$$\frac{\partial \beta_{I,\psi}}{\partial x} > 0$$

which implies, from (8), that

$$\begin{aligned} xG'(x) - \psi'(x) - G'(x)\beta^{I,\psi}(x) &> 0 \\ \Rightarrow G'(x)(x - \beta^{I,\psi}(x)) &> \psi'(x) \end{aligned} \quad (12)$$

i.e the rate at which the payoff increases is greater than the rate at which the privacy cost increases.

4.2 Symmetric Nash Equilibrium: Second-Price Auctions with Non-zero Privacy Cost

A second-price auction with privacy cost $\psi(x)$ does not possess a dominant strategy [9] and its expected payoff is (2), repeated here for convenience:

$$\Pi_{II}(x) = \begin{cases} 0 & b = 0 \\ (x - E[b_2|b = b_1])Pr[b = b_1] - \psi(\beta^{-1}(b)) & else \end{cases}$$

Differentiating, setting to zero and testing that the root is better than any other strategy gives us the result.

4.2.1 Formal Statements of Results

Theorem 2: For iid valuation x , distributed according to cumulative distribution function $F(x)$ over $[0, \omega]$, N risk neutral bidders, threshold symmetric equilibrium bidding strategy, and monotonic increasing privacy cost $\psi(x)$, if

$$\beta_{II,\psi}(x) = x - \frac{\psi'(x)}{G'(x)}$$

then

$$\beta^{II,\psi}(x) = \begin{cases} \beta_{II,\psi}(x) & x > x_t \geq 0 \\ 0 + & x_t > x > \frac{\psi(x_t)}{G(x_t)} \\ 0 & else \end{cases} \quad (13)$$

and

$$E[\Pi^{II,\psi}(x)] = \begin{cases} \int_{x_t}^x G(y)dy + x_t G(x_t) - \psi(x_t) & \beta_{II,\psi}(x) > 0 \\ \frac{xG(x_t) - \psi(x_t)}{1 - \frac{F(\psi(x_t))}{F(x_t)}} & x_t > x > \frac{\psi(x_t)}{G(x_t)} \\ 0 & else \end{cases} \quad (14)$$

As in the first price auction with privacy costs, there may be, in general, occasions when the bidder has a non-zero valuation but does not bid; similarly there may be occasions where $b = 0, 0+$ corresponds to more than one valuation, i.e. the strategy is not threshold. For a threshold strategy,

$$\frac{\partial}{\partial x} \left(\frac{\psi'(x)}{G'(x)} \right) < 1 \quad (15)$$

4.2.2 Example 2

Figure 2 illustrates the results described above on the same example used for Example 1 - $N = 2$, $F(x) = x$, and $\psi(x) = cx^k$. Hence $G(x) = x$, and $G'(x) = 1$, and (13) gives $\beta_{II,\psi}(x) = x - ckx^{k-1}$:

(a) $k = 3, c = \frac{2}{3}$

$\beta_{II,\psi}(x) = x - 2x^2 > 0 \Leftrightarrow x \in (0, \frac{1}{2})$, and hence is not a threshold bid.

(b) $k = 2, c = \frac{1}{3}$

$\beta_{II,\psi}(x) = x - \frac{2}{3}x = \frac{1}{3}x > 0 \quad \forall x \in (0, 1]$, hence monotonic increasing and a threshold strategy for $x_t = 0$. Further, (14) gives:




Figure 2: An illustration of strategies for second-price auctions with privacy costs

(c) $k = 1, c = \frac{1}{3}$

$\beta_{II,\psi}(x) = x - \frac{1}{3} > 0 \Leftrightarrow x \in (\frac{1}{3}, 1]$ is a threshold strategy for $x_t = \frac{1}{3}$. As $\frac{\psi(x_t)}{G(x_t)} = x_t$, other bids are zero:

$$\beta^c(x) = \begin{cases} x - \frac{1}{3} & x \in (\frac{1}{3}, 1] \\ 0 & \text{else} \end{cases}$$

The expected payoff is (see (14)):

$$E[\Pi^c(x)] = \begin{cases} \frac{x^2}{2} - \frac{1}{18} & x \in (\frac{1}{3}, 1] \\ 0 & \text{else} \end{cases}$$

4.3 Observations

Whether $x_t = 0$ or not, certain interesting results have been observed. These are also observed in [9]: bids are strictly smaller with privacy cost, hence so is the expected sale price, and the expected

revenue. When $x_t = 0$, [9] observes that revenue equivalence holds. When $x_t > 0$, we observe some very interesting results. This section states those results formally. Proof details are in the appendix.

