Profit Sharing Auction

Sandip Sen and **Teddy Candale** and **Susnata Basak**
Mathematical &Computer Sciences Department
University of Tulsa
{sandip, teddy-candale, susnata-basak}@utulsa.edu

Abstract

Auctions are a class of multi-party negotiation protocols. Classical auctions try to maximize social welfare by selecting the highest bidder as the winner. If bidders are rational, this ensures that the sum of profits for all bidders and the seller is maximized. In all such auctions, however, only the winner and the seller make any profit. We believe that "social welfare distribution" is a desired goal of any multi-party protocol. In the context of auctions, this goal translates into a rather radical proposal of profit sharing between all bidders and the seller. We propose a Profit Sharing Auction (PSA) where a part of the selling price paid by the winner is paid back to the bidders. The obvious criticism of this mechanism is the incentive for the seller to share its profit with non-winning bidders. We claim that this loss can be compensated by attracting more bidders to such an auction, resulting in an associated increase in selling price. We run several sets of experiments where equivalent items are concurrently sold at a First Price Sealed Bid, a Vickrey, and a PSA auction. A population of learning bidders repeatedly choose to go to one of these auctions based on their valuation for the good being auctioned and their learned estimates of profits from these auctions. Results show that sellers make more or equivalent profits by using PSA as compared to the classical auctions. Additionally, PSA always attracts more bidders, which might create auxiliary revenue streams, and a desirable lower variability in selling prices. Interestingly then, a rational seller has the incentive to share profits and offer an auction like PSA which maximizes and distributes social welfare.

Introduction

Designers of agent societies are interested in constructing environments and interaction protocols that promote stable systems serving all individuals. Characteristics of desirable outcomes of multi-party interactions include stability and efficiency. Negotiation and coordination mechanisms that produce desirable outcomes have therefore been extensively studied in multiagent literature (Kraus 2001). General multiagent interaction protocols typically address issues of distribution of profits between all parties concerned. Other researchers, e.g., in cooperative game theory (Eichberger 1993) use concepts like Shapley value to promote fair distribution of profits between members of a group even though a number of other distributions, including some with no profit for certain members, will maintain stability of the group[1].

One particular class of such multi-party negotiation protocols, viz., auctions, have been used in a variety of multiagent domains to facilitate transactions of goods and resources. Auctions have also received increased prominence with the success of popular Internet sites like eBay, uBid, etc. Multiagent researchers have both developed customizable auction sites, e.g., eMediator (Sandholm 2000), auctionBot (Wurman, Wellman, & Walsh 1998), etc., as well as investigated new auction protocols (Parkes 2001; Suyama & Yokoo 2004) and bidding strategies in different auction settings (Greenwald & Boyan 2005; Stone *et al.* 2003).

The premise of this work is that one of the preferred goals of any multi-party protocol is social welfare distribution. More specifically, this dictates that if an interaction of many entities create a profit or surplus in the system, that amount should be distributed among the entities.

Given that auctions are a particular class of multi-party interaction protocols, we can evaluate their effectiveness in view of the criterion of social welfare distribution. Now, one of the desired property accepted in auction theory is that of *social welfare maximization*. Let r be the reservation price of the seller of an item in any auction. Let $\mathcal{B}$ be the set of bidders in the auction. If v_h is the highest valuation for the item among any bidder in the auction, then the maximum surplus that can be created by the interaction of the seller and the bidders $\mathcal{B}$ is $v_h - r$, assuming all bidder valuations are drawn from the same distribution. This social welfare maximization is realized when the item is allocated to the highest bidder, winner, at any selling price, p, such that $r \leq p \leq v_h$. This results in a profit of $p - r$ for the seller and $v_h - p$ for the winner with all non-winning bidders receiving zero profit. The goal of social welfare distribution, however, would suggest, rather radically, some profit distribution among all bidders and the seller.

