

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

Sumit Joshi¹, Yu-An Sun², and Poorvi L. Vora²

¹ Dept. of Economics, George Washington University, Washington DC 20052
sumjos@gwu.edu

² Dept. of Computer Science, George Washington University, Washington DC 20052
ysun,poorvi@gwu.edu

Abstract. This paper approaches first and second-price sealed-bid auctions from the point of view that bidding reveals one's valuation, and hence comes with a privacy cost which depends on the valuation estimate. The privacy cost applies even if the sale is lost, and negative pay-offs, not possible in regular auctions, are possible. Hence auctions with privacy cost are not equivalent to regular auctions with lower valuations. This paper shows that second-price sealed-bid auctions with privacy costs do not have a dominant strategy. Further, it shows that privacy costs lower expected revenues in both first and second-price auctions, and that, when privacy costs are low enough, bidders pass on *all* privacy costs to the seller. It concludes that the seller might often benefit by providing 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 [22], is one of the most fundamental mechanisms in game theory. Its importance arises from two properties: (a) unlike most auctions and other economic games, it possesses a *dominant strategy* - i.e. a bidder's optimal strategy is independent of the actions of other bidders, and (b) this strategy is one of *truth revelation* - i.e. the optimal bid is the maximum a bidder would pay. Further, another fundamental result in game theory is that of *revenue equivalence* - many different auctions, including first and second-price ones, provide the same expected revenue to the seller.

An enhancement of the Vickrey model to include privacy costs would be very applicable to today's markets. Tracking and financial profiling are very common, and a seller may be re-encountered in another game as a competitor. Hence, revealing one's valuation through one's bid comes with a privacy cost. Cryptography may be used to reduce the bidder's privacy cost in auctions; cryptographic protocols have been developed to provide many strong provable results regarding privacy protection in auctions. However these are computationally expensive or

required a trusted third party, and are hence not yet practical for widespread use. The Vickrey model with privacy costs would hence prove useful in answering a number of interesting questions. 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? It would also enhance the existing computer science literature on quantitative privacy models [20, 5, 23], which, to our knowledge, have not been used to address privacy in auctions.

This paper shows that neither the first nor second-price sealed-bid auctions have a dominant strategy when privacy costs are taken into account. It derives symmetric Nash equilibria for both auctions and shows further that neither auction is truth-revealing and that privacy costs result in revenue loss. It also shows that if privacy costs are low enough, they do not affect revenue equivalence.

1.1 Information Revelation Comes With a Privacy Cost

Knowledge of an individual's valuation of an item enables price discrimination, the practice of charging different amounts for the same item to different customers. This could be useful if it allowed a fairer distribution of items and enabled an efficient estimation of demand - for example, knowledge of unsuccessful auction bids can be used to estimate latent demand [10]. In a stable market equilibrium, vendors would be motivated to reduce consumer privacy in order to improve the accuracy of price discrimination [2, 19], 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.

It might appear that customers do not value their privacy enough to tip the scales in favor of privacy. For example, 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. Yet, this is not the whole story. Certain consumers are willing to trade their privacy in exchange for something in return, for example, many trade their grocery shopping profile for a small discount [21], while others do not. 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. 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. 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 with a more complete picture of the real value customers place on privacy.

1.2 The Model

This paper takes the view that a privacy cost associated with revealing a bid in an auction, could affect the optimal strategy. In the model, a bidder's privacy costs are included in the payoff. Equilibrium bid strategies are derived and compared with classical ones for first and second-price sealed-bid auctions [15]. If $Payoff^{regular}$ is the payoff in an auction without privacy costs, the payoff in a similar auction with privacy cost ψ is $Payoff^\psi = Payoff^{regular} - \psi$. The additional (negative) component in the pay-off, the privacy cost, exists whether the bidder wins or loses, and is zero only if she does not bid.