Corollary 1: For a threshold strategy when $x_t > 0$, the following properties hold:

1. The expected payoff of the first-price auction, when $b \neq 0, 0+$, is identical to that when $x_t = 0$, and to that of the first-price auction without privacy cost.
2. The expected payoff of the second-price auction is strictly smaller than (a) that of the first-price auction, (b) that of the second-price auction when $x_t = 0$ and (c) that of the second-price auction without privacy cost.
3. The expected revenues for the two auctions are unequal. For both cases the revenues are strictly smaller than in the case with no privacy costs, and strictly larger than when $x_t = 0$.

Proof: 1 and 2 are obvious from Theorem 1 and the proof of Theorem 2 (Appendix) respectively. 3 is proved in the Appendix.

5 Comparison of the auctions

In this section we look at a single example and compare the two auctions when the strategies are threshold. We assume $F(x) = x$ and $\omega = 1$, i.e. valuations are uniformly distributed over $[0, 1]$, and $f(x) = 1$. Then $\frac{\int_0^x yG'(y)dy}{G(x)}$, the expected value of the highest valuation below x , is $\frac{N-1}{N}x$. We consider the example privacy cost function $\psi(x) = cx^k$ for $c \in (0, \frac{1}{N})$, $0 < k < N$. For context, Figure 3 shows this privacy function (labelled Case 2), compared to the function $\frac{N-1}{N}x^N$ (expected payoff in a classical auction), and to a privacy function that would result in a monotonic increasing strategy, $\psi(x) = cx^k$ for $c \in (0, \frac{N-1}{k(k-N+1)})$, $k \geq N$ (labelled Case 1).

The values of x_t differ in the two auctions with threshold strategies, for the same privacy cost. The bid values for $x > x_t$ are the same as in Case 1. Equilibrium bids for Case 2 are shown in Figure 4 for $\psi(x) = \frac{x^3}{10}$, $N = 5$.