While a philosophical debate about the desirability of social welfare distribution is beyond the scope of this paper, we briefly highlight two cases where such a distribution can be rationalized without axiomatic claims of the desirability

[1]Some payoff vectors in the *core* of the game may have zero payoff to some members.

of social welfare distribution:

- In the case of Vickrey's auction, the bid of the second highest bidder determines the selling price. Assuming this bid is unique, in the absence of that bidder the selling price would have been less, resulting in less profit to the auctioneer. One can then argue, using the analogy of cooperative game theory, that the auctioneer should share some of its profit with the second highest bidder.
- In a First Price Sealed Bid (FPSB) auction with N bidders, and bidder valuations drawn from an uniform distribution, the equilibrium strategy is to bid $\frac{N-1}{N}$ times ones valuation. With more bidders in the auction, the winning price, i.e., the sellers profit, increases. Here again, a case can be made that non-winning bidders are contributing to the profitability of the seller, and hence the total payoff to the "coalition" of the seller and non-winning bidders should be distributed, in some proportion, between them.

We propose a Profit Sharing Auction (PSA) where a part of the seller's profit is paid back to the bidders. Even if one agrees with the philosophy of social welfare distribution, there is still a more immediate and evident criticism of this proposal: what is the incentive for the seller to share its profit with the bidders? We posit that the loss from profit sharing can be more than compensated by attracting more bidders to such an auction, resulting in an associated increase in selling price. Attracting more bidders can also open up new revenue sources. While we do not delve into exploiting these additional income sources in this work, the pricing schemes in most popular web portals such as Yahoo and Google confirms that more visitors can lead to more revenues through auxiliary channels.

We empirically evaluate the claim of recouping loss from profit sharing by gains from additional participation. We experiment with a framework where where equivalent items are concurrently sold at a First Price Sealed Bid, a Vickrey, and a PSA auction. A fixed population of learning bidders repeatedly choose to go to one of these auctions based on their valuation for the good being auctioned and their learned estimates of profits from these auctions. We have varied the parameters of the experiment including the number of bidders in the population. The major findings from our experiments are as follows:

- For reasonable parameter values of the PSA auction, sellers make more or equivalent profits by using PSA as compared to the classical auctions. Thus sellers have the incentive, or at least have no disincentive, for offering PSA as compared to the classical auctions.
- PSA always attracts more bidders, which might create auxiliary revenue streams as suggested above.
- PSA delivers a lower variability in selling prices compared to the classical auctions. Thus PSA provides a more steady income stream for the seller which is desirable.

Interestingly then, we observe a win-win situation where a rational seller has the incentive to share profits and offer an auction like PSA which both maximizes and distributes social welfare.

Related work

Maximizing the seller's expected payoff from an auction is studied by a branch of auction theory known as *optimal auction*. The preferred approach is to show that the auction under consideration has *truth revelation* as property or that the seller revenue is equivalent to an auction where truth revealing is a dominant strategy. The assumption is that the seller will make the highest profit when the buyers have no incentive to speculate about other's valuations or lie about their own valuation of the item. Speculation, for example, can produce underbidding, which can cause both reduction in selling price, and at times, select someone other than the bidder with highest valuation as the winner, thereby failing to maximize social welfare.

The optimal auction literature dates back to the work of Vickrey (Vickrey 1961). His *revenue equivalence principle* implies that the four standard auctions: English ascending, Dutch, Vickrey, and FPSB, provides the same expected revenue to the seller. Only one of them, namely the second-price auction or Vickrey auction, is truth revealing. This theorem is applicable to every settings where *risk-neutral buyers share the same probability distribution regarding their private valuations, and the item is allocated to the buyer with the highest valuation*. Myerson (Myerson 1981) extends the theorem to the case where buyers do not share the same valuation distributions but these distributions are common knowledge.

