

Price Discrimination and Privacy: A Note*

Sumit Joshi

Department of Economics, The George Washington University

Yu-An Sun

Xerox Corporation[†]

Poorvi L. Vora[‡]

Department of Computer Science, The George Washington University

Final Revision: June 2010

Abstract

In many instances of price discrimination, a seller of an item is in possession of signals from competing buyers regarding their private valuation for the item. If the seller uses this information to price discriminate against the buyer, buyers would correspondingly modify their signalling strategy. Our paper shows that the seller can gain by sometimes strategically ignoring the information contained in the signals and pricing the item in a non-discriminatory way. This “mixed” strategy induces buyers to send more informative signals in equilibrium than if the seller were to always price discriminate. Thus the seller can offset any revenue loss in states where he ignores information by the gains made in states where he can price discriminate more effectively due to the larger amount of information now communicated in the signals.

JEL Classification: D42, D82, L12

Key words and phrases: Price discrimination, noisy signaling, partition strategies, Bayesian-Nash equilibrium

*We would like to thank the Associate Editor and the referee for their comments and suggestions.

[†]This work was performed while the author was with Department of Computer Science, The George Washington University

[‡]Corresponding author. Department of Computer Science, Academic Center, 801 22nd Street NW, Washington D.C. 20052. Phone: (202) 994-1864; Fax: (202) 994-4875; Email: poorvi@gwu.edu

1 Introduction

Textbook models of first-degree price (or perfect) discrimination assume the seller has perfect information regarding buyers' preferences, and uses this information to make a take-it-or-leave-it price offer (for example, Varian 1989, Section 2.2). However, as noted by Varian (Section 2.2.2), the perfect information assumption is not likely to hold in the real world. A seller will typically price an item based on some prior beliefs regarding the distribution of buyers' preferences. Between the textbook model based on perfect information, and pricing based on prior beliefs alone, lies an interesting class of examples that have not hitherto been analyzed in the literature. These instances of (first-degree) price discrimination are distinguished by the fact that a seller is in possession of signals from the buyers reflecting their private valuation. While the information is not perfect, it is also not as coarse as implied by some unconditional prior regarding the distribution of buyers' valuations. Some examples of this phenomenon include the following:

1. **The second-chance offer of eBay:** Sellers on the auction site eBay collect bids from potential buyers and then assign the item to the highest bidder. Additional units of the item may then be sold to the losing set of bidders at an estimate of their valuations (Salmon and Wilson 2006). Thus sellers on eBay use the information contained in the bids (signals) to price discriminate among losing bidders.
2. **Resale in an open cry auction:** The winner of an open cry auction can use information from losing bids to resell the item to a losing bidder at a discriminatory price (Krishna 2002, Section 4.4).
3. **Phantom signalling in second price auctions:** Price discrimination can take a more subtle form in second-price auctions as noted by Lucking-Reiley (2000). Since submitting a bid equal to private valuation is a weakly dominant strategy in second-price auctions, the bids are an accurate signal of each buyer's private valuation. The seller then price discriminates by interspersing a phantom second bidder with a fictitious valuation just below the winning bid. The winning bidder thus pays a price higher than the legitimate second-highest bid.

It seems intuitive that an optimizing seller should strategically exploit all available information. However, if buyers are cognizant that any signal they provide to the seller to determine a winner amongst them will actually be used to price discriminate to their detriment, then they will incorporate this information into their signalling behavior. Thus, in competitive situations where there is one item and many buyers, each buyer will confront the following trade-off: the dictates of competition will demand that any signal sent be an accurate

reflection of the actual valuation; on the other hand, the proclivity of the seller to price discriminate will demand that signals be sufficiently noisy so that a seller cannot capture all the surplus from the exchange. This tension between information revelation and privacy raises a number of interesting issues. How do buyers address the conflict between information revelation due to competition and the concern for privacy due to price discrimination? How much information do buyers actually transmit to sellers in equilibrium? Is it always optimal to price discriminate? Or, should the seller sometimes behave in a non-discriminatory manner by strategically ignoring information contained in the signals and committing to a fixed pricing rule?

