Strategic Free Information Disclosure for a Vickrey Auction

Shani Alkoby[1] and David Sarne[2]

Bar-Ilan University, Ramat-Gan, 52900 Israel

Abstract. In many auction settings we find a self-interested information broker, that can potentially disambiguate the uncertainty associated with the common value of the auctioned item (e.g., the true condition of an auctioned car, the sales forecast for a company offered for sale). This paper extends prior work, that has considered mostly the information pricing question in this archetypal three-ply bidders-auctioneer-information broker model, by enabling the information broker a richer strategic behavior in the form of anonymously eliminating some of the uncertainty associated with the common value, for free. The analysis of the augmented model enables illustrating two somehow non-intuitive phenomena in such settings: (a) the information broker indeed may benefit from disclosing for free some of the information she wishes to sell, even though this seemingly reduces the uncertainty her service aims to disambiguate; and (b) the information broker may benefit from publishing the free information to the general public rather than just to the auctioneer, hence preventing the edge from the latter, even if she is the only prospective customer of the service. While the extraction of the information broker's optimal strategy is computationally hard, we propose two heuristics that rely on the variance between the different values, as means for generating potential solutions that are highly efficient. The importance of the results is primarily in providing information brokers with a new paradigm for improving their expected profit in auction settings. The new paradigm is also demonstrated to result, in some cases, in a greater social welfare, hence can be of much interest to market designers as well.

1 Introduction

Information disclosure is a key strategic choice in auctions and as such vastly researched both theoretically and empirically [8, 11]. One of the main questions in this context is the choice of the auctioneer to disclose information related to the common value of the auctioned item [10, 24, 12, 19, 20, 4]. For example, the board of a firm offered for sale can choose the extent to which the firm's client list or its sales forecast will be disclosed to prospective buyers. Various other examples are given in the literature cited throughout this paper. The disclosed information affects bidders' valuation of the auctioned item and consequently the winner determination and the auctioneer's profit.

In many cases, the information is initially not available to the auctioneer herself, but rather needs to be purchased by her from an external information broker. This is typically the case whenever generating the information requires some specific expertise or special equipment that the auctioneer does not possess. For example, in the firm selling case the information may pertain to the financial stability of key clients of the firm, hence typically offered for sale in the form of business analysts' reports. The auctioneer thus needs to decide both whether to purchase the information and whether to disclose it to bidders whenever purchased. The problem further complicates when the information broker herself is acting strategically, e.g., setting the price of the information offered in a way that maximizes her profit.

Prior work that dealt with uncertain auction settings with a self-interested information broker [33] allowed the information broker to control only the pricing of the information offered for sale. In this paper we extend the modeling of the information broker's strategy, enabling her also to disclose for free some of the information she holds. Specifically, we allow her to publicly eliminate some of the possible outcomes, narrowing the set of possible values that the common value may obtain. For example, prior to offering a firm the purchase of a market prediction report, the analyst can publicly publish its preliminary version that eliminates some of the possible outcomes. This behavior might seem intuitively non-beneficial, because now the information service disambiguates between less values, hence seemingly "worth" less. Nevertheless, our analysis of the augmented model enables demonstrating, numerically, that this choice can be sometimes beneficial. A second somehow surprising choice that we manage to illustrate is the one where the information broker finds it more beneficial to disclose the free information to both the auctioneer and the bidders rather than to the auctioneer only. The latter choice strengthens the auctioneer in the adversarial auctioneer-bidders interaction, allowing her to make a better use of the information offered for sale, if purchased, hence potentially enabling charging more for the service.

As explained in more details in the following paragraphs, the information brokers' problem of deciding what information to disclose for free is computationally extensive. Therefore another contribution of the paper is in presenting and demonstrating the effectiveness of two heuristics for ordering the exponential number of solutions that need to be evaluated, such that those associated with the highest profit will appear first in the ordering.

In the following section we provide a formal model presentation. Then, we present an equilibrium analysis for the case where the free information is disclosed to both the auctioneer and the bidders and illustrate the potential profit for the information broker from revealing some information for free, as well as the ordering heuristics and their evaluation. Next, we present the analysis of the case where the free information is disclosed only to the auctioneer. Finally we conclude with review of related work and discussion of the main findings.

2 The Model

Our basic auction model considers an auctioneer offering a single item for sale to n bidders using a second-price sealed-bid auction (with random winner selection in case of a tie). The auctioned item is assumed to be characterized by some value X (the "common value"), which is a priori unknown to both the auctioneer and the bidders [13, 14]. The only information publicly available with regard to X is the set of possible values it can obtain, denoted $X^* = \{x_1, ..., x_k\}$, and the probability associated with each value, $Pr(X = x)$ ($\sum_{x \in X^*} Pr(X = x) = 1$). Bidders are assumed to be heterogeneous in the sense that each is associated with a type T that defines her valuation of the auctioned item (i.e., her "private value") for any possible value that X may obtain. We use the function $V_t(x)$ to denote the private value of a bidder of type $T = t$ in case the true value of the item is $X = x$. It is assumed that the probability function of types, denoted $Pr(T = t)$, is publicly known, however a bidder's specific type is known only to herself.

The model assumes the auctioneer can obtain the value of X from an outer source, denoted "information broker" (for the rest of the paper will be called "broker"), by paying a fee C set by the broker. Similar to prior models (e.g., [33]), and for the same justifications given there, it is assumed that this option of purchasing the information is available only to the auctioneer, though the bidders are aware of this possibility.

If purchasing the information, the auctioneer, based on the value obtained, can decide either to disclose this information to the bidders or keep it to herself (hence disclosing $\emptyset$). If disclosing the information, then it is assumed that the information received from the broker is disclosed as is (i.e., truthfully and symmetrically to all bidders), e.g., in case the auctioneer is regulated or has to consider her reputation. Finally, it is assumed that all players (auctioneer, bidders and the broker) are self-interested, risk-neutral and fully rational agents, and acquainted with the general setting parameters: the number of bidders in the auction, n, the cost of purchasing the information, C, the discrete random variables X and T, their possible values and their probability functions.