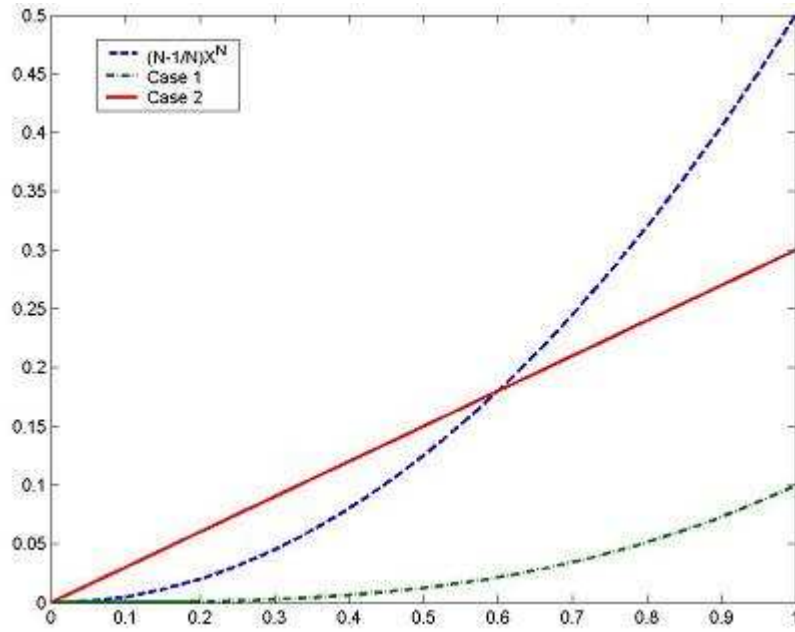


Figure 3: Simple possibilities for $\psi(x)$

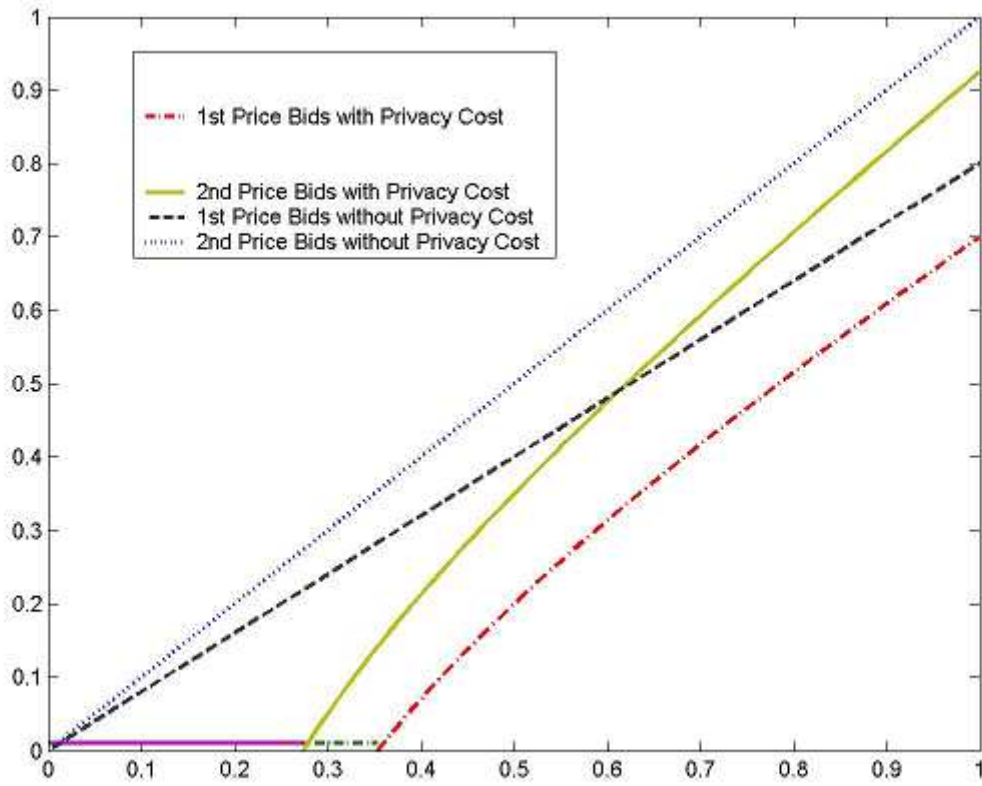


Figure 4: Case 2: Equilibrium Bids. $\psi(x) = cx^k$ for $c < \frac{1}{N}$ and $k < N$

6 Initial Experimental Results and Discussion

The website eBay.com contains a vast amount of public data on public auctions for all sorts of objects. The auctions conducted on eBay are all English auctions (increasing bid, open cry, first-price auctions). In addition, some of these auctions may be chosen to be *private*. A private auction is an auction where the bidder's email address and account name are hidden from the market. After the auction ends, only the seller is told who the buyer is. The choice for a private auction is made by the seller. eBay.com requests the seller not to use the private auction unless he has "a specific reason, such as potential embarrassment for bidders and the buyer" [22]. Because of the availability of sale prices with and without privacy, we made some initial examinations of the eBay data.

We selected the eBay-motors category for data collection because the "blue book" [23] would provide an independent estimate of market value of the cars sold. We used all the available data on successful auctions held from January 29th to February 28th 2005 for the Ford Mustang 2003 model. We chose this car because it is among the most popular used car models [24]. We analyzed those auctions that contained more than two bids during the auction period - this resulted in 7 private and 15 non-private auctions. For the private auctions, the average sale price is \$15,748.14, and the average blue book price is \$16,130.71 dollars. The difference is 1.8% of the sale price. For the non-private auctions, the average sale price is \$16,577.67 and the average blue book price \$18,389.67. The difference is 10.9% of the average sale price. We also observe that the average number of bids submitted per private auction is 23.85 and the average number of bids submitted per non-private auction is 22.06, which is not very different.