In this paper, we show that the case where the seller uses all the information is not optimal, and that seller revenue is greater when he sometimes behaves in a non-discriminatory manner. The model is described in Section 2. The equilibrium of the price discrimination game is characterized in Section 3. Section 4 allows non-discriminatory behavior. Section 5 discusses the connection of our work to the literature on optimal auctions. Section 6 concludes.

2 The Model

There are N ex-ante symmetric risk-neutral buyers for an object being offered by a seller. Buyer i has a private valuation V_i for the object which is iid on the interval $[\underline{v}, \bar{v}]$ according to a distribution function F with associated density f . For a buyer with realized valuation v , we denote by $G(v) = F^{N-1}(v)$ the probability that all $N - 1$ other buyers will have a lower valuation. Buyer i sends a signal b_i to the seller. Let $\mathbf{b} = (b_1, b_2, \dots, b_N)$ denote the vector of signals and b_{-i} denote signals from buyers other than i . The seller assigns the object to the buyer with the highest signal (and randomly allocates to one of the highest signals in the event of a tie).

A discriminating seller will extract all information from the signals and choose a (take-it-or-leave-it) price to maximize expected revenue. Let $h(v_i|b_i)$ denote the conditional expectation of V_i given b_i . For winning bid g , the seller chooses a price $p = p^*(b)$ that maximizes:

$$R(p, b) = p \int_p^{\bar{v}} h(x|b) dx \quad (1)$$

We believe that we would not lose much generality under the restriction that the price depends only on the winning bid. This is because, for the discriminating seller, signals other than the winning signal do not provide information on the valuation of the winning buyer.

A buyer i with realized valuation v_i and signal b_i receives a utility of $v_i - p^*(\mathbf{b})$ from winning the object and zero otherwise. Therefore we can write the expected payoff of buyer

i as:

$$u(b_i, v_i) = W(b_i, b_{-i})(v_i - p^*(b_i)) \quad (2)$$

where $W(b_i, b_{-i})$ is the probability of i winning the object given (b_i, b_{-i}) .

All these facts are common knowledge among the players. We refer to this game as the *price discrimination* (PD) game. Given the ex-ante symmetry of the buyers, we restrict attention to symmetric Bayesian-Nash equilibria. Such an equilibrium consists of a signalling strategy for the buyers and a pricing rule for the seller which are mutual best responses taking into account the players' private information. We first note an important property of equilibrium signalling strategies.

Lemma 1 *Symmetric Bayesian-Nash equilibrium signalling strategies of the PD game cannot be invertible.*

Proof. Suppose $\alpha(v)$ is an invertible symmetric equilibrium strategy. For winning bid $b = b_i$, $h(v_i|b) = 1$ if $v_i = \alpha^{-1}(b)$ and 0 otherwise. Hence the seller's price from maximizing (1) is $p^*(b) = \alpha^{-1}(b) = v_i$, and, with strategy α , each buyer's expected utility is zero. Consider a single deviating buyer j using the signalling strategy $\hat{\alpha}(v) = \alpha\left(\frac{v}{2}\right)$. The probability of winning is $G\left(\frac{v_j}{2}\right)$, and the expected utility $\frac{v_j}{2} \times G\left(\frac{v_j}{2}\right) > 0$ when $v_j \neq \underline{v}$. Thus the buyer has an incentive to unilaterally deviate from α . ■

