Manipulating Information Providers Access to Information in Auctions

Shani Alkoby [a], David Sarne [a], and Esther David [b]

[a] Department of Computer Science
Bar-Ilan University, Ramat-Gan, Israel.
[b] Department of Computer Science
Ashkelon Academic College, Ashkelon, Israel.
shani.alkoby@gmail.com
sarned@cs.biu.ac.il
astrdod@acad.ash-college.ac.il

Abstract. *Information purchasing is a crucial issue that auctioneers have to consider when running auctions, in particular in auction settings where the auctioned item's value is affected by a common value element. In such settings it is reasonable to assume the existence of a self-interested information provider. The main contribution of the information provider may be the elimination of some uncertainty associated with the common value of the auctioned item. The existence of an information provider does not necessarily impose the use of its services. Moreover, in cases in which the auctioneer decides to purchase information, it is not always beneficial for him to disclose it. In this work, we focus on environment settings where the information that may purchased still involves some uncertainty. The equilibrium analysis is provided with illustrations that highlight some non-intuitive behaviors. In particular, we show that in some cases it is beneficial for the auctioneer to initially limit the level of detail and precision of the information he may purchase. This can be achieved, for example, by limiting the information provider's access to some of the data required to determine the exact common value. This result is non-intuitive especially in light of the fact that the auctioneer is the one who decides whether or not to use the services of the information provider; hence having the option to purchase better information may seem advantageous.*

Keywords: auction, common value, self interested auctioneer, information provider

1 Introduction

One of the main crucial issues that an auction mechanism designer should take into account is information disclosure. Namely, what part of the information should be revealed in order to maximize the auctioneer's target utility which can either be related to the auctioned good's expected revenue in case of a self-interested designer or social welfare in cases in which the auction designer acts as the social planner. Many researchers have explored this issue in both theoretical and empirical manners [9, 12]. In particular, this issue becomes more relevant in auction settings where the auctioned item involves an uncertain common value element [11, 29, 14, 13, 20, 21, 4]. For example, the board of a firm for sale can choose which part of the firm's client list or its sales forecast will be disclosed to the potential buyers. The decision regarding the information disclosure directly affects bidders' valuation of the auctioned item and consequently also the winner's determination and the auctioneer's expected revenue.

More often, information regarding the common value element is not available to the auctioneer before the auction. However, the auctioneer may use some relevant expert services termed external information provider. This situation is common in scenarios where information discovery involves special expertise or equipment the auctioneer does not own. Specifically, in the scenario of a firm for sale, 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. In such situations the auctioneer's responsibility is to decide both whether to purchase the information and whether to disclose it fully or partially to bidders when purchased. Such scenarios become much more complex when the information provider acts strategically, controlling the accuracy of the information provided and its price.

Prior work in such settings assumed strategic behavior on the auctioneer and the information provider sides. However, the auctioneer's strategy was limited to the choice of the information to be disclosed to the buyers [4, 11] while the information provider's strategy was limited to setting the price of the information provided (i.e., assume the information provided is fully certain and captures the exact common value [34]).

In this paper we extend the model to the more realistic case, where the information provider cannot guarantee the identification of the true common value, but rather can offer a more precise estimate of this variable. In particular we focus on the case in which the information provider can only eliminate some of the possible values and cannot fully distinguish between others. For example in the example of the firm, it is possible that the information provider will be able to classify customers as "good" and "bad" where each category spans a wide range of possible values. Similarly, it is possible that the information provider will be able to distinguish between strong and weak sales forecasts, but will not be able to differentiate between a wide range slightly above or below the average sales figures.

To this end, the paper's contribution is twofold:

– We augment the three-ply equilibrium analysis (considering the strategic behavior of the information provider, the auctioneer and the bidders) to cases where the information provider can reduce the uncertainty associated with the common value rather than provide its true value.
– We illustrate a beneficial, yet somewhat non-intuitive, strategic behavior of the auctioneer. In particular, this behavior is the auctioneer's choice to intentionally limit the information provider's (e.g., the analyst) ability to distin-

guish between values. This becomes possible when the information provider's ability to provide accurate information depends on inputs received from the auctioneer. For example, in the sale of the firm example, the board can decide to provide the analyst with accurate, yet aggregative, information, such that the information provider can estimate future sales as weak or strong rather than a certain figure from a wider range of values. The non-intuitiveness of doing this is attributed to the fact that at the end of the day the information provider's information is offered for sale to the auctioneer herself, thus by restricting the information provider's ability to distinguish between values the auctioneer restricts himself by not having the choice of purchasing more accurate information.