If a bidder bids as she would have in a similar auction without privacy cost, the privacy cost will be completely absorbed in her expected payoff, and entirely borne by her. The seller will see no change in expected revenue. On the other hand, if the bidder is able to decrease the bid so as to maintain the value of her expected payoff, the seller will bear the entire cost and see a decrease in the expected revenue. Thus the problem of determining the equilibrium strategy for the bidder is one of determining how to distribute the privacy cost, i.e. whether privacy cost should reduce the bid or the expected payoff or both.

1.3 The Results

This paper 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 (Theorem 1).
2. Privacy costs always influence symmetric Nash equilibrium strategies for both auctions (Theorems 2 and 3); they result in lower bids and lower expected revenue (Corollary 1).
3. The bidder passes on all the privacy cost to the seller, and her expected payoff remains as without privacy cost (Corollary 1). 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 (Corollary 1).
4. 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 (Corollary 1).

The authors have also observed other interesting results, such as inefficiency, the existence of zero-valued bids for non-zero valuations, and the absence of revenue equivalence, in ongoing work [12]. These are observed when strategies are not required to be monotonic increasing, and are beyond the scope of this paper.

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

2 Related Work

Cryptographic auction schemes tend to address the following security issues: secrecy of the bids [8, 18], bidder anonymity [9], non-repudiation [8] and correctness of the auction result [18, 8, 9, 16, 3]. They are either computationally expensive or require a trusted third party, and are hence not yet widely deployed. For example, the FCC has deployed a new automated auction system [14] with neither encryption (except a Secure ID card for authentication purposes), nor anonymous protocols. There is also other work common to cryptography and game theory, for example [7] contains a cryptographic solution to a game theory problem. The computer science literature contains some recent work on the quantitative modelling and measurement of privacy [5, 20, 23], but, as mentioned in section 1, we are not aware of the application of these ideas to auctions. There are also other interesting papers on where markets are headed with respect to the value of personal information and the ease with which it is obtainable [19, 21, 2].

Two common types of auctions [15] are relevant to our work: the first and second-price sealed bid auctions. In both auctions, each bidder submits his own bid without knowing any other bidder's offer, and the bidder who makes the highest offer wins the item. In the first-price auction, the bidder pays her own bid as the sale price; in the second-price one, she pays the second highest bid. It can be shown that both auctions are efficient - the highest valuation obtains the bid; and revenue equivalent. As mentioned earlier, the second-price sealed-bid auction, has a dominant strategy and is a truth-revealing auction [22]. 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 [11]. There has been work in the literature on efficiently learning functions, in particular on learning functions with minimum expense when determining information comes with a cost [4]; such algorithms can easily be adapted to efficiently - and economically - learning financial profiles of individuals on encountering them in many interactions, with costs associated with the interactions. For example, it is possible to classify competitors (high or low valuation, for example) after a series of auctions is over. With knowledge of a competitor's classification, other players (bidders, negotiators, vendors) are able to adjust their strategies. As described earlier, we assume that this results in a privacy cost associated with bidding, and that the

cost depends on the bidder's valuation. There has been some literature on costs in auctions, such as entry cost in fixed interval [13] and identical transaction cost [6], neither of which is valuation dependent.

3 Preliminaries

Equilibrium bid strategies in the first and second-price sealed bid auctions with assumed zero privacy cost are well-known to be $E[x_2|x = x_1]$ and x respectively for valuation x , highest valuation x_1 , and second highest valuation x_2 [15]. In both cases, the bid is an injective function of the valuation and the valuation can be determined (exactly) from the bid. We denote the resulting privacy cost by $\psi(x)$; it seems reasonable 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. The payoff is always non-negative, and non-zero valuations always result in non-zero bids, i.e. $x \neq 0 \Rightarrow b \neq 0$ [15]. The payoff from a won sale may be expressed as $Payoff f^{regular}(b) = x - \text{saleprice}$. On the other hand, the payoff resulting from a sale in an auction with privacy cost $\psi(x)$ is $Payoff f^\psi(b) = Payoff f^{regular}(b) - \psi(x)$. Not winning results not in a zero payoff, but in a payoff of value $-\psi(x)$. As dominant strategies do not exist, we seek the symmetric equilibrium strategy that maximizes the expressions for the payoff with privacy costs (i.e. we seek the Nash equilibrium). We focus only on the cases where $x \neq 0 \Rightarrow b \neq 0$; some other cases are being studied in [12].