What kind of non-invertible strategies are candidates for equilibrium? We draw on the work of Crawford and Sobel (1982) and consider *partition* strategies for buyers. Buyers partition the valuation space into a finite number of disjoint intervals and submit an identical signal for valuations belonging to the same interval. Let $X(M) \equiv \{x_0, x_1, \dots, x_M\}$ denote a partition of the valuation space $[\underline{v}, \bar{v}]$, $x_0 = \underline{v}$, $x_M = \bar{v}$, $x_{J-1} < x_J$, $J = 1, 2, \dots, M$. An M -partition strategy is a family of signalling functions $\{\gamma(b|v) : v \in [\underline{v}, \bar{v}]\}$ such that $\gamma(\cdot|v)$ is any distribution on the signalling interval $[\beta(x_{J-1}), \beta(x_J)]$ when $v \in [x_{J-1}, x_J)$, $J = 1, 2, \dots, M$, for distinct (increasing) values $\beta(x_J)$, $J = 1, 2, \dots, M$. With some abuse of notation, we will also let $X(M)$ denote the M -partition strategy. One may construe the partition strategy as an attempt by buyers to distinguish themselves from other (lower valuation) buyers while shielding their true valuation. The requirement that $\gamma(\cdot|v)$ be any distribution on the interval ensures, as noted in Gibbons (1992, pp.216-217), that there are no signals that are "out of equilibrium". We denote by s_J the equivalence class of all signals in the interval $[\beta(x_{J-1}), \beta(x_J))$, as the seller treats them exactly the same.

3 The Discriminating Seller

The seller collects all signals and offers the object to the buyer from the highest interval. If there is more than one buyer in the highest interval, then one of them is selected randomly by the seller. If the buyer refuses the seller's price, then the seller retains the object and the game ends. If the highest signal b corresponds to the class s_J , then the conditional density for the highest buyer's valuation is:

$$h(v|b \in s_J) = \begin{cases} \frac{f(v)}{F(x_J)-F(x_{J-1})}, & x_{J-1} \leq v < x_J \\ 0, & \text{else} \end{cases} \quad (3)$$

The seller computes the price, p_J , according to (1). It is clear that $p_J \geq x_{J-1}$. Note also that $p_J < x_J$, $J \neq 0$. This is because $p_J \geq x_J$ gives zero revenue; however a price p_J , such that $x_{J-1} \leq p_J < x_J$ makes a non-zero revenue of $p_J \times \frac{F(x_J)-F(p_J)}{F(x_J)-F(x_{J-1})}$. We are now ready to characterize the partition equilibria of the PD game:

Theorem 1 $(X^*(M^*), p^*)$ is a partition equilibrium in the PD game if and only if:

- (1) $M^* = 1$.
- (2) The seller's price p^* is a constant and satisfies $p^* = \max(\underline{v}, \frac{1-F(p^*)}{f(p^*)}) = \frac{1}{\lambda(p^*)}$, where λ denotes the hazard rate.

Remark: Theorem 1 states that when buyers are certain that the seller will price discriminate, then their signals convey no information to the seller. The partition strategy is the coarsest one given by $X^*(M^*) = \{0, 1\}$.

Proof. (Necessity). Consider a partition equilibrium $(X^*(M^*), p^*)$ where p^* is not necessarily constant, but depends on the signals. We first establish that the probability of winning is greater with a signal in an interval with a higher index. Denote by W_J the probability of winning with a signal in the equivalence class s_J . One of the events leading to winning is the event that all other buyers have a valuation not greater than x_{J-1}^* ; its probability is $G(x_{J-1}^*)$. However, this is not the only event, as a buyer can win the item even if another buyer is in the same winning interval (in which case the winning buyer is chosen at random). Hence, $W_J > G(x_{J-1}^*)$. Further, the event that all bidders have valuations smaller than x_J does not ensure a win, hence $W_J < G(x_J^*)$. Now consider two indices, $J^{*'} < J^*$. Then $W_{J^{*'}} < G(x_{J^{*'}}^*) \leq G(x_{J^*-1}^*) < W_{J^*}$.

