

## PRIVACY IN A MULTI-STAGE GAME-- AN EVOLUTIONARY PROGRAMMING APPROACH

SUMIT JOSHI

*Department of Economics, The George Washington University, 1922 F St NW  
Washington DC, 20052, U.S.A.*

YU-AN SUN, POORVI VORA

*Department of Computer Science, The George Washington University, 800 22<sup>nd</sup> St NW  
Washington DC, 20052, U.S.A.*

The *second chance offer* is a common seller practice on eBay. It consists of price discrimination against the losing bidder, who is offered an identical item at the value of his or her highest bid. This paper approaches this as a privacy problem, which is posed as one of information revelation in a multi-stage game. [7] has shown that, if bidders are rational, they change their bidding strategies to compensate for the price discrimination so that the seller is not motivated to violate privacy in this manner. [8] has shown this is true even if the seller price discriminates with a probability  $\alpha$ , when there are as many items as there are bidders. In this paper, we examine the cases when the items are insufficient for all bidders – to determine whether privacy violation is more acceptable to the bidder, in the presence of scarcity. Using both Bayesian techniques and evolutionary programming simulations, we determine optimal strategies for bidders and sellers. We present the Nash equilibria obtained, and determine that there is at least one case when sellers are motivated to perform randomized price discrimination because of the second mover advantage. Our use of evolutionary programming simulations is unique in the privacy literature.

### 1. Introduction

We treat interactions among entities as stages in a multi-stage game. Information revealed in an earlier stage can reduce (or increase) a player's utility in a later stage. An example is the case of differential pricing following an auction, when goods are priced higher for high bidders. In this paper, we focus on differential pricing in general, and eBay's *second-chance offer* in particular.

Several papers have suggested that differential pricing is the goal of privacy invasion and usage tracking [1, 10, 12]. eBay.com's second chance offer mechanism is an example of differential pricing. It allows sellers to offer non-

winning bidders an identical item at the price of their highest failed bid. While meant to be used when the winning bidder does not pay, or when the Reserve Price has not been met [5], this mechanism can be abused to provide a powerful tool for differential pricing: a bidder's valuation of the item is learnt from her bid(s) in the auction itself, and used as a price if she loses the auction.

We examine whether the seller can benefit by offering the second chance offer, not with certainty, but with a probability  $\mathbf{a}$ . That is, after one object is sold in a first-price auction in Stage I, the seller offers, in Stage II, either (a) (with probability  $\mathbf{a}$ ) a privacy invasive second chance offer to all individual bidders, or (b) (with probability  $1-\mathbf{a}$ ) a privacy-protecting uniform-price offer  $P$ . *We assume that the items are scarce, that is, that the number of items is smaller than the number of bidders.* Our contributions are as follows:

**Case I**  $\mathbf{a}$  and  $P$  form the reputation of the seller and are known to the bidder. The values of  $\mathbf{a}$  and  $P$  may be seen as representing the probability distribution on the seller's type – his offer in the second stage – which is unknown *a priori* to the bidder. We examine the Bayesian Nash equilibrium and show that a monotonic increasing optimal strategy for the bidders does not always exist. In these cases, we find that evolutionary programming experiments provide optimal strategies that may not be monotonic increasing.

**Case II:**  $\mathbf{a}$  and  $P$  form the seller's second-mover strategy and are unknown to the bidder. There are  $n$  bidders participating and  $k$  items available, where  $1 < k < n/2 + 1$ . We use evolutionary programming experiments to derive equilibria for  $n=8$ . We find that when the number of available items is less than or equal to four (half of the number of bidders), it is the seller's best response to always price discriminate in the second stage. In these cases, the bidding strategy found in simulations is consistent with the Bayesian strategy derived in Case I, with  $\mathbf{a}=1$ , and no second mover advantage for the seller is shown. This important finding also demonstrates the ability of evolutionary programming simulations to obtain Nash equilibrium for this problem. When the number of available items equals to five ( $n/2 + 1$ ), we find that the seller's value of  $\mathbf{a}$  does not converge. With further numerical analysis, we confirm that  $\mathbf{a}$  does not converge because the seller's total revenue is the same for any  $\mathbf{a}$ .