Up to this point our model resembles those found in prior literature. For example, it generalizes the one found in [10, 24] in the sense that it requires the auctioneer to decide on purchasing the external information rather than assuming she initially possesses it. It is also equivalent to the one found in [33] where the broker is self-interested agent that controls C, the price of purchasing the information. Our model, however, extends prior work in the sense that it allows the broker also to anonymously publish some of the information for free before the auctioneer makes her decision of whether to purchase the information. The anonymity requirement in this case is important as discussed later on in the analysis section. Yet, there are numerous options nowadays for publishing such information anonymously, e.g., through an anonymous email, uploading the information to an electronic bulletin board or anonymous file server, sending the information to a journalist or an analyst. The typical case, which we use for our analysis, is the one where the broker, knowing the true value $x \in X^*$, eliminates a subset of values $D \subset X^*$ (where $x \notin D$), leaving only the values $X^* - D$ as applicable values the common value may obtain. Doing so, our model distinguishes between the case where the free information is disclosed to all and the one where it is disclosed to the auctioneer only (allowing the latter to decide what parts of it to disclose further to the bidders prior to starting the auction).

3 Disclosing Information for Free

Consider the case where the true common value is x. In this case, if the broker publicly eliminates (i.e., anonymously publishes that the common value is not part of) the subset $D \subset X^*$ then the auctioneer and bidders are now facing the problem where the common value may receive only the subset $X^* - D$ and the a priori probability of each value in the new setting is given by $Pr'(X = x) = \frac{Pr(X=x)}{\sum_{x_i \in X^*-D} Pr(X=x_i)}$. Since the auctioneer needs to decide both whether to purchase the true value $x \in X^* - D$ and if so whether to disclose it to the bidders, her (mixed) strategy can be characterized using $R^{auc} = (p^a, p^a_1, ..., p^a_k)$ where p^a is the probability she purchases the information from the broker and p^a_i ($1 \leq i \leq k$) is the probability she discloses to the bidders the value x_i if indeed $X = x_i$. The dominating bid of a bidder of type t, when the auctioneer discloses that the true value is x, denoted $B(t, x)$, is given by $B(t, x) = V_t(x)$ [36]. If no information is disclosed ($x = \emptyset$) then the dominating strategy for each bidder is to bid her expected private value, based on her belief of whether information was indeed purchased and if so, whether the value received is intentionally not disclosed by the auctioneer [10]. The bidders' strategy, denoted R^{bidder}, can thus be compactly represented as $R^{bidder} = (p^b, p^b_1, ..., p^b_k)$, where p^b is the probability they assign to information purchase by the auctioneer and p^b_i is the probability they assign to the event that the information is indeed disclosed if purchased by the auctioneer and turned to be x_i.[1]

The bid placed by a bidder of type t in case the auctioneer does not disclose any value, $B(t, \emptyset)$, is therefore:

$$B(t, \emptyset) = \sum_x V_t(x) \cdot Pr^*(X = x) \tag{1}$$

where $Pr^*(X = x)$ is the posterior probability of x_i being the true common value, based on the bidders' belief R^{bidder} and is being calculated as:

[1] Being rational, all bidders hold the same belief in equilibrium.

$$Pr^*(X = x_i) = \frac{Pr(X = x_i)(p^b(1 - p_i^b) + (1 - p^b))}{(1 - p^b) + p^b \sum(1 - p_i^b)Pr(X = x_i)} \tag{2}$$

The term in the numerator is the probability that x_i indeed will be the true value and will not be disclosed. If indeed x_i is the true value (i.e., with a probability of $Pr(X = x_i)$) then it will not be disclosed either if the information is not purchased (i.e., with a probability of $(1 - p^b)$) or if purchased but not disclosed (i.e., with a probability of $p^b(1 - p_i^b)$). The term in the denominator is the overall probability that the information will not be disclosed. This can happen either if the information will not be purchased (i.e., with a probability of $(1 - p^b)$) or when the information will be purchased however the value will not be disclosed (i.e., with probability of $p^b \sum(1 - p_i^b)Pr(X = x_i)$).

Consequently, the auctioneer's expected profit when using R^{auc} while the bidders use R^{bidder}, denoted $EB(R^{auc}, R^{bidder})$, is given by:

$$\begin{aligned} EB(R^{auc}, R^{bidder}) &= p^a \sum Pr'(X = x_i)p_i^a \cdot ER_{auc}(x_i) \\ &+ ((1 - p^a) + p^a \sum(1 - p_i^a)Pr'(X = x_i)) \cdot ER_{auc}(\emptyset) - p^a \cdot C \end{aligned} \tag{3}$$

where $ER_{auc}(x_i)$ is the expected second highest bid if disclosing the true value x_i ($x_i \in \{X^* - D, \emptyset\}$). The broker's expected profit is $p^a \cdot C$. The first row of the equation deals with the case where the auctioneer discloses the true value to the bidders (i.e., p^a is the probability that the information was purchased and $\sum Pr'(X = x_i)p_i^a \cdot ER_{auc}(x_i)$ is the probability that x_i is the true value multiplied by the auctioneer's expected profit for this case). The second row deals with the case where the information was not disclosed to the bidders (i.e., when the information is not purchased by the auctioneer (with probability $(1 - p^a)$) and when the information is purchased but not discloses (with probability $p^a \sum(1 - p_i^a)Pr'(X = x_i)$)).

A stable solution in this case (for the exact same proof given in [33]) is necessarily of the form $R^{auc} = R^{bidder} = R = (p, p_1, ..., p_k)$ (as otherwise, if $R^{auc} = R' \neq R^{bidder}$, the bidders necessarily have an incentive to deviate to $R^{bidder} = R'$), such that [33]: (a) for any $0 < p_i < 1$ (or $0 < p < 1$): $ER_{auc}(\emptyset, R) = ER_{auc}(X_i)$ (or $ER_{auc}(\emptyset, R^{bidder}) = ER_{auc}((1, p_1, ..., p_k), R^{bidder})$); (b) for any $p_i = 0$ (or $p = 0$): $ER_{auc}(\emptyset, R^{bidder}) \geq ER_{auc}(X_i)$ (or $ER_{auc}(\emptyset, R^{bidder}) \geq ER_{auc}((1, p_1, ..., p_k), R^{bidder})$; and (c) for any $p_i = 1$ (or $p = 1$): $ER_{auc}(\emptyset, R^{bidder}) \leq ER_{auc}(X_i)$ (or $ER_{auc}(\emptyset, R^{bidder}) \leq ER_{auc}((1, p_1, ..., p_k), R^{bidder})$. Therefore one needs to evaluate all the possible solutions of the form $(p, p_1, ..., p_k)$ that may hold (where each probability is either assigned 1, 0 or a value in-between). Each mixed solution of these $2 \cdot 3^k$ combinations (as only one solution where $p = 0$ is applicable) should be first solved for the appropriate probabilities according to the above stability conditions. Since the auctioneer is the first mover in this model (deciding on information purchase), the equilibrium used is the stable solution for which the auctioneer's expected profit is maximized.