At first glance, it appears that bidders with privacy appear to pay more than those without. What is not clear, however, is whether there is a reason other than privacy cost for the two sale prices to be so different. Is a certain type of customer more likely to ask for a private auction, and is the higher sale price a reflection of some other property of that type of customer? Definitive answers require further work.

7 Conclusions and Future Directions

Our initial experimental observations indicate that auctions that provide privacy might result in higher revenue. Our theoretical results indicate that a non-zero privacy cost changes a number of auction properties. In particular, revenue equivalence does not hold, nor are the mechanisms efficient in general, and equilibrium strategies are not necessarily injective. Privacy costs typically decrease seller revenue, and often result in bidders not bidding even when their valuations are non-

zero. Because privacy costs result in different payoffs and expected revenues for the two auctions, it is often possible to determine whether one auction is better or worse than the other from the point of view of the bidder (based on expected payoff) and the seller (expected revenue). This fact, combined with the inefficiency and the lower revenues, could motivate sellers to provide privacy protection while holding auctions.

While we have determined that privacy costs result in several issues for both sellers and bidders, a number of problems remain unaddressed. For example, this paper does not analyze equilibrium strategies that are non-injective. It also does not address the issue of a single bidder with a privacy cost bidding in an equilibrium of bidders without privacy costs. Where it is not possible to come up with analytical results, simulations and computational optimization could provide additional direction. Finally, this paper does not address the possibility of designing other auction mechanisms that would naturally present lower privacy costs; for example, auctions in which it is difficult to estimate valuations from bids.

References

- [1] Acquisti, A. *Privacy in Electronic Commerce and the Economics of Immediate Gratification*. In Proceedings of the ACM Electronic Commerce (EC 04), New York, NY, ACM Press, 2004, 21-29.
- [2] Acquisti, A. and Varian, H. *Conditioning Prices on Purchase History*. Forthcoming, Marketing Science, 2005.
- [3] Cachin, C. *Efficient Private Bidding and Auctions with an Oblivious Third Party*. In Proceeding of 6th ACM Conference on Computer and Communications Security, 1999
- [4] Daniel, K. D. and Hirshleifer, D. A. *Theory of Costly Sequential Bidding*. University of Michigan Business School Working Paper, July 1998.
- [5] Franklin, M. K. and Reiter, M. K. *The Design and Implementation of a Secure Auction Service*. IEEE Transactions on Software Engineering, 1996, 302-312
- [6] Harkavy, M., Tygar, J. D. and Kikuchi, H. *Electronic Auctions with Private Bids*. In Proceeding of the Third USENIX Workshop on Electronic Commerce, 1998, 61-74
- [7] Huberman, B. A., Hogg, T. and Swami, A. *Using Unsuccessful Auction Bids to Identify Latent Demand*. In Proceedings of SMC, 2001 IEEE International Conference on System, Man and Cybernetics, Tucson, AZ, 2001, 2911-2916.

- [8] Jeitschko, T. D. *Learning in Sequential Auctions*. Southern Economic Journal, Vol. 65, No.1, July 1998, 98-112
- [9] Sumit Joshi, Yu-An Sun, Poorvi Vora, *Privacy Cost and Monotonic Increasing Strategies in Sealed-Bid First and Second-Price Auctions*. In review. At: <http://www.seas.gwu.edu/~poorvi/auctions.pdf>
- [10] Kaplan, T. and Sela, A. *Auctions with Private Entry Costs*. C.E.P.R Discussion Papers, 2003.
- [11] Klemperer, P. *Auction Theory: A Guide to the Literature*. Journal of Economic Surveys, 13(3), July 1999, 227-286.
- [12] Knowles, R. *Overview of Auction Process and FCC Automated Auction System*. FCC Auction Seminar, December 2003. Available at: <http://wireless.fcc.gov/auctions/55/resources/AuctionProcessFCCAutomatedAuctionSystem.pdf>
- [13] Krishna, V. *Auction Theory*. Academic Press. 2002.
- [14] Lipmaa, H., Asokan, N. and Niemi, V. *Secure Vickery Auctions without Threshold Trust*. In Annual Conference on Financial Cryptography, 2002, LNCS 2357. Springer, 87-101
- [15] Lopez, N., Nunez, M., Rodriguez, I. and Rubio, F. *Improving Privacy in Vickrey Auction*. ACM SIGEcom Exchanges Vol. 5.1, July 2004.
- [16] McAfee, R. P. and Mcillan, J. *Auctions and Bidding*. Journal of Economic Literature, June 1987, 25(2), 699-738
- [17] Milgrom, P. R. *Auction Theory*. Advanced in Economic Theory Fifth World Congress. Cambridge University Press, 1987, 1-32
- [18] Naor, M., Pinkas, B. and Sumner, R. *Privacy Preserving Auctions and Mechanism Design*. In Proceeding of 1st ACM Conference on E-Commerce, ACM Press, 1999, 129-139
- [19] Odlyzko, A. M. *Privacy, Economics, and Price Discrimination on the Internet*. ICEC 2003: Fifth International Conference on Electronic Commerce, ACM Press, 2003, 355-366
- [20] Varian, H. *Economic Aspects of Personal Privacy*. In U.S. Department of Commerce, Privacy of Information Age, 1996.
- [21] Vickrey, W. *Counterspeculation, Auctions, and Competitive Sealed Tenders*. Journal of Finance, 16, 1961, 8-37
- [22] <http://pages.ebay.com/help/sell/private.html>