Thus we determine that, in the presence of sufficient scarcity (when the number of items is smaller than or equal to half the number of bidders) the seller is motivated to price discriminate. Because of the scarcity, bidders do not lower their bids sufficiently. The paper is organized as follows. Section 2 presents related work, section 3 preliminaries, and section 4 analytical results. Section 5 presents the results of an evolutionary programming simulation of the same game,

played with sequential moves and with the seller playing second. Section 6 contains conclusions.

## 2. Related Work

Knowledge of the valuations of several individuals for a particular item can enable a fairer distribution of the item and an efficient estimation of demand. It also enables price discrimination – the practice of charging different amounts for identical items to different customers. Vendors are motivated to reduce consumer privacy in order to improve the accuracy of price discrimination [1, 10]. On the other hand, if customers are rational, they would anticipate the availability of information and adjust their strategies to obtain higher payoffs [1, 7], making price discrimination a suboptimal strategy for the seller. Though experimental results indicate that customers are not rational or do not value their privacy enough [2], it is possible that consumers would behave differently if the costs of behaving rationally were smaller [7] (if, for example, automated rational agents were easily available and easy to use).

Previous work [7] by the authors has shown that rationality does enable bidders to protect privacy when the seller's strategy is deterministic, that is, when  $\alpha = 1$  and the seller always price discriminates. They have also shown [8] that rationality protects bidders when the seller uses a randomized strategy and the number of items is identical to the number of bidders – that is, in the absence of scarcity. To the authors' knowledge, [8] was the first to use evolutionary programming to study privacy.

The use of evolutionary programming was first devised by John Holland in 1975 [6]. A large number of studies have applied genetic algorithms to well-known economic problems. For example, [4] addresses the prisoner's dilemma and shows that the tit-for-tat strategy evolves in a genetic algorithm simulation of the prisoner's dilemma game, predicting cooperation among prisoners, and a simple cobweb model [3] is used to analyze supply and demand.

## 3. Preliminaries

We model a two stage price discrimination game as follows:

- Stage I: The seller has  $k$  available identical items.  $n$  bidders join a first-price sealed-bid auction. All bidders simultaneously and independently make bids. The bidder with the highest bid wins the auction. On the occurrence of a tie, the winner is chosen at random. The remaining  $N-1$  bidders enter Stage II.
- Stage II: The seller offers  $k-1$  identical items to the  $k-1$  highest of the remaining  $N-1$  bidders: (1) The *failed bid (privacy-infringing) option*: With

probability  $\mathbf{a}$ , the price is the bidder's highest bid in Stage I. (2) The *uniform price(privacy-protecting) option*: With probability  $1-\mathbf{a}$ , the price is a uniform price for all bidders. The bidders can reject or accept either offer.

Each bidder's payoff is computed as the difference between the prices paid for the item and the bidder's valuation. The valuations of all the bidders are modeled as being uniformly and independently distributed between  $0$  and  $1$ .

We examine two cases. The number of available items,  $k$ , is known to the bidders in both cases. In Case I,  $\mathbf{a}$  and  $P$  represent the seller's reputation and are known to the bidder. The seller's second move is not strategy, but determined by his type, which is characterized by probability  $\mathbf{a}$  and fixed price  $P$ . In Case II,  $\mathbf{a}$  and  $P$  represent the seller's second-mover strategy and are unknown to the bidder. We denote the (private) valuation by  $x$ , the bid by  $b$ , the optimal bidding function by  $\mathbf{b}$ , and the number of bidders by  $N$ . For the symmetric Nash equilibrium we assume, as is standard, that bidder strategies are identical, and all strategies are known to all players.  $G(x)$  denotes the probability that a given valuation  $x$  is the highest among  $N$  bidders.  $H(x)$  denotes the probability that a given valuation  $x$  is among the  $k$  highest, but not the highest.