If the information is provided for free ($C = 0$) then information is necessarily obtained and the resulting equilibrium is equivalent to the one given in [10] for the pure equilibrium case and [24] for the mixed equilibrium case.

Being able to extract the equilibrium for each price C she sets, the broker can now find the price C which maximizes her expected profit. Repeating the process for all different sets $D \subset X^*$, enables extracting the broker's expected-profit maximizing strategy (D, C).

Figure 1 depicts the expected profit of the auctioneer (vertical axis) as a function of the information cost C (horizontal axis), for five of the possible D sets. The setting used is given in the table at the bottom of the figure. It is based on four possible values the common value may obtain: $X^* = \{x_1, x_2, x_3, x_4\}$, where x_3 is the true value. The subset D that is used for each curve is marked next to it. For each set D the information provider discloses, the auctioneer chooses whether to purchase the information and what values to disclose, if purchasing, according to the auctioneer's expected-profit-maximizing equilibrium. For example, the lowest curve depicts the auctioneer's expected profit when the broker initially eliminates the values $\{x_1, x_4\}$ and the auctioneer's strategy is to disclose to the bidders the value x_2 in case it is the true value of the auctioned item. Since equilibria in this example are all based on pure strategies, the expected-profit-maximizing price C, and hence the expected profit, equals the highest price at which information is still purchased (marked by circles in the graph, as in this specific example the last segment of each curve applies to an equilibrium by which the information is not being purchased at all). From the figure we see that indeed in this sample setting, anonymously eliminating some of the applicable values is highly beneficial - for example, the elimination of x_1 results in a profit of 3.7, compared to a profit of 1.2 in the case no information is being a priori eliminated (i.e., $D = \emptyset$).

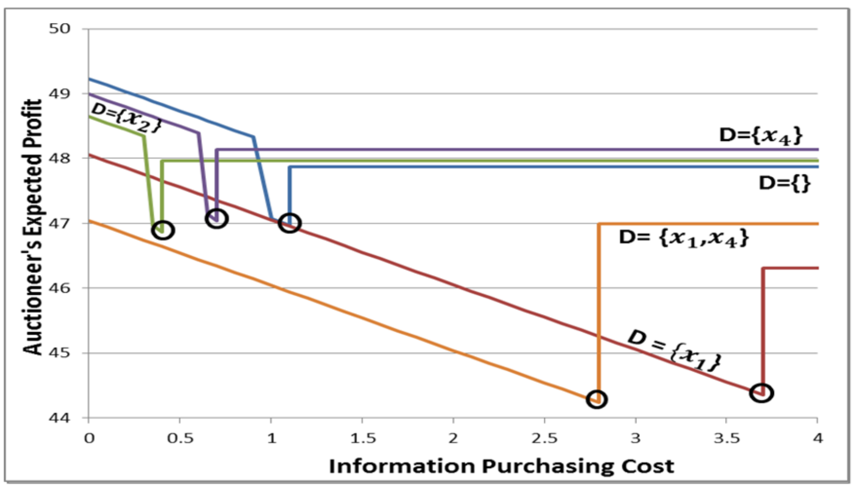

n=3		private values			
		x1	x2	x3	x4
Type	p(type) \ p(value)	0.626	0.159	0.143	0.072
1	0.252	54	25	84	10
2	0.33	79	56	0	53
3	0.386	21	59	46	69
4	0.032	25	94	69	93

Fig. 1. Auctioneer's expected profit as function of information purchasing cost, for different a priori eliminated subsets.

As discussed in the introduction, benefiting from providing some of the information for free may seem non-intuitive at first—seemingly the broker is giving away some of her ability to disambiguate the auctioneer's and bidders' uncertainty. Yet, since the choice of whether the information is purchased or not at any specific price derives from equilibrium considerations, rather than merely the auctioneer's preference, it is possible that providing information for free becomes a preferable choice for the broker.

The benefit in free information disclosure does not necessarily come at the expense of social welfare. For exemplifying this we introduce Figure 2. The setting used for this example

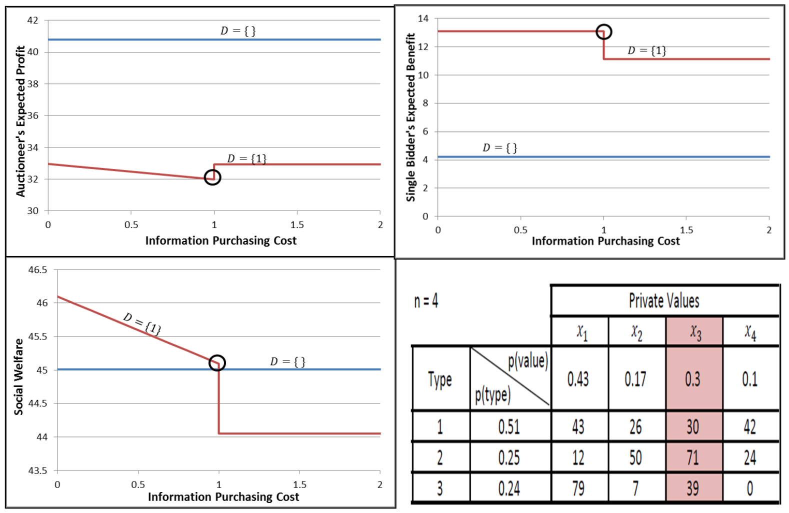

n = 4

		Private Values			
		x_1	x_2	x_3	x_4
Type	p(value) / p(type)	0.43	0.17	0.3	0.1
1	0.51	43	26	30	42
2	0.25	12	50	71	24
3	0.24	79	7	39	0

Fig. 2. An example of an improvement both in the broker's expected profit and the the social welfare as a result of free information disclosure. The true common value of the auctioned item in this example is x_3.

is given in the bottom right side of the figure. Again, the auctioneer's strategy is to disclose the set which will benefit her the most. In this example the broker's expected profit increases from 0 to 1 by publicly eliminating the value x_1 (the information is not purchased otherwise), and at the same time the social welfare (sum of the bidders' and auctioneer's profit) increases from 45 to 45.2, due to the substantial increase in the bidder's profit (from 4.2 to 13.1). If including the broker's expected profit in the social welfare calculation, the increase is even greater.