Sealed-bid auctions

We consider an auction environment with one seller and N bidders. The seller wants to sell a single-unit item for which it has a reserve price equals to $r \geq 0$. We assume risk-neutral bidders have a private valuation for the item independently drawn from a common known cumulative distribution F with derivative $f = F'$ (*symmetric model*). Bids are drawn from the range $[\underline{v}, \overline{v}]$ where $0 \leq \underline{v} < \overline{v}$. In the remainder of the paper, s will denote the seller and i the i^{th} bidder. When there will be no ambiguity v_i and b_i will denote respectively the valuation and the bid of the i^{th} bidder. We assume, without loss of generality, $v_1 \geq v_2 \geq \cdots \geq v_N$.

Classical sealed-bid auctions

In classical sealed-bid auctions, each bidder sends its bid to the auctioneer without access to bids from other bidders. Two well-known sealed-bid auctions are the *first-price sealed-bid* (*FPSB*) auction and the *second-price sealed-bid* auction or *Vickrey* auction. The bidder with the highest bid is chosen as the winner. The winner pays its own bid in FPSB but only the second highest bid in the Vickrey auction. Rational bidders speculate and underbid in FPSB auction to obtain non-zero profit. It is proven the symmetric equilibrium bidding strategy consists in bidding the expected second highest valuation. If $Y_i^{(N)}$ is the random variable representing the highest valuation of all bidders but i, then the equilibrium bid strategy, for valuation v, is to bid $\beta(v) = E\left[Y_i^{(N)} | Y_i^{(N)} < v\right]$. Vickrey auction prevents speculation since *truth revelation* is one of its properties.

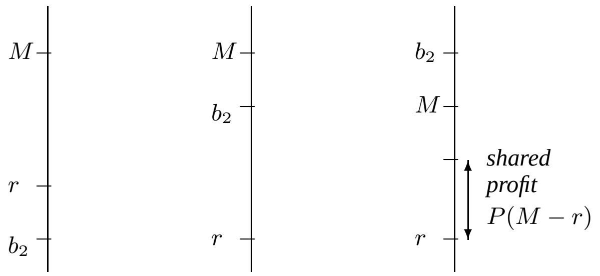

(a) No trade (b) Vickrey outcome (c) Profit sharing

Figure 1: Possible outcomes in PSA.

Bidding one's valuation is a dominant strategy in a Vickrey auction.

Profit Sharing Auction (PSA)

We now introduce the Profit Sharing Auction where the seller shares a portion of its profit with bidders. The auction is a modified second-price sealed-bid auction. We assume the seller has a sharing price threshold, M. It will return a portion P of a part, $M - r$, of its profit to all bidders whose bid is above M, where $M > r$. Qualified bidders, then, receive an extra profit of $\frac{P\cdot(M-r)}{N_M}$ where N_M is the number of bidders whose bid was greater than or equal to M. In the case where at most one bid is above M the outcome is identical to a Vickrey auction. Figure 1 shows the outcome of the auction in different situations.

PSA properties

The following presents desirable properties of an auction scheme (Jackson 2000):

Efficiency: The outcome of an auction is efficient when the sum of the seller and every bidder's utilities is maximized. If truth revelation is a dominant strategy, the bidder with the highest valuation wins the auction.

Individual rationality: Individuals weakly prefer participation in the auction to not participating, i.e., the expected utility of the seller and every bidder is non-negative.

Budget balance: The payment made by bidders equals the payment received by the seller.

Proposition 1 (Dominant strategy) *Truth revealing is a dominant strategy in PSA for a bidder with valuation greater than M.*

Proof: The proof is quite similar to the one for second-price auction (Vickrey 1961). Additionally, the fact the winner receives a part of the shared seller profit ensures that the winner has the highest profit of all bidders. Consequently, there is no reason to underbid to receive the shared profit without winning the item. □

Proposition 2 (Equilibrium strategy) *Bidding one's valuation when it is greater than M, and M otherwise, is an equilibrium strategy if $F(M)^N \leq P$.*