#### 4. Case I: $\mathbf{a}$ and $P$ Known to Bidders

Consider a simplified version of the game, where the bidder knows the sellers' values of the probability  $\mathbf{a}$  and the uniform price  $P$  prior to the start of the auction. This may be thought of as a steady state setting in a repeated game, where the seller has chosen optimal values of  $\mathbf{a}$  and  $P$ , and the bidder has learnt them over repeated interactions.  $\mathbf{a}$  and  $P$  represent the reputation of the seller, and we may consider them as representing the distribution on the type of the seller. Just as the distribution on the bidder's valuation is known to the seller, so also the parameter  $\mathbf{a}$  is known to the bidder, along with  $P$ . The bidders submit bids in the first stage; after the auction ends, the system flips a coin, biased according to the value of  $\mathbf{a}$ , to choose between the uniform price option and the failed bid option. The result of the coin flip may be viewed as the type of the seller – much as the valuation  $x$  represents the type of the bidder – and is similarly unknown, *a priori*, to the other players. Because of the insufficient number of items available for all bidders, only with probability  $H(x)$  will a bidder make a profit entering the second stage.

It can be shown that there is no dominant deterministic strategy for bidders. The symmetric Bayesian Nash equilibrium strategy is stated in equation (1) but not derived here due to space constraints. We find that, for certain  $\mathbf{a}$  and  $P$ , the bidding strategy  $\mathbf{b}(x)$  is not monotonic increasing for  $x > P$  and thus contradicts the assumption that bids are invertible. Thus, any attempts to obtain the strategy

in these cases would require the assumption that the strategy is not invertible; optimal strategies may be obtained using evolutionary programming experiments.

$$b(x) = \begin{cases} \frac{\frac{n-1}{n}x + \frac{\int xH'(x)}{G(x)} - (1-a)P}{1 + a \frac{H(x)}{G(x)}} & x \geq P \\ \frac{\frac{n-1}{n}x + a \frac{\int xH'(x)}{G(x)}}{1 + a \frac{H(x)}{G(x)}} & x < P \end{cases} \dots\dots\dots(1)$$

## 5. Case II: $a$ and $P$ Unknown to Bidders

In a less constrained situation, the two-stage game would be played sequentially, with the seller moving later than the bidders. That is, the seller would determine the values of  $\alpha$  and  $P$  after receiving all the submitted bids. The sequential game may be approached as an asymmetric, incomplete information game; we approach it using evolutionary programming simulation. Assuming the game is played repeatedly, we examine how the seller and different types of bidders evolve over generations.

### 5.1. Evolutionary Programming Simulation

#### 5.1.1. Simulation Method

We choose chromosomes consisting of bidder's bids for the fitness function *expected payoff*, and the seller's values of  $a$  and  $P$  for *expected revenue*. The bidder's chromosome is coded as a lookup table of 20 entries of real numbers; the bidder's bidding function is a piecewise linear interpolation of this lookup table. We use the lookup table so that the bidding function is not restricted to be in a particular parameterized form. The seller's chromosome is coded as 2 real numbers  $a$  and  $P$ , representing the seller's action after Stage I.