3.1 Notation

As far as possible, we follow the notation of Krishna [15], and 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 $E[.]$, 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 symmetric Nash equilibrium values. Thus, the symmetric 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.

3.2 Assumptions

We make the following assumptions:

1. $\beta(x)$ is monotonic increasing³, the same for everyone, and known, i.e. strategies are symmetrical. Thus x can always be determined from the bid b : $\hat{x} = \beta^{-1}(b) = x$, where \hat{x} is an estimate of x .
2. Bidders are risk neutral, i.e. they seek to maximize their payoff. Bidders are symmetric, i.e. 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, $G(x) = [F(x)]^{N-1}$.
3. The privacy cost is denoted $\psi(\hat{x})$, and depends only on the estimated valuation.
4. Revealing a larger valuation has a larger privacy cost, i.e. $\psi(x)$ is monotonic increasing as a function of x : $\psi'(x) > 0$. 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 $\psi(x)$.

4 Our Results

This section contains most of our mathematical analysis with formal statements and proof sketches of all our results; more detailed proofs are in the Appendix and may not appear in the final conference proceedings version if so dictated by space constraints.

4.1 Dominant Strategies do not Exist

Dominant strategies do not exist in first-price auctions without privacy costs. Therefore, as we would expect, they also do not exist in first-price auctions with privacy costs. The interesting result in this section is that they also do not exist in second-price auctions with privacy costs.

Theorem 1: Dominant strategies do not exist in first or second-price auctions with non-zero privacy costs, unless, trivially, if $\psi(x) > x \forall x$, for which the dominant strategy is $\beta(x) = 0 \forall x$.

Proof: Suppose $\psi(x) > x \forall x$. Then the payoff of a win, $x - \text{saleprice} - \psi(x)$, is negative independent of the sale price. So is the payoff of a loss, $-\psi(x)$. Hence the zero payoff of not bidding is preferred, and $\beta(x) = 0$ is the optimal strategy independent of other bids, i.e. it is a dominant strategy.

Now suppose $\psi(x) < x$ for some x . Consider any strategy $\beta(x) \neq 0$. If the highest of the other bids is b_1 , then, if $b_1 > \beta(x)$, the sale is lost and the payoff

³ We find that a number of strategies are not monotonic increasing, these are beyond the scope of this paper; we address some in [12].

is $-\psi(x) < 0$. Hence, if $b_1 > \beta(x)$, a zero bid is preferable to $\beta(x)$. However, if $0 < b_1 < x - \psi(x)$, and $b_1 < \beta(x)$, the payoff would be larger than zero, and $\beta(x)$ is preferred to the zero bid. Hence there is no strategy that is optimal independent of other bids, and a dominant strategy does not exist in either auction.

In this paper we study the Nash equilibrium strategies of first and second-price auctions with privacy costs when the strategies are monotonic increasing. The equilibrium strategy in both cases is determined by (i) differentiating the expected payoff wrt b , setting to zero and solving for b , and (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).

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

For the first-price auction without privacy costs, it is well known that the symmetric Nash equilibrium strategy that maximizes the expected payoff:

$$E[\Pi_I(x)] = (x - b)G(x) \quad (1)$$

is the expected value of the second-highest valuation, conditional to x being the highest one:

$$\beta^I(x) = E[x_2|x = x_1] = \frac{\int_0^x yG'(y)dy}{G(x)} \quad (2)$$

where $G(x) = Pr[x = x_1] = Pr[b = b_1]$ because $\beta(x) = b$ is assumed monotonic increasing. Substituting (2) in (1) gives the maximum expected payoff, which happens to be identical to the maximum expected payoff of the second-price auction without privacy cost:

$$E[\Pi^{I,II}(x)] = \int_0^x G(y)dy \quad (3)$$

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:

$$E[\Pi_I(x)] = \begin{cases} 0 & b = 0 \\ (x - b)G(\beta^{-1}(b)) - \psi(\beta^{-1}(b)) & else \end{cases} \quad (4)$$

where $G(\beta^{-1}(b)) = Pr[b = b_1]$. We observe the following.