Finally, we note the importance of disclosing the information anonymously or without leaving a trace of a strategic behavior from the broker's side. If the auctioneer and bidders suspect that the broker may disclose free information strategically, then the equilibrium analysis should be extended to accommodate the probabilistic update resulting from their reasoning of the broker's strategy. This latter analysis is left beyond the scope of the current paper—as discussed in the previous section there are various ways nowadays for anonymous disclosure of information, justifying this specific modeling choice.

4 Sequencing Heuristics

The extraction of the broker's expected-profit-maximizing subset D is computationally exhausting due to the exponential number of subsets for which equilibria need to be calculated — the broker needs to iterate over all possible $2^{|X^*|-1} - 1$ D subsets (as there are $|X^*| - 1$ values that can be eliminated, and eliminating all but the true value necessarily unfolds the latter as the true one). Therefore, in this section we present two efficient heuristics—Variance-based (Vb) and Second-Price-Variance-based ($SPVb$)—that enable the broker to predict with much success what subsets D are likely to result, if eliminated for free, with close to optimal expected profit. The heuristics can be considered sequencing heuristics, as they aim to determine the order according to which the different subsets should be evaluated. The idea is to evaluate early in the process those subsets that are likely to be associated with the greatest expected profit. This way a highly favorable solution will be obtained regardless of how many subsets can be evaluated in total.

Variance-based (Vb) The value of the information supplied by the broker derives from the different players' (auctioneer and bidders) ability to distinguish the true common value from others, i.e., to better identify the worth of the auctioned item to different bidders. Therefore this heuristic relies on the variance between the possible private values that the information purchased will disambiguate as the primary indicator for its worth. Specifically, if the broker a priori eliminates the subset D, we first update the probabilities of the remaining applicable values, i.e., $Pr^*(x \in X^* - D) = \frac{Pr(X=x)}{\sum_{y \in X^*-D} Pr(X=y)}$. The revised probabilities are then used for calculating the variance of the private values in the bidder's type level, denoted $Var(T = t)$: $Var(T = t) = \sum_{x \in X^*-D} Pr^*(x)(V_t(x) - B(t, \emptyset))^2$, where $V_t(x)$ is the private value of a bidder of type $T = t$ if knowing that the true common value is x, as defined in the model section, and $B(t, \emptyset)$ is calculated according to (1), based on a setting $X^* - D$. The overall weighted variance is calculated as the weighted sum of the variance in the bidder's type level, using the type probabilities as weights, i.e., $\sum_{t \in T} Pr(T = t) \cdot Var(T = t)$. The order according to which the different subsets $D \subset X^*$ should be evaluated is thus based on the overall weighted variance, descending.

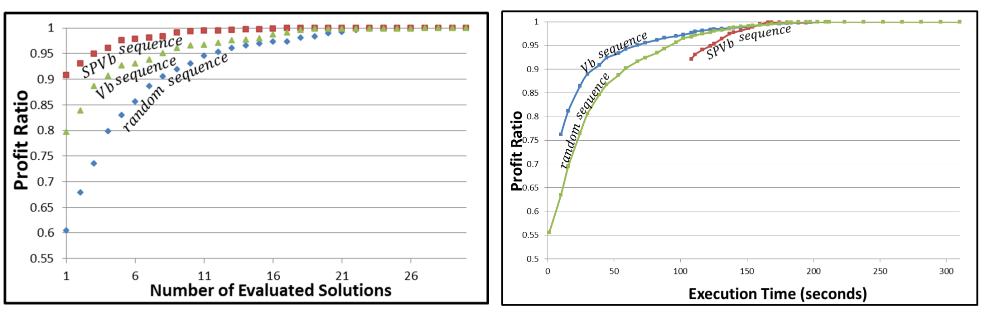

Fig. 3. Performance (ratio between achieved expected profit and maximal expected profit): (a) Vb and $SPVb$ versus random ordering; and (b) all three methods as a function of running time. All data points are the average over 2500 random settings with 6 possible values the common value obtains.

Figure 3(a) illustrates the performance of Vb (middle curve) as a function of the number of evaluated free disclosed subsets (horizontal axis). Since the settings that were used for producing the graph highly varied, as detailed below, we had to use a normalized measure of performance. Therefore we used the ratio between the broker's expected profit if following the sequence generated by the heuristic and the expected profit achieved with the profit-maximizing subset (i.e., how close we manage to get to the result of brute force) as the primary performance measure in our evaluation. The graph depicts also the performance of random ordering as a baseline. The set of problems used for this graph contains 2500 randomly generated settings where the common value may obtain six possible values, each assigned with a random probability, normalized such that all probabilities sum to 1. Similarly, the number of bidders and the number of bidder types in each setting were randomly set within the ranges (2-10) and

(2-6), respectively. Finally, the probability assigned to each bidder type was generated in the same manner as with the common value probabilities. For each setting we randomly picked one of the values the common value may obtain, according to the common-value probability function. Each data point in the figure thus represents the average performance over the 2500 randomly generated settings.