Condition 1 $F(M)^N \leq P$

Proof: Proposition 1 shows that an agent with a valuation greater than M will bid its valuation. We now consider an agent i with valuation below M. If all other agents bid as described in Proposition 2, the expected utility of i is 0 if it bids below M. If it bids M it gets the shared profit but takes the risk of getting the item if all bids are equal to M. Its expected utility of bidding M is $P \cdot \frac{(M-r)}{N} + \frac{F(M)^N}{N} \cdot (v_i - M) > \left(P - F(M)^N\right) \cdot \frac{M-r}{N} \geq 0$, under Condition 1 and $v_i > r$. It cannot do better by bidding more than M. □

An agent with a valuation below M has two choices: (i) bidding M to ensure the gain of the shared profit while taking the risk of getting the item at price above its valuation, (ii) bidding its valuation and not qualifying for the shared profit. Increase in the number of bidders in the auction has two opposite effects: (-) it decreases the part of the shared profit received by every agent, (+) it increases the probability that at least two agents have their valuations above M. The latter reduces the risk of the bidders who are overbidding, i.e., those with valuation less than M. Condition 1 ensures that the positive effect dominates the negative one.

Our auction does not have the truth revelation property, as bidders with valuation below M overbid. This, however, does not change the winner or the selling price if $b_2 \geq M$, i.e., at least two bidders have valuation $\geq M$. This is almost always the case for a reasonably large N and an appropriately chosen M.

Proposition 3 (Efficiency) *PSA is efficient under the assumption of risk-neutral bidders.*

Proof: If p denotes the selling price ($p = v_2$ or $p = M$), the seller utility is $p - r - P \cdot (M - r)$, the winner utility is $v_1 - p + P \cdot \frac{M-r}{N}$, the utility of other bidders is $P \cdot \frac{M-r}{N}$. The sum of utilities is then $v_1 - r$. □

Proposition 4 (Individual rationality) *PSA is individually rational under Condition 1.*

Proof: The utility of the seller is always positive since the item will never be sold at a price below r. Risk-neutral agents will always bid their equilibrium strategy. Condition 1 ensures the shared profit will always be enough to compensate the expected loss. □

Proposition 5 (Budget balance) *PSA is budget balanced.*

Proof: Payment by the winner is shared by the seller and all bidders with bids M and above. □

Proposition 6 (Closing price) *Expected closing price is*

$$N(1-F(M))F^{N-1}(M)M + E\left[Y^{(N)} | M \leq Y^{(N)}\right] + F^{(N)}(M)M$$

where $Y^{(N)}$ is the random variable representing the second highest bid under Condition 1.

Proof: PSA bidding strategy: $\beta^{\mathrm{PSA}}(v) = \max\{v, M\}$, ex post (after knowing valuation) expected payment:

$$m^{\mathrm{PSA}}(v) = \left| \begin{array}{ll} \frac{1}{N} F^{N-1}(M)\, M & \text{if } v < M \\ F^{N-1}(M)\, M + \int_M^v (N-1)\, v\, f(v)\, F^{N-2}(v) \mathrm{d}v \end{array} \right. ,$$

ex ante (before knowing valuation) expected payment:

$$\begin{aligned} E[m^{\text{PSA}}(v)] &= \int_{\underline{v}}^{\overline{v}} m^{\text{PSA}}(v)\, f(v)\mathrm{d}v \\ &= \frac{F^{\text{N}}(M)}{N}\, M + (1 - F(M))\, F^{\text{N-1}}(M)\, M + \\ & \int_{\text{M}}^{\bar{v}} \left(\int_x^{\bar{v}} f(v)\mathrm{d}v \right) (N-1) x f(x) F^{\text{N-2}}(x)\mathrm{d}x \\ &= \frac{F^{\text{N}}(M)}{N}\, M + (1 - F(M))\, F^{\text{N-1}}(M)\, M + \\ & \int_{\text{M}}^{\bar{v}} x\, (1 - F(x))\, f_1^{(\text{N})}(x)\mathrm{d}x, \end{aligned}$$

where $Y_1^{(\text{N})}$ is the random variable representing others' highest bid with probability distribution $f_1^{(\text{N})}$.
Expected closing price:

$$\begin{aligned} E[R^{\text{PSA}}] &= N\, E[m^{\text{PSA}}] \\ &= F^N(M)\, M + N\, (1 - F(M))\, F^{N-1}(M)\, M \\ &\quad + \int_M^{\bar{v}} N\, x\, (1 - F(x))\, f_1^{(N)}(x)\mathrm{d}x \\ &= F^N(M)\, M + N\, (1 - F(M))\, F^{N-1}(M)\, M + \\ &\quad E\left[Y^{(N)} | Y^{(N)} \in [M, \bar{v}] \right] \end{aligned}$$