#### 5.1.2. Simulation Steps

The evolutionary programming simulation has the following basic steps: (1) create the initial population; (2) evaluate the current population; (3) reproduce: selection, crossover and mutation; (4) repeat steps 2-3 until the stop condition is reached. We stop our simulation when the number of generations is 40,000. The initial population is set at 500 sellers, and 8 distinct bidders are used for each seller. We conduct simulations for available item numbers  $K=2$  to 5. Bidder chromosomes are randomly generated at the beginning of the simulation. Bidder

valuations are randomly generated *each generation*. The tournament selection is used for reproduction: two chromosomes will be randomly drawn from the population pool. The chromosome with the higher fitness score will be copied to the new population representing the next generation. It will stay in the pool for further tournaments, and the process continues until the new population has the same size as the previous one. The crossover process combines chromosomes from two parents, at random. The optimal mutation rate is set to be equal to  $1/22$  as suggested in [9] because of 22 different variables in the sellers and bidders chromosomes. We also implement the mutation operator suggested in [9].

After the population is initialized, each bidder submits a bid by substituting its valuation into its lookup table. For example, a bidder with valuation  $0.80$  will submit a bid of  $L(16)$  where  $L$  is the lookup table and its  $16^{\text{th}}$  entry is the bid. Notice that, if the bidder's bidding strategies are correct, this is indeed an optimal bid. The seller's fitness score is the total revenue over both stages; that of the bidder is the payoff – the difference between its purchasing price and its valuation. If the seller offers a fixed price higher than the bidder's valuation, the offer is rejected. The fitness functions take inputs from the chromosomes of both sellers and bidders. Sellers and bidders can be viewed as two species that affect each other while evolving.

## 5.2. Simulation Results

The simulation results for sellers' chromosomes are summarized in table 1. Note that  $P$  does not converge because  $\mathbf{a}$  converges to 1; meaning sellers never choose a fixed price offer.

Table 1. Sellers Chromosomes

|       | Item Number |               | P             |
|-------|-------------|---------------|---------------|
| Exp 1 | 2           | 1.0           |               |
| Exp 2 | 3           | 1.0           | Not Converged |
| Exp 3 | 4           | 1.0           |               |
| Exp 4 | 5           | Not Converged | 0.5           |

Results for experiments 1, 2 and 3 are consistent with the Bayesian bidding strategy. That is, in these cases, an invertible strategy happens to be optimal. We plot the Bayesian bidding function for  $\mathbf{a}=1$  and the function obtained from the bidders' lookup table in figure 1. This optimal strategy is not invertible and is hence not obtainable through Bayesian analysis. We further conduct a numerical analysis that confirms the simulation result is the optimal bidding strategy. Bidders have to bid higher than Case I in response to the sellers' second mover advantage.




Figure 1. Experiment Four

## 6. Observation and Conclusion

We have demonstrated that evolutionary programming is a powerful approach to discover Nash equilibria in multi-stage games, particularly when non-invertible optimal strategies cannot be obtained through Bayesian analysis. At the same time, simulation results are consistent with the Bayesian strategy, when the optimal strategies are invertible. We have further shown that randomizing between fixed price offer and price discrimination is optimal for sellers. However, due to the second mover advantage of the sellers; it results in bidders' payoff loss, which is also viewed as the cost of privacy invasion.

## References

1. A. Acquisti and H. Varian, *Market Science*, **24(3)**, 1-15 (2005).
2. A. Acquisti, *Pro. of the ACM Electronic Commerce*. (2004).
3. J. Arifovic, *Journal of Economic Dynamics and Control*, 18(1994).
4. R. Axelrod, *Princeton University Press*, (1997).
5. [http://pages.ebay.com/help/sell/second\\_chance\\_offer.html](http://pages.ebay.com/help/sell/second_chance_offer.html)
6. J. H. Holland, *University of Michigan Press*. (1975).
7. S. Joshi, Y. Sun and P. Vora, *ACM WPES*,. (2005).
8. S. Joshi, Y. Sun and P. Vora, *ACM WPES*,. (2006).
9. H. Mühlenbein and D. Schlierkamp-Voosen, *Evol. Com.*, **1(1)**, 25-49 (1993).
10. A. M. Odlyzko, *Pro. of the ACM ICEC*. 355-366(2003).
11. H. Varian, *National Telecommunications and Information Administration*,. (1996).