As can be seen from the graph, Vb dominates the random sequencing in the sense that it produces substantially better results for any number of subsets being evaluated. In particular, the improvement in performance with the heuristic is most notable for relatively small number of evaluated solutions, which is the primary desirable property for such a sequencing method, as the goal is to identify highly favorable solutions within a limited number of evaluations. As expected, the performance of both Vb and random ordering monotonically increase, converging to 1 (and necessarily reaching 1 once all possible solutions have been evaluated). This is because as the number of evaluated subsets increases the process becomes closer to brute force

Second-Price-Variance-based ($SPVb$) This heuristic is similar to Vb in the sense that it orders the different subsets according to their weighted variance, descending. It differ from Vb in the sense that instead of depending on the variance in bidders' private values it uses the variance in the worth of information to the auctioneer, i.e., in the expected second price bids. The variance of the expected second price bids if disclosing D for free, denoted $Var(D)$, is calculated as: $Var(D) = \sum_{x \in X^* - D} Pr^*(x)(ER_{auc}(x) - ER_{auc}(\emptyset|D))^2$, where $Pr^*(x)$ is calculated as in Vb, $ER_{auc}(x)$ is the expected second highest bid if disclosing to the bidders that the true value is x, as given in the former section. $ER_{auc}(\emptyset|D)$ is the expected second highest bid if the auctioneer discloses no information to the bidders however the bidders are aware of the elimination of the subset D by the broker, i.e., bid according to $B(t, \emptyset) = \sum_{x \in X^* - D} V_t(x) Pr(X = x) / \sum_{x \in X^* - D} Pr(X = x)$.

Figure 3(a) also illustrates the performance of $SPVb$ (upper curve) as a function of the number of evaluated subsets D using a similar evaluation methodology and the same 2500 settings that were used for evaluating Vb, as described above. As can be seen from the graph, $SPVb$ dominates random sequencing and produces a substantial improvement, especially when the number of evaluated subsets is small. In fact, comparing the two upper curves in Figure 3(a) we observe that $SPVb$ dominates Vb in terms of performance as a function of the number of evaluated sets. One impressive finding related to $SPVb$ is that even if choosing the first subset in the sequence it produces a relatively high performance can be obtained—91% of the maximum possible expected profit, on average. This means that even without evaluating any of the subsets (e.g., in case the broker is incapable of carrying the equilibrium analysis) but merely by extracting the sets ordering, the broker can come up with a relatively effective subset of values to disclose for free.

This dominance of $SPVb$ is explained by the fact that it relies on the variance between the winning bids rather than the bidders' private values. Meaning it relates to the true worth of the information to the auctioneer and consequently to the broker's profit. While this is $SPVb$'s main advantage, compared to Vb, it is also its main weakness: from the computational aspect, the time required for calculating the expected second-price variance of all applicable subsets D is substantially greater than the time required for Vb to calculate the variance between the possible private values. The expected profit of the auctioneer when disclosing the information $X = x$, denoted $ER_{auc}(X = x)$, equals the expected second-best bid when the bidders are given x, formally calculated as:

$$ER_{auc}(X=x) = \sum_{w \in \{B(t,x)|t \in T\}} w(\sum_{k=1}^{n-1} n \binom{n-1}{k}$$
$$\sum_{B(t,x)>w} Pr(T=t)(\sum_{B(t,x)=w} Pr(T=t))^k (\sum_{B(t,x)<w} Pr(T=t))^{n-k-1} \tag{4}$$
$$+ \sum_{k=2}^{n} \binom{n}{k} (\sum_{B(t,x)=w} Pr(T=t))^k (\sum_{B(t,x)<w} Pr(T=t))^{n-k})$$

The calculation iterates over all of the possible second-best bid values, assigning for each its probability of being the second-best bid. As we consider discrete probability functions, it is possible to have two bidders placing the same highest bid (in which case it is also the second-best bid). For any given bid value, w, we therefore consider the probability of having either: (i) one bidder bidding more than w, $k \in 1, ..., (n-1)$ bidders bidding exactly w and all of the other bidders bidding less than w; or (ii) $k \in 2, ..., n$ bidders bidding exactly w and all of the others bidding less than w. Notice that (4) also holds for the case where $x = \emptyset$ (in which case bidders use $B(t, \emptyset)$ according to (1)).

The mentioned calculation results in a combinatorial (in the number of values the common value may obtain) run time. The $SPVb$ method thus requires more time to run for producing the sequence according to which sets need to be evaluated, however the ordering it produces is substantially better than the one produced by Vb. Similarly, random sequencing does not require any "setup" time and the different subsets can be evaluated right away.

In order to weigh in this effect in the heuristics' evaluation we present Figure 3(b). Here, the performance is depicted as a function of the actual run-time (in seconds, over the horizontal axis) rather than the number of subsets evaluated once the ordering is completed.[2] Here, we can see the tradeoff between the initial calculation required for the ordering itself and the improvement achieved within the first few evaluated subsets. The shift of each curve over the horizontal axis, till its first data point, is the time it took to generate the sequence of subsets. From the graph we see that if the amount of time allowed for running is relatively small then one should choose to use a random sequence for evaluation. If the broker is less time-constrained, the best choice is to use Vb and then evaluate subsets according to the generated sequence. We notice that the same typical behavior was observed for the case of five and seven possible values that the common value may obtain. Evaluating for settings with more than six values is impractical, as it requires solving for thousands of such settings each, as seen from the Table 1, takes substantial time to solve.

Table 1 depicts the average time it took to extract the equilibrium solution for a setting according to the number of values in X^*. Each data point is the average for the 2500 problems described above. This justifies our use of six values settings in the numerical evaluation, and generally motivates the need for the sequencing heuristics we provide by showing that evaluating all possible sets is in many cases impractical — indeed in many cases the total number of values in X^* is moderate,[3] however, even with 8 values it takes more than 10 minutes to extract the broker's equilibrium profit for a single instance.

[2] Our evaluation framework was built in Matlab R2011b and run on top of Windows7 on a PC with Intel(R) Xeon(R) CPU E5620 (2 processors) with 24.0 GB RAM.

[3] For example, in oil drilling surveys, geologists usually specify 3-4 possible ranges for the amount of oil or gas that is likely to be found in a given area. Similarly, when requesting an estimate of the amount of traffic next to an advertising space, the answer would usually be in the form of ranges rather than exact numbers.

# of Possible Values	3	4	5	6	7	8
Execution Time (seconds)	0.16	0.58	3.57	20.07	103.19	708.46

Table 1. Average time in seconds for extracting the broker's equilibrium profit in a single setting as a function of $|X^*|$.