□

The expected closing price corresponds to that of an auction with reserve price M plus an extra payment of $F^N(M) \cdot M$. A PSA auctioneer can, then, use M to "simulate" a higher reservation price compared to the actual reservation price, r. Also, the expression for the closing price is a monotonically increasing function of N. Therefore, if more bidders are attracted to PSA because of profit sharing, the closing price, and hence the seller's profit, will increase.

Proposition 7 (Asymptotic closing price) *The closing price of PSA tends to $\overline{v}$ in probability.*

Proof: It is enough to show that the second highest bid tends to $\overline{v}$ in probability. More precisely, we have to show that $\forall \varepsilon > 0,\ P\left[Y^{(N)} \in [\overline{v} - \varepsilon, \overline{v}]\right] \to 1$. It suffices to show that P_N, the probability that at least two bid is in $[\overline{v} - \varepsilon, \overline{v}]$ tends to 1. Now, $P_N = 1 - \big((F(\overline{v}) - F(\overline{v} - \varepsilon))^N + (F(\overline{v}) - F(\overline{v} - \varepsilon))^{N-1}\big)$. □

Simulation environment

We needed to evaluate our hypothesis that sellers may have profit incentives of offering PSA auctions over FPSB or Vickrey auctions. To do this we chose an environment with N bidders and three auctions: an FPSB, a Vickrey auction and a PSA. The three auctions repeatedly offer the same item for sale with identical seller reservation prices. In each period, each bidder chooses a valuation from a uniform distribution in the range $[\underline{v}, \overline{v}]$ for the item being offered, and, based on this valuation, chooses to bid in the auction which is expected to return most profit. Agents bid their dominant strategy in Vickrey and bid their equilibrium strategy in PSA (we ensured that the condition $F(M)^N \leq P$ is always satisfied) and FPSB auction. We assume that the valuation distribution is common knowledge while the number of agents participating in an auction is unknown since bids are sealed. Thus, bidders in FPSB have to estimate this number. We provide details of the estimation process later.

Bidders use *Q-learning* (Watkins & Dayan 1992) to estimate the utility of attending different auctions given their valuation. The state is the agent's valuation for the item. As valuations are continuous, we discretized the valuation range in K isometric intervals $I_k = [V^{k-1}, V^k]$ where $\underline{v} = V^0 < V^1 < \cdots < V^K = \overline{v}$. The action set is the set of auctions, $\mathcal{A}$. The profit, R, a bidder receives from the auction $a \in \mathcal{A}$ when its valuation was v is used to update the Q-estimate as follows: $Q(I_{x(v)}, a) \leftarrow (1-\alpha) Q(I_{x(v)}, a) + \alpha R$, where $0 \leq \alpha < 1$ is a learning parameter and $I_{x(v)}$ is such that $v \in I_{x(v)}$. When choosing an auction, a bidder randomly picks an auction with probability ε and with probability $(1-\varepsilon)$ it will choose the auction with the highest Q-estimate for that valuation. This ε-greedy exploration policy allows the bidder to continually search for more profitable auctions.