The paper is structured as follows. In the following section we provide a formal presentation of the model. Then, we present an equilibrium analysis and illustrate the potential profit for the auctioneer from influencing the accuracy of the information that can be provided by the information provider. Finally we conclude with a review of related work and a discussion on the main findings.

2 The Model

Our 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 [16, 17]. 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 $T = t$ if the true value of the item is $X = x$. It is assumed that the probability distribution 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 information related to the value of X from an outside source, denoted "information provider", by paying a fee C that is set by the information provider. Similar to prior models (e.g., [34]), and for the same justifications given there, it is assumed that the option of purchasing the information is available only to the auctioneer, though the bidders are aware of the auctioneer's option to purchase such information. In its most general form, the information provided by the information provider is a subset $X' \subset X^*$, ensuring that one of the values in X' is the true common value. This is usually the case when the information provider cannot distinguish between some of the possible outcomes however can eliminate others. Therefore, the information provider will provide a subset $X' \in D = \{X_1, ..., X_l\}$ where D is the set of possible subsets of X^*, each containing values between which the information provider cannot distinguish, such that $\cup_{X_i \in D} X_i = X^*$ and $X_i \cap X_j = \emptyset, \forall i, j$.

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

The above model generalizes the one found in [11, 29] in the sense that it requires that the auctioneer decide whether or not to purchase the external information rather than assume that she initially possesses it. Similarly, it generalizes the work in [34] in the sense that it allows the information provider to provide a subset of values rather than the specific true value.

3 Analysis

Our analysis uses the concept of mixed Bayesian Nash Equilibrium. Since the auctioneer needs to decide both whether to purchase the information and if so whether to disclose the information received, we can characterize her strategy using $R^{auc} = (p^a, p_1^a, ..., p_l^a)$ where p^a is the probability she will purchase the information from the information provider and p_i^a $(1 \leq i \leq l)$ is the probability she will disclose to the bidders the subset received if that subset is X_i. The dominating bid of a bidder of type t, when subset X' is received (including the case where $X' = \emptyset$, i.e., no information is disclosed), denoted $B(t, X')$, is the expected private value calculated by weighing each private value $V_t(x)$ according to the post-priori probability of x being the true common value given the information X', denoted $Pr(X = x|X')$ [11], i.e.: $B(t, X') = \sum_{x \in X^*} V_t(x) \cdot Pr(X = x|X')$. If the auctioneer discloses a subset $X' \subset X^* \neq \emptyset$ then $Pr(X = x|X') = \frac{Pr(X=x)}{\sum_{y \in X'} Pr(X=y)}$ for any $x \in X'$ and $Pr(X = x|X') = 0$ otherwise. If no information is disclosed $(X' = \emptyset)$ then $Pr(X = x|X' = \emptyset)$ needs to be calculated based on the bidders' belief of whether information was indeed purchased and if so, whether that value is intentionally not disclosed by the auctioneer. Assume the bidders believe that the auctioneer has purchased the information from the information provider[1] with a probability of p and that if indeed purchased then if the information received was the subset X_i then it will be disclosed to the bidders with a probability of p_i. In this case the probability of any value $x \in X_i$ being the true common value is given by:

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

$$Pr(X{=}x|X'{=}\emptyset) = \frac{Pr(X = x)(p(1 - p_i) + (1 - p))}{(1 - p){+}p\sum_{X_j}(1 - p_j)\sum_{y \in X_j} Pr(X = y)} \quad (1)$$

The term in the numerator is the probability that x is indeed the true value however the subset it is in is not disclosed. If indeed x is the true value (i.e., with a probability of $Pr(X = x)$) then it is not disclosed either when the information is not purchased (i.e., with a probability of $(1 - p)$) or when purchased but not disclosed (i.e., with a probability of $p(1{-}p_i)$). The term in the denominator is the probability information will not be disclosed. This happens when the information is not purchased (i.e., with a probability $(1 - p)$) or when the information is purchased however the auctioneer does not disclose the subset received (i.e., with a probability of $p\sum(1 - p_j)\sum_{y \in X_j} Pr(X = y)$). Further on in the paper we refer to the strategy where information is not disclosed as an empty set. 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 purchased and p^b_i is the probability they assign to the event that the information is indeed disclosed if purchased and becomes X_i.