5 The Influence of Bidders' Awareness

Next we consider the case where instead of revealing the information for free to all, only the auctioneer receives it (e.g., using anonymous email). In this case the auctioneer needs to decide whether to reveal this information (or part of it) to the bidders. This complicates a bit the structure of the game: (a) First, the broker needs to decide on the set D of values to be eliminated for free and the price C of her service of disambiguating the remaining uncertainty; (b) then, she needs to transfer D anonymously to the auctioneer; (c) next, the auctioneer needs to decide what part $D' \subseteq D$ to further disclose to the bidders; (d) then, the auctioneer needs to decide whether to purchase the true value from the broker, and if purchasing, upon receiving the value, whether to disclose it to the bidders or leave them uncertain concerning the true value; (e) finally, the bidders need to bid for the auctioned item.

The analysis of this case relies heavily on the analysis given in the former sections. The resulting adversarial setting if using D and D' is one where bidders bid $V_t(x)$ whenever the information is purchased and disclosed by the auctioneer, and otherwise $B(t, \emptyset)$ according to (1), except that this time the probabilities $Pr^*(X = x_i)$ used by bidders result from the equilibrium of a setting where the original values are $X^* - D'$. Therefore, upon receiving the information D from the anonymous source, the auctioneer needs to calculate her expected profit from disclosing any subset $D' \subseteq D$ and choose the one that maximizes it. The auctioneer's expected profit calculation in this case is, however, a bit different, due to the asymmetry in information. When initially disclosing D' to bidders, the auctioneer needs to calculate the expected second best bid from disclosing any value $x \in X^* - D$, based on the bidders' type distribution and their bidding strategy as given above. The auctioneer should choose to disclose any value x for which the expected second best bid if disclosed is greater than the expected second best bid when no information is disclosed (i.e., when bidders bid $B(t, \emptyset)$ according to the equilibrium for the $X^* - D'$ instance of the original problem, as explained above). This allows the broker deciding what subset D to disclose, such that her expected profit is maximized.

Figure 4 is an example of a case where the information broker discloses the free information only to the auctioneer and it is to the auctioneer's choice which parts of the information (if at all) to disclose to the bidders prior to the start of the auction. It relies on a setting of three bidders, two possible types and four different values the common value may obtain ($x_1, ..., x_4$), out of which x_4 is the true common value. The full setting details are given in the table in the right hand side of the figure. The leaf nodes provide the expected profit of the auctioneer (inside the rectangle) and the broker (below the rectangle) for each combination of selections made by these two players (the subset D disclosed for free and the subset $D' \subseteq D$ disclosed to the bidders), according to the resulting equilibrium as analyzed above. The yellow colored leafs are therefore those corresponding to the auctioneer's best response given the subset D picked by the broker, hence the expected-profit maximizing strategy for the broker is to anonymously disclose to the auctioneer the subset $\{x_2, x_3\}$ as in this case the auctioneer will choose not to disclose any of these two values to the bidders, resulting in expected profit of 0.9 (compared to 0.8,0.6,0.6,0.8,0.4 and 0.4 if eliminating $\{\emptyset\}, \{x_1\}, \{x_2\}, \{x_3\}, \{x_1, x_2\}$ and $\{x_1, x_3\}$, respectively).

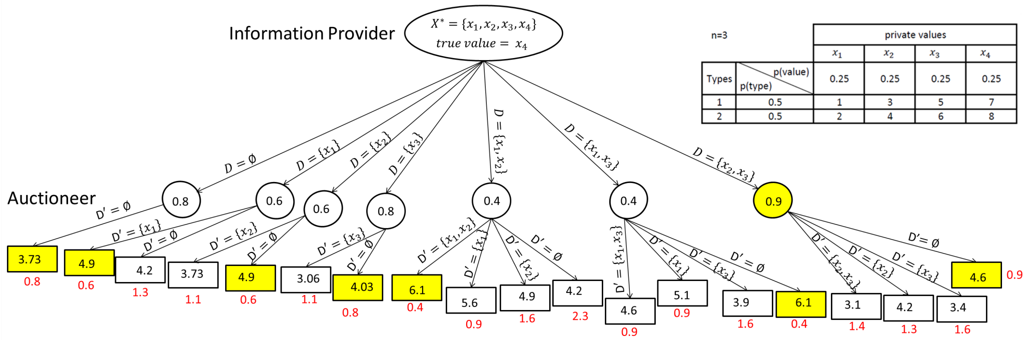

Fig. 4. Disclosing the free information to the auctioneer only: the broker needs to decide on the subset D to eliminate and then the auctioneer needs to decide on the subset $D' \subset D$ to disclose to bidders.

Interestingly, if the broker chooses to anonymously disclose to both the auctioneer and the bidders that x_2 and x_3 can be eliminated, her expected profit, calculated based on the analysis given in former sections, is 1.4. This is substantially greater than in the case where the bidders are unaware of the information that was disclosed for free. Furthermore, eliminating x_2 and x_3 for free is not necessarily the broker's expected-profit-maximizing strategy for the scenario where the free information reaches both the auctioneer and bidders. It is possible that there is another subset which elimination results in an even greater improvement in profit when compared to disclosing the elimination of x_2 and x_3 to the auctioneer only. This outcome, as discussed in the introduction is quite non-intuitive because by eliminating the asymmetry in the information disclosed to the different players the broker seemingly reduces the auctioneer's power against the bidders in this adversarial setting. Indeed, when the choice is given to the auctioneer she would rather not disclose this information to the bidders and increase her profit. Since the auctioneer is the potential purchaser of the broker's service information offered by the broker, it might seem that by disclosing the free information only to her, she will have a greater flexibility in making use of the remaining information (that is offered for sale) hence will see a greater value in purchasing it. Yet, the improvement in the auctioneer's competence by disclosing the free information to her only does not translate to an improvement in the broker's profit—eventually the broker's profit depends on the range of prices and the corresponding probabilities at which her information is indeed purchased. These latter factors result from the equilibria considerations, leading to behaviors such as in the example above.

Even for this case, the sequencing heuristics Vb and $SPVb$ are of much importance. Figure 5 presents the performance evaluation for these two heuristics, for settings with six values, demonstrating that highly efficient solutions can be extracted even with a small number of evaluations.