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

$$\beta^{I,\psi}(x) = \frac{\int_0^x yG'(y)dy - \psi(x)}{G(x)} = \beta^I(x) - \frac{\psi(x)}{G(x)} \quad x > 0 \quad (5)$$

VIII

$\beta(0) = 0$, and the corresponding expected payoff is:

$$E[\Pi^{I,\psi}(x)] = \int_0^x G(y)dy \quad (6)$$

Proof Sketch: Differentiating (4) wrt b , setting to zero, and substituting $b = \beta(x)$ at equilibrium, gives:

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

the right hand side of which is $\frac{\partial G(x)\beta(x)}{\partial x}$. Integrating the above wrt x and rearranging gives (5). Substituting the equilibrium bid (5) in (4) gives the corresponding expected payoff, (6), using integration by parts. It can easily be shown that this is an equilibrium strategy by using techniques such as those used in [15] for deriving the classical first-price equilibrium strategy⁴.

Remark: Notice that a single bid value, b , could correspond to, in general, many values of x , (roots of $\beta^I(x) - \frac{\psi(x)}{G(x)} = b$) and the assumption that $\beta(x)$ is monotonic increasing, section 3.2, may not, in general, hold for all $\psi(x)$. We assume, however, that privacy costs are small enough so that (5) is positive for positive x , and that privacy costs rise slowly enough so that (5) is monotonic increasing (the necessary condition is $xG'(x) > \psi'(x)$). Some cases of larger privacy costs are considered in ongoing work, [12].

4.3 Second-Price Auctions

The sealed-bid second-price auction without privacy costs possesses a dominant strategy

$$\beta^{II}(x) = x \quad (7)$$

which maximizes the payoff:

$$\Pi_{II}(x) = \begin{cases} 0 & b \neq b_1 \\ (x - b_2) & \text{else} \end{cases} \quad (8)$$

In the presence of privacy costs, there is no dominant strategy as shown in section 4.1. Hence we study the symmetric Nash equilibrium which maximizes the expected payoff. In the second-price auction with privacy cost, 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 the interval $[x - b - \psi(x), x - \psi(x)]$ is possible if the sale is won. Hence, the expected payoff, in addition to being the expected value over winning and

⁴ Consider the payoff corresponding to any other strategy, $\beta_{I,\psi}(x) = \beta^{I,\psi}(z)$, $z \neq x$. When all other bidders use strategy $\beta^{I,\psi}(x)$, it can be shown that $\beta^{I,\psi}(z)$ provides a lower payoff than $\beta^{I,\psi}(x)$, $\forall z \neq x$.

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:

$$E[\Pi_{II}(x)] = \begin{cases} 0 & b = 0 \\ (x - E[b_2|b = b_1])G(\beta^{-1}(b)) - \psi(\beta^{-1}(b)) & \text{else} \end{cases} \quad (9)$$

Differentiating (9), setting to zero and testing that the root is better than any other strategy gives us the following result:

Theorem 3: For iid valuation x , distributed according to cumulative distribution function $F(x)$ over $[0, \omega]$, N risk neutral bidders, and monotonic increasing privacy cost $\psi(x)$, if the equilibrium strategy is assumed monotonic increasing, it is:

$$\beta^{II, \psi}(x) = x - \frac{\psi'(x)}{G'(x)} \quad (10)$$

and the corresponding expected payoff is

$$E[\Pi^{II, \psi}(x)] = \int_0^x G(y)dy \quad (11)$$

Proof Sketch: Differentiating⁵ (9) wrt b and setting to zero, with $b = \beta(x)$ at equilibrium:

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

gives the solution, for which the expected payoff follows easily and can be shown to be optimal as in Theorem 1.

Remark: 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 a single bid corresponds to more than one valuation, i.e. Assumption 1 is not valid. However, we consider only those cases where the privacy costs and its derivatives are small enough so that the strategy is a monotonic increasing function.