Let $p_{J^*}^*$ denote the price that is charged if the highest signal is from equivalence class s_{J^*} . We claim that in equilibrium $p_{J^*}^* = p_{J^{*'}}^* = p^*$ for all $1 \leq J^{*'} < J^* \leq M^*$. Consider buyers with valuations $v \in [p_{J^*}^*, \min(p_{J^*}^* + w, x_{J^*}^*)]$, where

$$w = \frac{p_{J^*}^* - p_{J^{*'}}^*}{W_{J^*} - W_{J^{*'}}} W_{J^{*'}}$$

Since $W_{J^*} > W_{J^{*'}}$, w is positive. For these buyers, the payoff from sending a signal from the lower interval $(v - p_{J^{*'}}^*) \times W_{J^{*'}}$ is strictly greater than that obtained from sending one from the higher interval, $(v - p_{J^*}^*) \times W_{J^*}$. Hence these buyers have an incentive to deviate, contradicting the hypothesis that $(X^*(M^*), p^*)$ is an equilibrium. Because $x_J > p_J \geq x_{J-1}$, $p_{J^*}^* = p_{J^{*'}}^* = p^*$ is only possible if $M^* = 1$, $X^*(M^*) = \{0, 1\}$.

(Sufficiency). Given that $X^*(M^*) = \{0, 1\}$, no information is conveyed to the seller. Thus the seller chooses the price p^* that maximizes $p(1 - F(p))$. A unilateral deviation by any buyer does not change payoffs because all signals convey the same coarsest information that valuations are distributed on $[0, 1]$ according to F . ■

It is clear that the equilibrium is in general inefficient. There are two sources of inefficiency. First, the seller may (randomly) allocate the object to a buyer who does not have the highest valuation. Second, since p^* may be strictly greater than \underline{v} , with positive probability the seller may choose a buyer who will reject the object. Thus the seller will retain the object even though there are buyers with positive valuations.

4 The λ -Discriminating Seller

In this section we demonstrate that the seller can induce the buyers to divulge more information if he chooses not to price discriminate with positive probability. We say that the seller's pricing strategy is *non-discriminatory* if the seller charges a fixed price that is *not* based on exploiting the information contained in the winning signal. This is what distinguishes a non-discriminatory price from a discriminatory one based on extracting information on valuation from the signals and charging a price to capture all the surplus. Adherence to a non-exploitative fixed price rule is also precisely what is required from the seller in order to induce buyers to reveal more private information. The main thrust of the paper is that such a fixed rule, even though it ignores information contained in the buyers' signals, can in fact be revenue enhancing when compared to the case of always utilizing the information.

We will let the non-discriminatory price charged be equal to $\beta(x_{J-1})$ when the highest signalling interval observed is $[\beta(x_{J-1}), \beta(x_J)]$. The rationale is that all signals in the interval $[\beta(x_{J-1}), \beta(x_J)]$ are considered equivalent by the seller. The seller accordingly chooses a fixed

price that does not exclude any buyer signalling in this interval, thus following the model of the first-price auction when buyer strategies are partition strategies. It is important to note that such a non-discriminatory price is *not* arbitrary. The buyers' signalling strategy, and the seller's prices (both discriminatory and non-discriminatory), are endogenously determined as mutual best responses in the Bayesian-Nash equilibrium of the game described below.

We consider a modified λ -PD game in which the seller once again offers the object to those with signals in the highest interval. With probability λ , the price charged is discriminatory and maximizes the seller's revenue. With probability $1 - \lambda$, the price is non-discriminatory and equal to the lower end of the signalling interval. We do not view the seller as committing to the value of λ . Instead, we think of the buyers as being able to accurately estimate λ as part of the seller's reputation. (For example, eBay might include the fraction of sales that were second-chance offers in the seller reputation it currently provides.) This may be formalized by considering λ a seller type known to the buyers. The buyer providing a signal in class s_J commits to price $\beta(x_{J-1})$, hence $\beta(x_{J-1}) \leq x_{J-1}$. Further, the seller will not sell the object for a price smaller than \underline{v} , hence $\beta(\underline{v}) \geq \underline{v}$. Notice that this implies that $\beta(\underline{v}) = \underline{v}$.