6 Related Work

Over the years auctions have focused much interest in research, mostly due to their advantage in effectively extracting bidders' valuations and the guarantee of many auction protocols

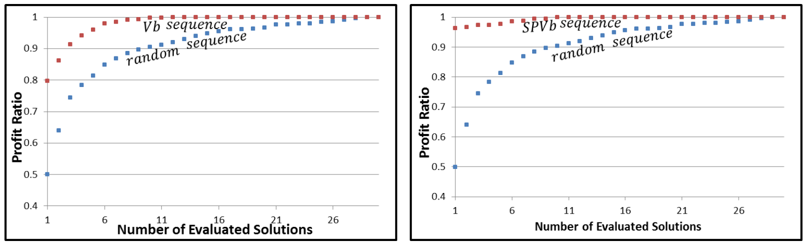

Fig. 5. Performance (ratio between achieved expected profit and maximal expected profit) when the information is disclosed for free only to the auctioneer and she chooses which information to disclose to the bidders : (a) Vb versus random ordering as a function of number of evaluated subsets; (b) $SPVb$ versus random ordering as a function of number of evaluated subsets. All data points are the average of 2500 random settings with 6 possible values the common value obtains.

to result in efficient allocation [21, 18, 7, 5, 35, 34]. The case where there is some uncertainty associated with the value of the auctioned item is quite common in auctions literature. Most commonly it is assumed that the value of the auctioned item is unknown to the bidders at the time of the auction and bidders may only have an estimate or some privately known signal, such as an expert's estimate, that is correlated with the true value [14, 22]. Many of the works using uncertain common value models assumed asymmetry in the knowledge available to the bidders and the auctioneer regarding the auctioned item, typically having sellers more informative than bidders [1, 10]. As such, much recent emphasis was placed on the role of information revelation [8, 11, 12, 19]. In particular, several works have considered the computational aspects of such models where the auctioneer needs to decide on the subsets of non-distinguishable values to be disclosed to the bidders [10, 24, 9]. Still, all these works assume the auctioneer necessarily obtains the information and that the division into non-distinguishable groups, whenever applicable, is always given to the bidders a priori. Our problem, on the other hand, does not require that the auctioneer possesses (or purchases) the information in the first place, and allows not disclosing any value even if the information is purchased. Recent work that does consider an auction setting with a strategic broker, and in fact provides the underlying three-ply equilibrium analysis for this case [33, 3], limits the strategic behavior of the broker to price-setting only. In this paper we extend that work to include an additional strategic dimension for the broker, in the sense of anonymously disclosing some of the information for free. Furthermore, unlike this prior work, in this paper we deal with the computational aspects of extracting the broker's strategy.

Models where agents can disambiguate the uncertainty associated with the opportunities they consider exploiting through the purchase of information have been studied in several other multi-agent domains, e.g., in optimal stopping domains [38, 29, 27, 30, 28]. Here, the main questions studied were how much costly information it makes sense to acquire before making a decision [25, 31], in particular when additional attributes can be revealed at certain costs along the search path [23, 37]. Relaxation of the perfect signals assumption has also been explored in models of economic search [6, 2]. Alas, mediators in such models usually take the form of matchmakers rather than information brokers. Those that do consider a self-interested information broker in these domains, e.g., Nahum et al.[26], focused on the way it should set the price for the information it provides and did not consider the option of free information disclosure.

Other related work can be found in the study of platforms that bring together different sides of the market (e.g., dating, or eCommerce platforms). Here, there is much work on the impact on selective information disclosure [15], strategic ordering of the disclosed information [16] and having the platform charging only one of the two participating sides [17] and even cases where consumers are in effect paid to use the platform were studied [32]. Our work can be viewed in a similar vein, especially in the context of the information broker subsidizing information provisioning, although the intuitions behind our results are quite different and grounded in the transition between different equilibria rather than in the profit of potentially increasing participation overall.

7 Conclusions and Future Work

Information brokers have become an integral part of many multi-agent systems. These range from individuals with specific expertise, offering their services for a fee (e.g., analysts), to large information services, such as `Carfax.com` or credit report companies. The model and analysis given in the paper adds an important strategic dimension to prior work in the form of influencing the auctioneer's and bidders' strategic interaction through the anonymous revelation of some of the information that is offered for sale. As discussed throughout the paper, this behavior may seem a bit unnatural. We show, however, that this strategy can actually be highly beneficial to the broker. In fact, as demonstrated in the paper, it can even lead to an overall improvement in the social welfare. Furthermore, if given the option to disclose the free information to both the bidders and the auctioneer or to the auctioneer only, the broker may benefit from choosing the first, despite the fact that the auctioneer is the one to decide about purchasing the information.

The paper presents two sequencing heuristics aiming to reduce the computation time of the broker's expected-profit maximizing strategy. The results of an extensive evaluation of these are quite encouraging - the generated sequences, with both heuristics, are quite effective, as the very few initial subsets placed first in the sequence offer expected profit very close to the expected-profit-maximizing one. Both methods use the variance as a measure for the profit in disclosing a given set, differing in the values based on which the variance is calculated—the bidder's private valuations and the expected second price bids. Interestingly, we find that while the use of the expected second-price produces a substantially more efficient sequence, it is better to rely on the raw values (i.e., bidders' valuations) as the execution time of generating the sequence using the latter method is substantially shorter, leading to better performance overall.

We note that, much like prior work, our model makes several assumptions that can be relaxed in future research. For example, one can think of settings where the information is provided to the bidders not just based on the auctioneer's decision to disclose it. Here, numerous variations can be considered. For example, the bidders can purchase the information, whether symmetrically or asymmetrically, either directly from the broker or indirectly from the auctioneer. These of course require extending the analysis to include all the different dynamics that will be formed. Another natural extension of our model would be one where the auctioneer and the bidders are aware to the fact that the broker is the one that disclosed the information for free (i.e., the free disclosure is not anonymous anymore) as discussed in the analysis section.