In order to formalize the expected second-best bid if the auctioneer discloses the subset X' we apply the calculation method given in [34] but replace the exact value X with a subset x'. We first define two probability functions. The first is the probability that given that the subset disclosed by the auctioneer is X' , the bid placed by a random bidder equals w, denoted $g(w, X')$, given by: $g(w, X') = \sum_{B(t,X')=w} Pr(T = t)$. The second is the probability that the bid placed by a random bidder equals w or below, denoted $G(w, X')$, given by: $G(w, X') = \sum_{B(t,X')\leq w} Pr(T = t)$.

The auctioneer's expected profit when disclosing the subset X', denoted $ER_{auc}(X')$, equals the expected second-best bid:

$$\begin{aligned} ER_{auc}(X') = &\sum_{w \in \{B(t,X')|t \in T\}} w(\sum_{k=1}^{n-1} n\binom{n-1}{k} \\ &(1{-}G(w, X'))(g(w, X'))^k(G(w, X') - g(w, X'))^{n-k-1} \\ &+\sum_{k=2}^{n}\binom{n}{k}(g(w, X'))^k(G(w, X') - g(w, X'))^{n-k}) \end{aligned} \quad (2)$$

The calculation iterates over all of the possible second-best bid values, assigning to each its probability of being the second-best bid. As we consider discrete probability functions, it is possible to have two bidders place 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 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.

Consequently, the auctioneer's expected revenue from the auction itself (i.e., excluding the payment C to the information provider), when the auctioneer uses

$R^{auc} = (p^a, p^a_1, ..., p^a_k)$ and the bidders use R^{bidder}, denoted $ER(R^{auc}, R^{bidder})$, is given by:

$$ER(R^{auc}, R^{bidder}) = p^a \sum_{i=1}^{l} \sum_{x \in X_i} Pr(X = x) p^a_i \cdot ER_{auc}(X_i) \\ + ((1-p^a) + p^a \sum_{i=1}^{l} \sum_{x \in X_i} Pr(X = x)(1 - p^a_i)) \cdot ER_{auc}(\emptyset) \tag{3}$$

where $ER_{auc}(X_i)$ is calculated according to (2) (also in the case where $X_i = \emptyset$). Consequently the auctioneer's expected benefit, denoted $EB(R^{auc}, R^{bidder})$, is given by $EB(R^{auc}, R^{bidder}) = ER(R^{auc}, R^{bidder}) - p^a * C$.

A stable solution in terms of the mixed Bayesian Nash Equilibrium in this case is necessarily of the form $R^{auc} = R^{bidder} = R = (p, p_1, ..., p_l)$ (because otherwise, if $R^{auc} = R' \neq R^{bidder}$ then bidders necessarily have an incentive to deviate to $R^{bidder} = R'$), such that: (a) for any $0 < p_i < 1$ (or $0 < p < 1$): $ER_{auc}(\emptyset, R) = ER_{auc}(X_i)$ (or $ER_{auc}(\emptyset, R^{bidder}) = ER_{auc}((1, p_1, ..., p_l), R^{bidder})$); (b) for any $p_i = 0$ (or $p = 0$): $ER_{auc}(\emptyset, R^{bidder}) \geq ER_{auc}(X_i)$ (or $ER_{auc}(\emptyset, R^{bidder}) \geq ER_{auc}((1, p_1, ..., p_l), R^{bidder})$; and (c) for any $p_i = 1$ (or $p = 1$): $ER_{auc}(\emptyset, R^{bidder}) \leq ER_{auc}(X_i)$ (or $ER_{auc}(\emptyset, R^{bidder}) \leq ER_{auc}((1, p_1, ..., p_l), R^{bidder})$. The proof for this derivation is similar to the proof given in [34] (see page 39), with the exception that instead of referring to individual values of X we refer to subsets of values X_i. Therefore we need to evaluate all the possible solutions of the form $(p, p_1, ..., p_l)$ 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 (because there is 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 whether or not to purchase information), the equilibrium used is the stable solution for which the auctioneer's expected profit is maximized.

We note that 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 [11] for the pure Bayesian Nash Equilibrium case and in [29] for the mixed Bayesian Nash Equilibrium case. Similarly, if $|X_i| = 1\ \forall i$ is enforced (i.e., the information provider provides the exact value of X) then the resulting equilibrium is the same as the one given in [34].

4 Influencing the Information Provider's Capabilities to distinguish between values

As discussed in the introduction, in various settings the auctioneer can influence the information provider's ability to distinguish between different values the common value obtains. In this section we consider the case where the auctioneer has full control over the structure of D, i.e., the division of X^* into disjoint subsets, each composed of values which the information provider cannot distinguish between.