Recall that in the PD game ($\lambda = 1$), the buyers provide no information at all in their signals, i.e. there is a single partition in the signalling strategy. However, when $\lambda \neq 1$, we show that a partition signalling strategy with *two* partitions exists for all $\lambda \in (0, 1)$. This suggests that the non-zero probability of choosing a non-discriminatory price motivates the buyers to provide more information to the seller. Let the partition be $(\underline{v}, v_a, \bar{v})$, where $v_a = \underline{v} + a(\bar{v} - \underline{v})$ for some $a \in (0, 1)$. Let $p^*(s_1)$ and $p^*(s_2)$ be the discriminatory prices charged when the highest signals are s_1 and s_2 respectively. Note that $p^*(s_1)$ and $p^*(s_2)$ are functions of the partition, parameterized by a . With some abuse of notation, we denote $p_1(a) = \frac{p^*(s_1) - \underline{v}}{\bar{v} - \underline{v}}$ and $p_2(a) = \frac{p^*(s_2) - \underline{v}}{\bar{v} - \underline{v}}$. Also, we explicitly note the dependence of W_J on the parameter a . The strategy $(\beta(\cdot), a)$ will be determined endogenously.

Theorem 2 *A symmetric equilibrium signalling strategy with partition $M = 2$ exists for all $\lambda \in (0, 1)$.*

Proof. Let $\phi(s_i, a, v)$ denote the buyer's payoff for signal s_i , partition $(\underline{v}, v_a, \bar{v})$ and valuation v . Consider the function $\psi(a) = \phi(s_2, a, v_a) - \phi(s_1, a, v_a)$, the difference in payoff between signalling a higher valuation (one in $[v_a, \bar{v}]$) and a lower valuation (one in $[\underline{v}, v_a]$), at the common point of these intervals, v_a . Then:

$$\begin{aligned}\phi(s_2, a, v_a) &= (1 - \lambda)(v_a - \beta(v_a))W_2(a) + \lambda(0) \\ \phi(s_1, a, v_a) &= \{(1 - \lambda)(v_a - \underline{v}) + \lambda(v_a - p^*(s_1))\}W_1(a)\end{aligned}$$

and it follows that:

$$\psi(a) = (1 - \lambda)(v_a - \beta(v_a))W_2(a) - (v_a - (1 - \lambda)\underline{v} - \lambda p^*(s_1))W_1(a)$$

Equating $\psi(a) = 0$ gives us

$$\beta^*(v_a) = \underline{v} + (a - \frac{(a - \lambda p_1(a))W_1(a)}{(1 - \lambda)W_2(a)})(\bar{v} - \underline{v})$$

Clearly, $\beta^*(v_a) \leq v_a$. Now we show that $\exists a = a^* \in (0, 1)$ such that $\beta^*(v_{a^*}) \geq \underline{v}$. Consider:

$$\rho(a) = (1 - \lambda)W_2(a) - W_1(a)$$

the difference in winning probabilities, for the buyer with valuation v_a , between signalling in the higher valuation and the lower valuation. Because $W_2(v_a) = \frac{\sum_{i=0}^{N-1} F^i(v_a)}{N}$, and $W_1(v_a) = \frac{F^{N-1}(v_a)}{N}$, $\rho(0) = \frac{1-\lambda}{N}$ and $\rho(1) = (1 - \lambda) - \frac{1}{N}$. As ρ is a continuous function, $\exists a = a^* \in (0, 1)$ such that $\rho(a^*) > 0$, and $\beta^*(v_{a^*}) > \underline{v}$.

To see that the signalling strategy with partition $(\underline{v}, \beta^*(v_{a^*}), \beta^*(\bar{v}))$, corresponding to the valuation partition $(\underline{v}, v_{a^*}, \bar{v})$ is an equilibrium strategy, consider a deviating buyer. The difference in expected payoffs of a buyer with valuation $v_{a^*} + \delta(\bar{v} - \underline{v})$ is

$$\phi(b_2, a^*, a^* + \delta) - \phi(b_1, a^*, a^* + \delta) \geq 0, \quad \delta \geq 0$$