[23] <http://www.kbb.com>

[24] <http://pages.ebay.com/sellercentral/sbc/scmotors/whatshot.html>

A Notation

x : valuation

x_t : threshold valuation

x_1 : highest valuation

x_2 : second-highest valuation

b : bid

b_1 : highest bid

b_2 : second-highest bid

$\psi(x)$: privacy cost as a function of valuation

$\beta(\cdot)$: optimal bidding function

Π : the payoff

R : the revenue

$E[\cdot]$: the expectation operator

N : the number of bidders

$\beta^I(\cdot)$: Nash equilibrium strategy for first-price sealed-bid auction

$\beta^{II}(\cdot)$: Nash equilibrium strategy for second-price sealed-bid auction

$\beta^{I,\psi}(\cdot)$: Nash equilibrium strategy for first-price sealed-bid auction with privacy cost function $\psi(x)$

$\beta^{II,\psi}(\cdot)$: Nash equilibrium strategy for second-price sealed-bid auction with privacy cost function

$\psi(x)$

$\beta_{I,\psi}(\cdot)$: An arbitrary strategy for the first-price sealed-bid auction with privacy cost function $\psi(x)$

$\beta_{II,\psi}(\cdot)$: An arbitrary strategy for the second-price sealed-bid auction with privacy cost function

$\psi(x)$

$E[\Pi_I(x)]$: expected payoff in the first-price sealed-bid auction

$E[\Pi_{II}(x)]$: expected payoff in the second-price sealed-bid auction

$F(x)$: the probability that the valuation of a single bidder is smaller than or equal to x

ω : the maximum valuation

$f(x) = F'(x)$: the probability distribution function of the valuation of a single bidder

$G(x) = F(x)^{N-1}$: the probability that the highest valuation of $N - 1$ bidders is smaller than or equal to x

z : an arbitrary valuation

B Optimal Strategies for Auctions with Privacy Costs

Our assumption that β is a threshold strategy implies that $Pr[b = b_1] = G(\beta^{-1}(b))$ for $b \neq 0, 0+$, and $Pr[b = 0, 0+] = G(x_t)$ where $G(x)$ is the probability that all other valuations will be smaller than x .

B.1 First-Price Auctions

The expected pay-off (1) for a first-price auction with injective strategy $\beta(x)$ for $x \geq x_t$ is:

$$E[\Pi] = (x - b)G(\beta^{-1}(b)) - \psi(\beta^{-1}(b)) \quad x \geq x_t \quad (16)$$

Theorem 1, Proof:

To find the Nash equilibrium (threshold) strategy for a first-price auction with privacy cost, differentiating wrt b and setting to zero in (16) gives:

$$\frac{\partial E[\Pi]}{\partial b} = -G(x) + (x - b)\frac{G'(x)}{\beta'(x)} - \frac{\psi'(x)}{\beta'(x)} = 0$$