4.4 An Example

An example in which the equilibrium bid is monotonic increasing is $\psi(x) = cx^k$ for $c \in (0, \frac{N-1}{k(k-N+1)})$, $F(x) = x$, and $k \geq N$. $F(x) = x$ corresponds to the uniform distribution on valuations, $f(x) = 1$, over interval $[0, 1]$, where $\frac{\int_0^x yG'(y)dy}{G(x)}$, the expected value of the highest valuation below x , is $\frac{N-1}{N}x$. Figure 1 illustrates the bids.

⁵ Observe that $E[b_2|b = b_1] = \frac{\int_0^{\beta^{-1}(b)} \beta(y)G'(y)dy}{G(\beta^{-1}(b))}$, and that, hence, $\frac{\partial E[b_2|b=b_1]G(\beta^{-1}(b))}{\partial b} = \frac{\beta(\beta^{-1}(b))G'(\beta^{-1}(b))}{\beta'(\beta^{-1}(b))} = \frac{\beta(x)G'(x)}{\beta'(x)}$

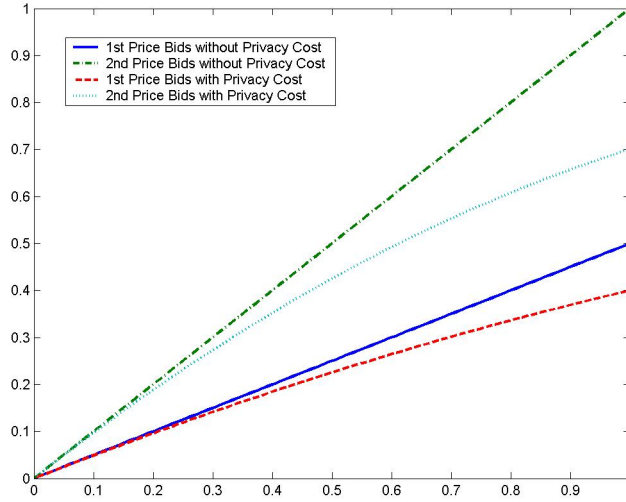


Fig. 1. Case 1: Equilibrium Bids. $\psi(x) = cx^k$ for $c < \frac{N-1}{k(k-N+1)}$ and $k > N$

4.5 Properties

The following properties follow immediately from Theorems 2 and 3 and comparisons of (7) with (10), (2) with (5) and (3) with (6) and (11).

Corollary 1: For iid valuation x , distributed according to cumulative distribution function $F(x)$ over $[0, \omega]$, N risk neutral bidders, and monotonic increasing privacy cost $\psi(x)$, for monotonic increasing first and second price equilibrium strategies, the following properties hold:

1. Both equilibrium bids are strictly smaller in the presence of privacy cost: $\beta^{I,\psi}(x) < \beta^I(x)$ and $\beta^{II,\psi}(x) < \beta^{II}(x)$.
2. The expected payoffs in both auctions are identical and equal to the expected payoffs of the first and second-price auctions without privacy costs.

$$E[\Pi^{I,\psi}] = E[\Pi^I] = E[\Pi^{II}] = E[\Pi^{II,\psi}]$$

3. The expected revenues in both auctions are identical - i.e. revenue equivalence holds - but strictly smaller in the presence of privacy cost. The revenue is:

$$E[R^\psi] = E[R] - NE[\psi(x)]$$

Proof Sketch: 1 and 2 are straightforward. The expected revenue in either auction is smaller by the amount lost because each bid is smaller. It can be

shown that the expected revenue loss in the second-price auction for a fixed highest valuation x_1 , where the expectation is taken over all possible values of the second-highest bid (i.e. the sale price), is identical to the revenue loss in the first-price auction for the same highest valuation x_1 . This proves revenue equivalence because the distribution of the valuations is identical in both auctions. The exact value of the revenue loss follows easily.

5 Conclusions and Future Directions

We have shown that dominant strategies do not exist in first and second-price sealed-bid auctions with privacy costs. Further, a non-zero privacy cost decreases seller revenue but does not affect expected payoff in symmetric Nash equilibria, when privacy costs are low and bid strategies monotonic increasing. This fact could motivate sellers to provide privacy protection while holding auctions in order to reduce privacy costs.