Hence a buyer has no incentive to deviate unilaterally. ■

In order to explore the implications of a two-partition equilibrium, we will now restrict ourselves to the case of uniformly distributed valuations. Let $F(v) = \frac{v-\underline{v}}{\bar{v}-\underline{v}}$. Further, suppose that $\underline{v} \geq \frac{\bar{v}}{2}$. Then $p^*(s_i) = x_{i-1}$. That is, the discriminatory price is the low end of the signaled valuation interval. In particular, the discriminatory and non-discriminatory prices for the lowest interval are both \underline{v} . For example, in the two buyer case, $W_2(a) = \frac{a+1}{2}$ and $W_1(a) = \frac{a}{2}$. $\beta^*(v_a) = v_a - \frac{a\frac{a}{2}}{(1-\lambda)\frac{1+a}{2}}(\bar{v} - \underline{v})$ which is greater than \underline{v} when $a \leq \frac{1-\lambda}{\lambda}$ (and perhaps also for other values of a). The buyer does not have an incentive to deviate for these values of a .

We can now compare the seller's expected revenue from the 2-partition equilibrium in the λ -PD game to that in the babbling equilibrium of the PD game. In particular, we show that the seller gains when he sometimes strategically ignores the information contained in the signals.

Proposition 1 *Let $F(v) = \frac{v-\underline{v}}{\bar{v}-\underline{v}}$, $v \in [\underline{v}, \bar{v}]$, $\underline{v} \geq \bar{v}/2$, and $N = 2$. For a 2-partition equilibrium the expected revenue of the seller in the 2-partition Bayesian-Nash equilibrium of the λ -PD game is independent of λ and exceeds that in the PD game where the seller always price discriminates.*

Proof. Consider the equilibrium bidding strategy with valuation partition $(\underline{v}, \underline{v} + a(\bar{v} - \underline{v}), \bar{v})$. Since $\underline{v} \geq \bar{v}/2$, $p^*(s_2) = v_a$ and $p^*(s_1) = \underline{v}$. The seller's revenue R is:

$$R = \underline{v} + (\bar{v} - \underline{v}) (1 - a^2) (\lambda a + (1 - \lambda) \left[a - \frac{(a)^{\frac{a}{2}}}{(1 - \lambda)^{\frac{1+a}{2}}} \right]) = \underline{v} + (\bar{v} - \underline{v}) a (1 - a)$$

which has a maximum at $a = \frac{1}{2}$. This maximum value is $R^* = \underline{v} + \frac{\bar{v} - \underline{v}}{4}$. In this case, the revenue of the PD game would have been \underline{v} . ■

5 Optimal Auctions

We now discuss our results in the context of the literature on optimal auctions which seek to maximize the seller's expected revenue.¹ The set of allocation rules over which the seller can search is in principle extremely large and complex because it can include virtually any manner of allocation and pricing. The seminal paper of Myerson (1981) simplified the problem considerably by restricting attention to direct mechanisms in which bidders directly report their valuations to the seller. The seller *commits* to allocation and pricing rules which are functions of the reports and which induce the bidders to voluntarily participate in the mechanism (individual rationality) and truthfully report their valuation in equilibrium (incentive compatibility, or IC). The solution with symmetric bidders (and some regularity restrictions) is a second-price auction with discriminatory reserve prices derived from the virtual valuations of the bidders (Krishna 2002, Proposition 5.4). The analogy with our framework arises because the optimal auction can be understood as optimal third degree price discrimination by the seller (Bulow and Roberts 1989, Cairns 1993). The good is assigned to the buyer with the highest positive marginal revenue who then pays a price corresponding to the larger of the reserve price and the second highest marginal revenue. The seller retains the item if the highest marginal revenue is negative. The amount paid by the winning buyer is less than the valuation, and this positive "informational rent" is conceded by the seller in order to induce truthful revelation of private valuations.