Limiting the information provider's ability to distinguish between values may seem non-intuitive in the sense that it limits the auctioneer's strategy space when it comes to disclosing this information to bidders, if it is purchased. Nevertheless, in many settings the strategy of constraining the information provider's input can actually play into the hands of the auctioneer and improve her expected profit. This phenomenon is illustrated in Figure 1, which depicts the auctioneer's expected profit (vertical axis) as a function of the information purchasing cost (horizontal axis), for several possible divisions of X^* into subsets of non-distinguishable values. The setting used for this example is given in the table below the graph. It is based on three bidders, where each can be of four different types. The first column of the table depicts the different bidder types and the second column gives their probability. Similarly, the second and third rows depict the different possible values of X (denoted x_1,x_2,x_3 and x_4) and their probabilities. The remaining values are the valuations that bidders of different types assign different possible values of the parameter X. For example, if a bidder is of type 3, then her valuation of x_2 is 59.

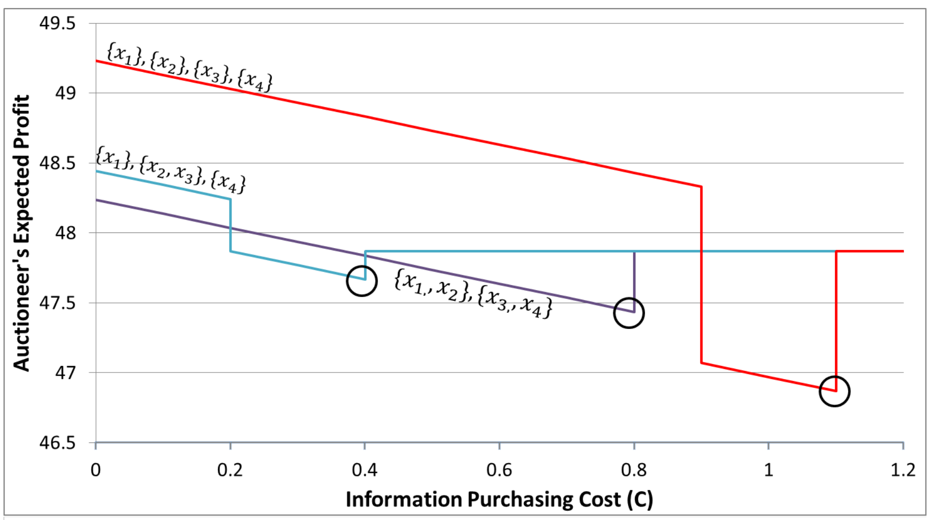

n=3		private values			
		x1	x2	x3	x4
Type	p(type) \ p(value)	0.626	0.159	0.143	0.072
1	0.252	54	25	84	10
2	0.33	79	56	0	53
3	0.386	21	59	46	69
4	0.032	25	94	69	93

Fig. 1. The auctioneer's expected profit as a function of the information purchasing cost for different divisions of X^* into subsets of non-distinguishable values.

Each of the three graphs given in the figure relates to different possible divisions, d of X^* (marked next to it), depicting the expected profit of the auctioneer in the equilibrium resulting in the specific cost of information on the horizontal axis. In this example the resulting equilibrium is always based on pure strategies (i.e., $p, p_i \in \{0, 1\}$) and the points of discontinuity in the curve represent the transition from one equilibrium to another. In particular, for C values in which the curve decreases, the equilibrium is based on always purchasing the information (though not necessarily disclosing all subsets). This happens when the cost of purchasing the information justifies its purchase, i.e., for relatively small C values. The non-decreasing part of the curve is associated with an equilibrium in which the information is essentially not purchased.

As can be seen from the figure, for any cost of purchasing the information $0.9 < C < 1.1$, the auctioneer is better off not allowing the information provider to distinguish between all values: the division $d = \{\{x_1\}, \{x_2\}, \{x_3\}, \{x_4\}\}$ is dominated by $d' = \{\{x_1\}, \{x_2, x_3\}, \{x_4\}\}$ and $d'' = \{\{x_1, x_2\}, \{x_3, x_4\}\}$. The explanation for this interesting phenomenon lies in the different costs of the transition between equilibria due to stability considerations. With fully distinguishable values, it is possible that a desired solution which yields the auctioneer a substantial expected profit is not stable (e.g., in our case when $0.9 < C < 1.1$ the solution is that the information is not purchased at all), whereas with inaccurate information the solution is stable and holds as the equilibrium.