To our knowledge, this is the first paper to address the equilibria of auctions with privacy costs and a number of problems remain unaddressed. For example, this paper does not analyze equilibrium strategies that are not monotonic increasing. Further, it would be interesting to apply the model to open-cry auctions, where, arguably, privacy costs would be higher than in sealed-bid auctions. Further still, we have not attempted to analyze auction data to estimate privacy costs. For example, a seller's expected revenue depends on the privacy cost; the analysis of auction data could provide estimates of whether the privacy cost is high, low or negligible. Some other questions, such as those concerning strategies that are not monotonic increasing, would benefit from simulations and computational optimization when it is not possible to provide analytical results. 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. We hope to continue working on this problem and pursuing answers to these and other questions.

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. A. Acquisti, and H. Varian, *Conditioning Prices on Purchase History*. Forthcoming, Marketing Science, 2005.
3. C. Cachin, *Efficient Private Bidding and Auctions with an Oblivious Third Party*. In Proceeding of 6th ACM Conference on Computer and Communications Security, 1999
4. Moses Charikar, Ronald Fagin, Venkatesan Guruswami, Jon Kleinberg, Prabhakar Raghavan and Amit Sahai. *Query Strategies for Priced Information*. Proceedings of the 32nd Annual ACM Symposium on Theory of Computing (STOC), 2000. Journal version in Journal of Computer and System Sciences, Special Issue 64(4): 785-819, 2002.

5. Shuchi Chawla, Cynthia Dwork, Frank McSherry, Adam Smith and Hoeteck Wee. *Towards Privacy in Public Databases*. To appear in Theory of Cryptography (TCC) 2005, February 2005.
6. K. D. Daniel and D. A. Hirshleifer, *Theory of Costly Sequential Bidding*. University of Michigan Business School Working Paper, July 1998.
7. Yevgeniy Dodis, Shai Halevi, Tal Rabin. *A Cryptographic Solution to a Game Theoretic Problem*. CRYPTO 2000: 112-130
8. M. K. Franklin and M. K. Reiter, *The Design and Implementation of a Secure Auction Service*. IEEE Transactions on Software Engineering, 1996, 302-312
9. M. Harkavy, J. D. Tygar, and H. Kikuchi, *Electronic Auctions with Private Bids*. In Proceeding of the Third USENIX Workshop on Electronic Commerce, 1998, 61-74
10. B. A. Huberman, T. Hogg, and A. Swami, *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.
11. T. D. Jeitschko, *Learning in Sequential Auctions*. Southern Economic Journal, Vol. 65, No.1, July 1998, 98-112
12. Sumit Joshi, Yu-An Sun, Poorvi Vora, *Privacy Cost and Threshold Strategies in Sealed Bid First and Second-Price Auctions*. In preparation.
13. T. Kaplan and A. Sela, *Auctions with Private Entry Costs*. C.E.P.R Discussion Papers, 2003.
14. R. Knowles, *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>
15. V. Krishna, *Auction Theory*. Academic Press. 2002.
16. H. Lipmaa, N. Asokan, and V. Niemi, *Secure Vickrey Auctions without Threshold Trust*. In Annual Conference on Financial Cryptography, 2002, LNCS 2357. Springer, 87-101
17. N. Lopez, M. Nunez, I. Rodriguez and F. Rubio, *Improving Privacy in Vickrey Auction*. ACM SIGEcom Exchanges Vol. 5.1, July 2004.
18. M. Naor, B. Pinkas, and R. Sumner, *Privacy Preserving Auctions and Mechanism Design*. In Proceeding of 1st ACM Conference on E-Commerce, ACM Press, 1999, 129-139
19. A. M. Odlyzko, *Privacy, Economics, and Price Discrimination on the Internet*. ICEC 2003: Fifth International Conference on Electronic Commerce, ACM Press, 2003, 355-366
20. L. Sweeney. *k-anonymity: a model for protecting privacy*. International Journal on Uncertainty, Fuzziness and Knowledge-based Systems, 10 (5), 2002; 557-570.
21. H. Varian, *Economic Aspects of Personal Privacy*. In U.S. Department of Commerce, Privacy of Information Age, 1996.
22. W. Vickrey, *Counterspeculation, Auctions, and Competitive Sealed Tenders*. Journal of Finance, 16, 1961, 8-37
23. Poorvi Vora. An Information-Theoretic Approach to the Measurement of the Privacy of Random Data Perturbation (RDP). *In review*. Parts of the material have appeared in *The channel coding theorem and the security of binary randomization*. Proc., 2003 IEEE International Symposium of Information Theory, and *Information Theory and the Security of Binary Data Perturbation*. Proc., INDOCRYPT 2004.