References

1. G. Akerlof. The market for "lemons": Quality uncertainty and the market mechanism. *The Quarterly Journal of Economics*, 84(3):488–500, 1970.
2. S. Alkoby, D. Sarne, and S. Das. Strategic free information disclosure for search-based information platforms. In *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems*, pages 635–643. International Foundation for Autonomous Agents and Multiagent Systems, 2015.
3. S. Alkoby, D. Sarne, and E. David. Manipulating information providers access to information in auctions. In *Technologies and Applications of Artificial Intelligence*, pages 14–25. Springer, 2014.
4. S. Board. Revealing information in auctions: the allocation effect. *Economic Theory*, 38(1):125–135, 2009.
5. J. Bredin, D. Parkes, and Q. Duong. Chain: A dynamic double auction framework for matching patient agents. *Journal of Artificial Intelligence Research (JAIR)*, 30:133–179, 2007.
6. S. Das and E. Kamenica. Two-sided bandits and the dating market. In *Proceedings of the Nineteenth International Joint Conference on Artificial Intelligence*, pages 947–952, Edinburgh, UK, August 2005.
7. S. Dobzinski and N. Nisan. Mechanisms for multi-unit auctions. *Journal of Artificial Intelligence Research (JAIR)*, 37:85–98, 2010.
8. M. Dufwenberg and U. Gneezyi. Information disclosure in auctions: an experiment. *Journal of Economic Behavior & Organization*, 48(4):431–444, 2002.
9. S. Dughmi, N. Immorlica, and A. Roth. Constrained signaling for welfare and revenue maximization. *CoRR*, abs/1302.4713, 2013.
10. Y. Emek, M. Feldman, I. Gamzu, R. Leme, and M. Tennenholtz. Signaling schemes for revenue maximization. In *Proceedings of the 12th ACM Conference on Electronic Commerce (EC-12)*, pages 514–531, 2012.
11. P. Eső and B. Szentes. Optimal information disclosure in auctions and the handicap auction. *The Review of Economic Studies*, 74(3):705–731, 2007.
12. J.J. Ganuza and J.S. Penalva. Signal orderings based on dispersion and the supply of private information in auctions. *Econometrica*, 78(3):1007–1030, 2010.
13. J. Goeree and T. Offerman. Efficiency in auctions with private and common values: An experimental study. *American Economic Review*, 92(3):625–643, 2002.
14. J. Goeree and T. Offerman. Competitive bidding in auctions with private and common values. *The Economic Journal*, 113:598–613, 2003.
15. C. Hajaj, N. Hazon, and D. Sarne. Improving comparison shopping agents? competence through selective price disclosure. *Electronic Commerce Research and Applications*, 14(6):563–581, 2015.
16. C. Hajaj, N. Hazon, and D. Sarne. Enhancing comparison shopping agents through ordering and gradual information disclosure. *to appear in Autonomous Agents and Multi-Agent Systems*, 2016.
17. C. Hajaj and D. Sarne. Strategic information platforms: selective disclosure and the price of free. In *Proceedings of the Fifteenth ACM conference on Economics and Computation*, pages 839–856. ACM, 2014.
18. M. Hajiaghayi, R. Kleinberg, M. Mahdian, and D. Parkes. Online auctions with re-usable goods. In *Proceedings of the Sixth ACM Conference on Electronic Commerce (EC-2005)*, pages 165–174, 2005.
19. I. Jewitt and D. Li. Cheap-talk information disclosure in auctions. Technical report, Working paper, 2012.
20. J.P. Johnson and D.P. Myatt. On the simple economics of advertising, marketing, and product design. *The American Economic Review*, pages 756–784, 2006.
21. A. Juda and D. Parkes. An options-based solution to the sequential auction problem. *Artificial Intelligence*, 173(7-8):876–899, 2009.
22. P. Klemperer. *Auctions: Theory and Practice*. Princeton University Press, 2004.
23. C. Lim, J.N. Bearden, and J.C. Smith. Sequential search with multiattribute options. *Decision Analysis*, 3(1):3–15, 2006.

24. P. Miltersen and O. Sheffet. Send mixed signals: earn more, work less. In *Proceedings of the 13th ACM Conference on Electronic Commerce(EC '12)*, pages 234–247, 2012.
25. G. Moscarini and L. Smith. The optimal level of experimentation. *Econometrica*, 69(6):1629–1644, 2003.
26. Y. Nahum, D. Sarne, S. Das, and O. Shehory. Two-sided search with experts. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, pages 754–771. ACM, 2012.
27. I. Rochlin, Y. Aumann, D. Sarne, and L. Golosman. Efficiency and fairness in team search with self-interested agents. *Autonomous Agents and Multi-Agent Systems*, 30(3):526–552, 2016.
28. I. Rochlin and D. Sarne. Utilizing costly coordination in multi-agent joint exploration. *Multiagent and Grid Systems*, 10(1):23–49, 2014.
29. I. Rochlin and D. Sarne. Constraining information sharing to improve cooperative information gathering. *Journal Artificial Intelligence Research (JAIR)*, 54:437–469, 2015.
30. I. Rochlin, D. Sarne, and M. Mash. Joint search with self-interested agents and the failure of cooperation enhancers. *Artificial Intelligence*, 214:45–65, 2014.
31. I. Rochlin, D. Sarne, and G. Zussman. Sequential multi-agent exploration for a common goal. *Web Intelligence and Agent Systems*, 11(3):221–244, 2013.
32. M. Rysman. The economics of two-sided markets. *The Journal of Economic Perspectives*, pages 125–143, 2009.
33. D. Sarne, S. Alkoby, and E. David. On the choice of obtaining and disclosing the common value in auctions. *Artificial Intelligence*, 215:24–54, 2014.
34. P. Stone, R. Schapire, M. Littman, J. Csirik, and D. McAllester. Decision-theoretic bidding based on learned density models in simultaneous, interacting auctions. *Journal of Artificial Intelligence Research*, 19(1):209–242, September 2003.
35. M. Tennenholtz. Tractable combinatorial auctions and b-matching. *Artificial Intelligence*, 140(1/2):231–243, 2002.
36. W. Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance*, 16(1):8–37, 1961.
37. D.D. Wiegmann, K.L. Weinersmith, and S.M. Seubert. Multi-attribute mate choice decisions and uncertainty in the decision process: A generalized sequential search strategy. *Journal of Mathematical Biology*, 60(4):543–572, 2010.
38. K. E. Wilson, R. Szechtman, and M. P. Atkinson. A sequential perspective on searching for static targets. *European Journal of Operational Research*, 215(1):218 – 226, 2011.