To estimate the number of bidders in the FPSB auction, a bidder assumes other bidders have formed the same Q-estimates as it has. On average, $\frac{\varepsilon}{3} \cdot N$ bidders will choose FPSB by exploration. If $Q(I_k, F) \geq \max\{Q(I_k, V), Q(I_k, P)\}$ (where F V, and P denotes FPSB, Vickrey and PSA respectively) then the bidder estimates $(1-\varepsilon) \cdot P[v \in I_k] \cdot N$ bidders with valuation in I_k will come to FPSB by choice. Hence the number of agents in the FPSB auction is estimated as

$$\left(\frac{\varepsilon}{3} + (1-\varepsilon) \sum_{k=1}^{K} \delta\big(Q_k^F \geq \max\{Q_k^V, Q_k^P\}\big)\, P[v \in I_k] \right) N$$

where $Q_k^A = Q(I_k, A)$ and $\delta(cond) = 1$ if $cond = true$ and 0 otherwise.

Experimental Results

We run simulations with following parameter values: $N = 20$ or 100, $\underline{v} = 0$, $\overline{v} = 100$, $K = 10$, $r = 10$, $\alpha = 0.3$, $\varepsilon = .2$, $M = 25$ or 75.

Auction attendance and valuations

Bidder percentages in the three auctions are plotted for N=100, M=25 in Figure 2. In these and other experiments, PSA consistently attracts significantly more bidders than the other two auctions. This was expected since an agent always has a strictly positive utility for going to PSA unlike other auctions. Note that with $\varepsilon = 0.2$ we expect on an average 6.67% percent of bidders in Vickrey and FPSB auctions due to exploration only. As only a small number of bidders learn to go to Vickrey and FPSB, winners in these auctions, and particularly those with high valuations, are likely to win with a high margin. Consequently, certain agents learn that with a high valuation it is profitable to go Vickrey or FPSB. But if more high-valuation bidders are drawn to Vickrey/FPSB, PSA becomes more profitable for the remaining high-valuation bidders. On closer investigation, we find that any given bidder will periodically prefer each of the three auctions when it has high valuations.

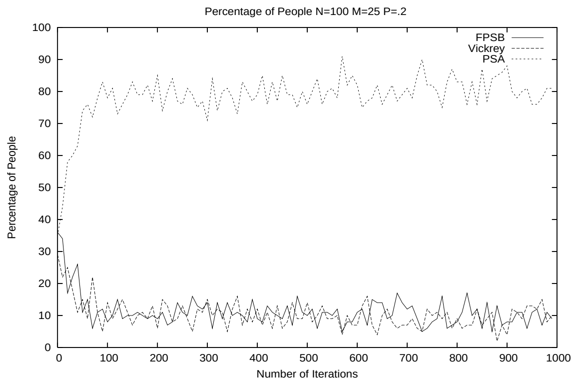

Figure 2: Percentage of bidders in each of the auctions.

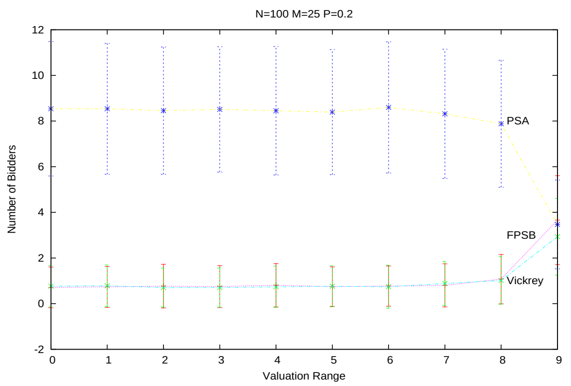

Figure 3: Average and standard deviation of the number of bidders in auctions for each I_k (N=100, M=25).

To get an accurate measure of the preferred choice of the bidders in each I_k, we plot their numbers in each of the auctions over a run (see Figure 3). For another perspective, we plot the valuation of the bidders in each of the three auctions over the course of a run in Figure 4. The plots clearly show that all bidders, except those with the highest valuation, prefer PSA over the other auctions after only a few iterations. In the highest valuation range, there are slightly more number of bidders in PSA and FPSB over Vickrey.