A First-Price Auctions

The expected pay-off (4) for a first-price auction with monotonic increasing strategy $\beta(x)$ is:

$$E[II] = \begin{cases} 0 & \beta(x) = 0 \\ (x - b)G(\beta^{-1}(b)) - \psi(\beta^{-1}(b)) & \text{else} \end{cases} \quad (12)$$

Theorem 1, Proof:

To find the symmetric Nash equilibrium strategy for a first-price auction with privacy cost, differentiating wrt b and setting to zero in (12) gives:

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

if x is the highest valuation, and $\beta(x) = b$. The above equation in turn gives:

$$G'(x)x - \psi'(x) = bG'(x) + G(x)\beta'(x) \quad (13)$$

Noticing that

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

and integrating (13) wrt x if $b \neq 0$ when $x \neq 0$ gives

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

where constants of integration are zero because $\psi(0) = \beta(0) = 0$. Hence,

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

Substituting the above equation in (12) gives a corresponding expected payoff:

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

using integration by parts.

To determine if this is an equilibrium strategy, we assume all other bidders use the strategy $\beta_{I,\psi}(x)$ 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:

$$E[II_{I,\psi,z}(x)] = xG(z) - \int_0^z yG'(y)dy \quad (14)$$

XIV

$$= xG(z) - zG(z) + \int_0^z G(y)dy$$

The difference between the payoff due to strategy $\beta_{I,\psi}$ and the above payoff is:

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

which is easily shown to be positive as $G(x)$ is monotonic increasing; hence $E[\Pi_{I,\psi,z}(x)] < E[\Pi_{I,\psi}(x)]$.

B Second-Price Auctions

Theorem 2, Proof:

For the second-price auction with privacy, the payoff is, from (8):

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

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 (9):

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

Assuming a bid b for a valuation x when others bid equilibrium strategy β :

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

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:

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

for which the expected payoff is:

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

where the final equality is obtained on integrating by parts.

To determine that this is the equilibrium strategy, suppose the bidder bids another strategy, $\beta_{II,\psi}(z)$, when others are bidding $\beta_{II,\psi}(x)$. Her payoff is:

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

As in Theorem 1, this value is strictly smaller than the payoff $E[\Pi_{II,\psi}(x)]$.

C Properties

Proof, Corollary 1: The loss in revenue is due to the smaller bids. The probability that the largest valuation is smaller than x is $[F(x)]^N$, and the probability distribution function of the largest valuation is $N[F(x)]^{N-1}f(x) = NG(x)f(x)$. Hence the expected value of the lost revenue, $E[R_L^{I,\psi}]$, for the first-price auction is:

$$E[R_L^{I,\psi}] = \int_0^\omega \frac{\psi(x)}{G(x)} NG(x)f(x) dx = N \int_0^\omega \psi(x)f(x) dx = NE[\psi(x)]$$

The expected value of the loss of revenue corresponding to a highest valuation of x in the second-price auction is

$$E[R_L^{I,\psi}(x)] = \frac{\int_0^x \frac{\psi'(y)}{G'(y)} G'(y) dy}{G(x)} = \frac{\psi(x)}{G(x)}$$

where expectation is over all possible values of the second-highest bid, i.e. sale price. The expected revenue loss over all possible values of the highest valuation is as for the first-price auction.