The key difference between our model and the optimal auction literature is that we do not impose IC. This can also be seen by the fact that IC is equivalent to the ex-ante probability of winning the object strictly increasing in the buyer's report. This is not true in

¹We are very grateful to a referee for pointing out the similarities between our model and optimal auctions.

our model where the seller picks any one buyer whose signal lies in the highest signalling interval. Our relaxation of IC is prompted by real world evidence on second-chance offers or phantom bidding through which sellers seek to dilute their commitment and reduce the informational rents that have to be paid out. Once IC is eliminated, buyers have an incentive to hide their true valuation. This leads to the next point of difference from optimal auctions: buyer strategies are no longer strictly increasing (invertible) in valuations. Drawing on the literature on strategic information transmission, we limit the buyers to partition strategies. The reason is that this seems to be the most general, yet tractable, way of modeling how the incentive of buyers to provide information varies with the degree of credibility that the seller will abide by the given price and allocation rule. Truthful revelation under IC is of course a special case when the partition is the finest.

Since we do not impose IC, the optimal auction in our model is not the second-price auction with discriminatory reserve prices. However some of our results bear similarity to classical optimal auctions. When the seller does not make any commitment to a fixed price (PD game), then the “optimal auction” we derive is one where the seller offers the object to a randomly chosen bidder at a price equal to the reserve price obtained in standard auctions (for instance, Krishna 2002, p. 26).² When the seller partially commits to a fixed price (λ -PD game), then in our “optimal auction” one of the bidders signalling in the highest interval is randomly chosen as the winner. When the seller is non-discriminatory with probability λ , the winner pays a price equal to the lower end of the signalling interval and earns informational rents. This is analogous to classical optimal auctions because the bidder obtaining the item pays the smallest amount that is required for being one of the winners. Also, similar to optimal auctions, the seller retains the item if the bid falls below this minimum amount.

6 Conclusion

This paper examined a first-degree price discrimination model in which buyers sent (possibly noisy) signals to the seller regarding their private valuation. We demonstrated that it was sometimes in the interest of the seller to strategically ignore price discrimination and instead adhere to a non-discriminatory pricing rule. An interesting unanswered question is whether there is any partition for which the seller’s revenue is greater than or equal to that implied

²In fact it is interesting to note that this outcome is identical to the one obtained for an all-inclusive cartel in a first-price auction when the inverse of the hazard rate is monotonically increasing. McAfee and McMillan (1992) show that in the optimal auction all bidders submit a bid equal to the reserve price and the seller randomly allocates the object to one of the bidders.

by revenue equivalence (which holds for invertible strategies). A complete characterization of partition equilibria, and the role of risk aversion and ex-ante buyer heterogeneity, remain productive avenues for future research.

References

- [1] J. Bulow and J. Roberts (1989) “The Simple Analytics of Optimal Auctions”, *Journal of Political Economy* 97(5), 1060-1090.
- [2] R.D. Cairns (1993) “The Optimal Auction: A Mechanism for Optimal Third-Degree Price Discrimination”, *Journal of Economic Behavior and Organization* 20, 213-225.
- [3] V.P. Crawford and J. Sobel (1982) “Strategic Information Transmission”, *Econometrica* 50(6), 1431-1451.
- [4] R. Gibbons (1992) *Game Theory for Applied Economists*, Princeton University Press.
- [5] V. Krishna (2002) *Auction Theory*, Academic Press.
- [6] P. McAfee and J. McMillan (1992) “Bidding Rings”, *American Economic Review* 82, 579-599.
- [7] D. Lucking Reiley (2000) “Vickrey Auctions in Practice: From Nineteenth-Century Philately to Twenty-First-Century E-Commerce”, *The Journal of Economic Perspectives* 14, 183-192.
- [8] R.B. Myerson (1981) “Optimal Auction Design”, *Mathematics of Operation Research* 6(1), 58-73.
- [9] T.C. Salmon and B.J. Wilson (2006) “Second Chance Offers versus Sequential Auctions: Theory and Behavior”, *Economic Theory* (forthcoming).
- [10] H.R. Varian (1989) “Price Discrimination”. In R. Schmalensee and R.D. Willig (ed). *Handbook of Industrial Organization*, Elsevier Science Publishers, 597-654.