Seller's profit and selling prices

Now, we turn to the important metric of seller's profit. In Table we present the average cumulative profit obtained by the sellers of each of the three auctions in different configurations. The cumulative profit is calculated over 1000 iterations. PSA provides better profits to the seller for reasonably low values of M and small values of N. If we increase M to 75, the seller profit in PSA drops below that of the other auctions as the seller is giving away too much of its profit. For $N = 100$ and $M = 25$, PSA and Vickrey provide almost equal profit to the seller, and this is more than the profit from FPSB. When increasing M the profit provided obviously decreases since the seller shares more. However, profits provided by other auctions follow the same trend. This is due

N	M	FPSB	Vickrey	PSA
20	25	53406.7	52644.6	63266.8
20	50	52044.8	49835.2	59576.5
20	75	51136.3	47928.8	58305
100	25	73889.5	81317.9	81475.1
100	50	73518.7	80350.4	76847.3
100	75	73371.5	79714.4	72111.6

Table 1: Average seller profits ($P = .2$).

to the fact that PSA becomes more appealing leaving less incentive for bidders to go to Vickrey of FPSB.

In Figure 5 we plot the selling prices over the course of a run when N=100, M=25. The sharp, periodic drops of the selling price in Vickrey stands out in contrast to the steady selling prices in the other two auctions. Thus, even though, the cumulative profit in PSA and Vickrey is approximately equal, PSA may be preferred by the seller because of a steady revenue stream.

We now analyze these periodic price drops in Vickrey. Apart from bidders choosing Vickrey while exploring, only a few high-valuation bidders prefer it over other auctions. Once in a while, all or most of these few high-valuation bidders may choose to explore other auctions. If one or less high-valuation bidders remained in Vickrey, and no other high-valuation bidder preferring another auction came to Vickrey by exploration, or vice versa, the selling price in Vickrey will drop. These drops were more frequent and more pronounced in the plots for N=20 (omitted for space constraints). These periodic drops, due to exploration necessary for learning, accounted for most of the seller profit difference between Vickrey and PSA for those configurations. For N=20, there are similar drops in the selling price for FPSB and these drops are significantly less pronounced for N=100. This is because as N increases, so does the baseline and estimated attendance, x, in FPSB. With higher x, the slope of the curve $\frac{x-1}{x}$, the fraction of the true valuation bid by bidders, decreases, and hence the bid variability is smaller, i.e., selling price drops are not as sharp.

Conclusion

We proposed a new auction, Profit Sharing Auction, motivated by the goal of social welfare distribution. A criticism of such an auction is the lack of seller incentive for sharing its profit. We hypothesized that the sharing of profit will attract more bidders, thereby increasing the selling price and compensating the seller for the loss from sharing profits. We derived theoretical properties of PSA, including the fact that it is efficient and individually rational. Though low-valuation bidders have incentive to overbid, we argued that this is not going to change the auction outcome if the profit sharing threshold, M, is chosen such that at least two bidders have higher valuations. We also showed that the closing price, and hence the seller's profit, is a monotonically increasing function of the number of bidders in the auction.

To evaluate the viability of PSA, particularly from the seller's perspective, we allowed a fixed population of bid-