The above equation in turn gives:

$$G'(x)x - \psi'(x) = bG'(x) + G(x)\beta'(x) \quad (17)$$

Noticing that

$$bG'(x) + G(x)\beta'(x) = \frac{\partial G(x)\beta(x)}{\partial x}$$

When $x_t = 0$, integrating (17) wrt x gives

$$\int_0^x yG'(y)dy - \psi(x) = G(x)\beta(x)$$

and

$$\beta_{I,\psi}(x) = \frac{\int_0^x yG'(y)dy - \psi(x)}{G(x)}$$

Assuming x_t is the only zero of the above expression when it is a threshold strategy, i.e. that $\int_0^{x_t} yG'(y)dy - \psi(x_t)$, and integrating (17) wrt x gives:

$$\int_{x_t}^x yG'(y)dy - (\psi(x) - \psi(x_t)) = G(x)\beta(x) - G(x_t)\beta(x_t)$$

As $\beta(x_t) = 0$, and $\int_0^{x_t} yG'(y)dy - \psi(x_t)$,

$$\int_0^x yG'(y)dy - \psi(x) = G(x)\beta(x)$$

Hence,

$$\beta_{I,\psi}(x) = \frac{\int_0^x yG'(y)dy - \psi(x)}{G(x)}$$

Substituting the above equation in (16) gives a corresponding expected payoff:

$$E[\Pi_{I,\psi}] = xG(x) - \int_0^x yG'(y)dy = \int_0^x G(y)dy$$

using integration by parts.

The above equations hold, of course, only when $x \geq x_t$, and $\beta_{I,\psi}(x) > 0$. When $\beta_{I,\psi}(x) \leq 0$, the bidder could choose to not bid. The expected payoff corresponding to this choice is zero. However, there is a range of non-zero payoffs, from $\int_0^x G(y)dy$ to 0 that is not achieved. All of these are not necessarily inaccessible to the bidder and her only choice is not between an expected payoff of $\int_0^x G(y)dy$ and one of zero. If possible, the bidder can choose a small non-negative bid such that her payoff is positive, and her bid close to zero. This would maintain continuity of both bid and payoff. If this is not possible, she bids zero.

To determine if this is an equilibrium strategy, we assume all other bidders use the strategy $\beta_{I,\psi}(x)$, except when $\beta_{I,\psi}(x) \leq 0$, and examine the payoff when a single bidder uses a bid $\beta_{I,\psi}(z)$, $z \neq x$ when her valuation is x . The probability that she wins is the probability that a bidder with bid $\beta_{I,\psi}(z)$ wins, i.e. it is the probability that valuation z wins. Similarly, her privacy cost is that of someone with valuation z because the valuation estimated from her bid is $\beta_{I,\psi}^{-1}(\beta_{I,\psi}(z)) = z$.

Her payoff is hence:

$$\begin{aligned} E[\Pi_{I,\psi,z}(x)] &= xG(z) - \int_0^z yG'(y)dy \\ &= xG(z) - zG(z) + \int_0^z G(y)dy \end{aligned} \tag{18}$$

which is positive when $x > \frac{\int_0^z yG'(y)dy}{G(z)} = \beta^I(z)$. The difference between the payoff due to strategy $\beta_{I,\psi}$ and the above payoff is:

$$\begin{aligned} E[\Pi_{I,\psi}] - E[\Pi_{I,\psi,z}(x)] &= \int_0^x G(y)dy - xG(z) + zG(z) \\ &\quad - \int_0^z G(y)dy \\ &= \int_z^x G(y)dy - (x-z)G(z) = \int_z^x (G(y) - G(z))dy \end{aligned}$$

which is easily shown to be positive as $G(x)$ is monotonic increasing; hence $E[\Pi_{I,\psi,z}(x)] < E[\Pi_{I,\psi}]$. However, $E[\Pi_{I,\psi,z}(x)]$ is non-negative and better than a payoff of zero when $x > \beta^I(z) > 0$.

Further, the derivative of $E[\Pi_{I,\psi,z}(x)]$ wrt z is $(x-z)G'(z)$, hence it is increasing wrt z if $x > z$, and decreasing if $x < z$. The best value of z is hence the one that is closest to x . As we assume a threshold strategy, the value of z closest to x for which $\beta_{I,\psi}(z)$ is non-negative is $z = x_t$.