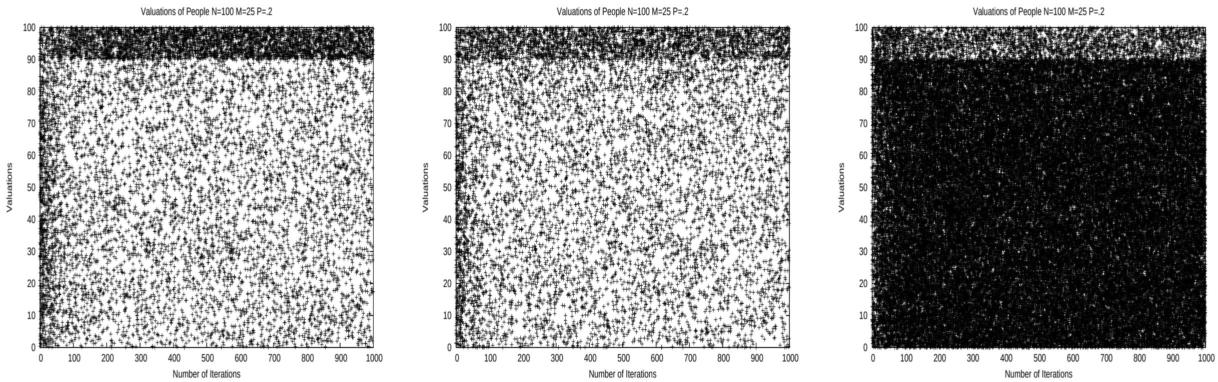

Figure 4: Bidder valuation distribution in the FPSB (left), Vickrey (middle), and PSA (right) over a run (N=100, M=25).

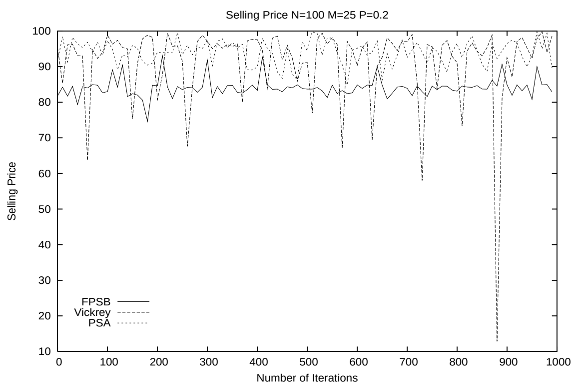

Figure 5: Selling Price of items in the three auctions.

ders to repeatedly choose between a PSA, Vickrey, and First Price Sealed Bid auction based on expected profits that is learned over repeated periods. Results show that PSA always attract more bidders (possible revenue source not discussed here) and, for reasonable settings of M, generates equal or more seller profits compared to the other auctions. Additionally, there is less variability in selling price with PSA. So, sellers have the incentive to offer PSA, which distributes welfare among a large number of bidders.

We need to evaluate the effect of many auctions of each type in an environment on seller profits. In our experiments, multiple bidders are learning simultaneously, and this possibly causes the high-valuation bidders to cycle between the auctions. This phenomenon needs to be studied in depth. We also need to derive formal guidelines for choosing PSA parameters given environmental configurations.

Acknowledgments: This work has been supported in part by an NSF award IIS-0209208.

References

Eichberger, J. 1993. *Game Theory for Economists*. Academic Press, Inc.

Greenwald, A., and Boyan, J. 2005. Bidding algorithms for simultaneous auctions: A case study. *AAMAS Journal* 10(1):67–89.

Jackson, M. 2000. Mechanism theory. In *Encyclopedia of Life Support Systems*.

Kraus, S. 2001. *Strategic Negotiation in Multi-Agent Environments*. MIT Press.

Myerson, R. B. 1981. Optimal auction design. *Mathematical Operations Research* 6:58–73.

Parkes, D. C. 2001. *Iterative Combinatorial Auctions: Achieving Economic and Computational Efficiency*. Ph.D. Dissertation, University of Pennsylvania.

Sandholm, T. 2000. emediator: a next generation electronic commerce server. In *Proc. AGENTS'00*, 341–348.

Stone, P.; Schapire, R. E.; Littman, M. L.; Csirik, J. A.; and McAllester, D. 2003. Decision-theoretic bidding based on learned density models in simultaneous, interacting auctions. *JAIR* 19:209–242.

Suyama, T., and Yokoo, M. 2004. Strategy/false-name proof protocols for combinatorial multi-attribute procurement auction. In *Proc. AAMAS'04*, 160–167.

Vickrey, W. 1961. Counterspeculation, Auctions and Competitive Sealed Tenders. *Journal of Finance* 8–37.

Watkins, C. J. C. H., and Dayan, P. D. 1992. Q-learning. *Machine Learning* 3:279 – 292.

Wurman, P. R.; Wellman, M. P.; and Walsh, W. E. 1998. The Michigan Internet AuctionBot: A configurable auction server for human and software agents. In *Proc. AGENTS'98*, 301–308.