Thus, when $\beta_{I,\psi}(x) < 0$, an optimal payoff is obtained when the bidder uses the strategy approaching that of $\beta_{I,\psi}(x_t)$, when $x > \beta^I(x_t)$, and bids zero otherwise. The expected payoff for a zero-plus bid is (18):

$$E[\Pi_{I,\psi}(x)] = \frac{xG(x_t) - \int_0^{x_t} yG'(y)dy}{\frac{F(x_t) - F(\beta^I(x_t))}{F(x_t)}N} = \frac{xG(x_t) - \psi(x_t)}{N(1 - \frac{F(\beta^I(x_t))}{F(x_t)})} \quad (19)$$

which is calculated by sharing the payoff of winning because all other bids are 0 or 0+ (probability $G(x_t)$) with all other possible winners, i.e. all other zero-plus bidders. The average number of zero-plus bidders is $N(1 - \frac{F(\beta^I(x_t))}{F(x_t)})$ when the highest bid is 0+.

B.2 Second-Price Auctions

Theorem 2, Proof:

For the second-price auction with privacy, the payoff is, from (6):

$$\Pi_{II}(x) = \begin{cases} 0 & \beta(x) = 0 \\ x - b_2 - \psi(\beta^{-1}(b)) & \beta(x) = b_1 \\ -\psi(\beta^{-1}(b)) & else \end{cases}$$

The expected value of the payoff takes into consideration not just the probability of winning, but also the expected value of the second-highest bid conditional on a win, and is (2).

$$\Pi_{II}(x) = \begin{cases} 0 & b = 0 \\ (x - E[b_2|b = b_1])Pr[b = b_1] - \psi(\beta^{-1}(b)) & else \end{cases}$$

Assuming a bid b for a valuation x when others bid equilibrium strategy β , and when β is a threshold strategy:

$$\begin{aligned} E[\Pi_{II}(x)] &= (x - \frac{\int_{x_t}^{\beta^{-1}(b)} \beta(y)G'(y)dy}{G(\beta^{-1}(b))})G(\beta^{-1}(b)) - \psi(\beta^{-1}(b)) \\ &= xG(\beta^{-1}(b)) - \int_{x_t}^{\beta^{-1}(b)} \beta(y)G'(y)dy - \psi(\beta^{-1}(b)) \end{aligned}$$

Differentiating wrt b and setting to zero, with $b = \beta(x)$ at equilibrium:

$$\frac{\partial E[\Pi_{II}(x)]}{\partial b} = \frac{xG'(x)}{\beta'(x)} - \frac{\beta(x)G'(x)}{\beta'(x)} - \frac{\psi'(x)}{\beta'(x)} = 0$$

The solution is:

$$\beta_{II,\psi}(x) = x - \frac{\psi'(x)}{G'(x)}$$

Assume that x_t is the only zero of the above function, and that $\beta_{II,\psi}(x) < 0$ for $x < x_t$. The expected payoff is:

$$\begin{aligned} E[\Pi_{II,\psi}(x)] &= \left(x - \frac{\int_{x_t}^x yG'(y)dy - \int_{x_t}^x \psi'(y)dy}{G(x)}\right)G(x) - \psi(x) \\ &= xG(x) - \int_{x_t}^x yG'(y)dy + \psi(x) - \psi(x_t) - \psi(x) \\ &= xG(x) - \int_0^x yG'(y)dy + \int_0^{x_t} yG'(y)dy - \psi(x_t) < \int_0^x G(x)dx \end{aligned}$$

where the inequality in the expression above is obtained on observing that $\psi'(x) > G'(x)x \forall x < x_t$, which implies that $\int_0^{x_t} \psi'(x)dx = \psi(x_t) > \int_0^{x_t} xG'(x)dx$. Further, on integrating the above expression by parts,

$$E[\Pi_{II,\psi}(x)] = \int_0^x G(y)dy + \int_0^{x_t} yG'(y)dy - \psi(x_t) = \int_{x_t}^x G(y)dy + x_tG(x_t) - \psi(x_t)$$

When $x - \frac{\psi'(x)}{G'(x)} < 0$, as in the first-price auction, the bidder may bid a non-zero value. Suppose she bids $\beta_{II,\psi}(z)$. Her payoff is:

$$\begin{aligned} E[\Pi_{II,\psi,z}(x)] &= \left(x - \frac{\int_{x_t}^z yG'(y)dy - \int_{x_t}^z \psi'(y)dy}{G(z)}\right)G(z) - \psi(z) \\ &= xG(z) - \int_{x_t}^z yG'(y)dy - \psi(x_t) \\ &= (x - z)G(z) + x_tG(x_t) + \int_{x_t}^z G(y)dy - \psi(x_t) \end{aligned}$$

Subtracting this from the payoff $E[\Pi_{II,\psi}(x)]$ gives:

$$E[\Pi_{II,\psi}(x)] - E[\Pi_{II,\psi,z}(x)] = \int_z^x G(y)dy - (x - z)G(z)$$

is greater than zero, hence, when greater than 0, $\beta_{II,\psi}$ is optimal. Further, when $x < x_t$, if

$$x > \frac{\int_{x_t}^z yG'(y)dy + \psi(x_t)}{G(z)} \Rightarrow E[\Pi_{II,\psi,z}(x)] > 0$$

As in the first-price auction, the optimal such z is the one closest to x , hence it is the value x_t . The optimal expected payoff for a 0+ bid is hence:

$$= xG(x_t) - \psi(x_t)$$

distributed over the many possible 0+ bidders, as with the first-price auction:

$$\frac{xG(x_t) - \psi(x_t)}{1 - \frac{F(\psi(x_t))}{F(x_t)}}$$

C Observations

Proof, Corollary 1:

The probability that the largest valuation is smaller than x is $[F(x)]^N$, and the probability distribution function of the largest valuation is $N[F(x)]^{N-1}f(x) = NG(x)f(x)$. Lost revenue is due to two sources: no bids for valuations smaller than x_t , and smaller non-zero bids for the same valuations. The lost revenue for the first-price auction is hence:

$$\begin{aligned} E[R_L^{I,\psi}] &= N \int_{x_t}^{\omega} \frac{\psi(x)}{G(x)} G(x) f(x) dx \\ &\quad + N \int_0^{x_t} \frac{\int_0^x y G'(y) dy}{G(x)} G(x) f(x) dx \\ &= N \int_{x_t}^{\omega} \psi(x) f(x) dx + N \int_0^{x_t} \int_0^x y G'(y) dy f(x) dx \\ &< N \int_{x_t}^{\omega} \psi(x) f(x) dx + N \int_0^{x_t} \psi(x) f(x) dx \end{aligned}$$

because $\psi(x) > \int_0^x y G'(y) dy$ for $x < x_t$, and $\psi(x_t) > \int_0^{x_t} y G'(y) dy$.

$$= N \int_0^{\omega} \psi(x) f(x) dx = NE[\psi(x)]$$

which is the lost revenue when $x_t = 0$. Similarly, the lost revenue for the second-price auction is:

$$\begin{aligned} E[R_L^{II,\psi}] &= N \int_{x_t}^{\omega} \frac{\int_{x_t}^x \frac{\psi'(y)}{G'(y)} G'(y) dy}{G(x)} G(x) f(x) dx \\ &\quad + N \int_0^{x_t} \frac{\int_0^x y G'(y) dy}{G(x)} G(x) f(x) dx \\ &= N \int_{x_t}^{\omega} [\psi(x) - \psi(x_t)] f(x) dx \\ &\quad + N \int_0^{x_t} \int_0^x y G'(y) dy f(x) dx \end{aligned}$$

$$= N \int_{x_t}^{\omega} \psi(x) f(x) dx + N \int_0^{x_t} \int_0^x y G'(y) dy f(x) dx - N \psi(x_t) (1 - F(x_t))$$

Now, $\psi'(x) > xG'(x)$ for $x \geq x_t$, hence the first two terms together are strictly smaller than $NE[\psi(x)]$, and

$$E[R_L^{II,\psi}] < NE[\psi(x)]$$

We do not compare $E[R_L^{I,\psi,2}]$ and $E[R_L^{II,\psi,2}]$ in general because the values of x_t for the two auctions are distinct.