In particular, in our example, when the information provider acts fully strategically, i.e., sets the price of information to the maximum possible price for which the information will still be purchased (the C value in which the equilibrium changes from purchase to not purchase the information, marked with circles in the graphs) the auctioneer will gain (and the information provider will essentially lose) from restricting the information provider's ability to distinguish between values. For example, with $\{\{x_1\}, \{x_2, x_3\}, \{x_4\}\}$ the information will be priced at $C = 0.4$ yielding the auctioneer an expected profit of 47.6 (compared to $C = 1.1$ and a profit of 46.8 in the "fully distinguishable" case).

5 Related Work

Auctions are an effective means of trading and allocating goods whenever the seller is unsure about buyers' (bidders') exact valuations of the sold item [24, 25]. The advantage of many auction mechanism variants in this context is in the ability to effectively extract the bidders' valuations [23, 32], resulting in the most efficient allocation. Due to its many advantages, this mechanism is commonly used and researched and over the years has evolved to support various settings and applications such as on-line auctions [22, 27, 19, 37, 36], matching agents in dynamic two-sided markets [5], resource allocation [31, 30, 7] and even for task allocation and joint exploration [15, 26]. In this context great emphasis has been placed on studying bidding strategies [40, 38, 3], the use of software agents to represent humans in auctions [6], combinatorial auctions [39] and the development of auction protocols that are truthful [5, 8, 7, 2] and robust (e.g., against false-name bids in combinatorial auctions [41]).

The case where there is some uncertainty associated with the value of the sold (auctioned) item is quite common in the literature on auctions. 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 [17, 24]. 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, 11]. As such, much emphasis was placed on the role of information revelation [28, 33, 9, 12, 14, 13, 20, 21]. In particular, several authors 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 [11, 29, 10]. Nonetheless, all these works assume the auctioneer necessarily obtains the information and that the division into non-distinguishable groups, whenever applicable, is always a priori given to the bidders. Furthermore, not disclosing any information (signal) is not allowed in these works. Our problem, on the other hand, does not require that the auctioneer possess (or purchase) the information in the first place, and allows the auctioneer the decision of whether or not to disclose any value even if the information is purchased. In particular, when no information is disclosed bidders cannot distinguish between the information not being purchased in the first place and the information is purchased but the value is not disclosed. More importantly, none of the prior work considers the option of influencing the ability of the information provider to distinguish between different values.

Prior work that considers a three-ply equilibrium in settings where information can be potentially purchased from an external information provider assumes the information provider can always supply the true common value [34]. Moreover, this work does not allow any influence whatsoever on the auctioneer's strategy over the ability to distinguish between different values. Work in other domains that did consider selective information disclosure, e.g., for comparison shopping agents [18] or for sharing data for user modeling [35] is very far in terms of the principles used, and cannot be applied in our case. On the whole, despite the many prior models that consider a subset of our model's characteristics, to the best of our knowledge, an analysis that addresses all of the different aspects included in our model does not exist in the literature.

6 Conclusions and Future Work

In this paper we advance the state of the art by providing a three player equilibrium analysis that allows the ability of influencing the auctioneer's expected profit through controlling the granularity and accuracy of the information offered for sale. The presence of information providers in multi agent systems has become substantial and consequently, enforces the reconsideration of the equilibrium where this time such options are taken into account. The information providers may be individuals with specific expertise who offer their services for a fee (e.g. an analyst) or large information service providers such as Carfax.con or credit report companies. It is commonly assumed that these information providers indeed can control the level of accuracy they offer their customers. Moreover, the accuracy of the information provided depends on the customer's cooperation and the level of the inputs she provides. Against this background, the importance of this equilibrium construction and analysis for auctioneers or the information providers is clear, especially, in terms of the ability to control the granularity in which information is provided.

Here, we show an interesting phenomenon where the auctioneer may benefit in cases where the information provider cannot fully identify the exact state of nature, even though the information is eventually offered exclusively to the auctioneer. This phenomenon is explained by the stability requirement – beneficial solutions that could not hold with the complete ("perfect") information scheme, because of stability considerations, are found to be stable once the information being offered for sale is constrained.

Acknowledgment

This work was supported in part by the Israeli Science Foundation grant no. 1083/13.

References

1. Akerlof, G.: The market for "lemons": Quality uncertainty and the market mechanism. The Quarterly Journal of Economics 84(3), 488–500 (1970), http://www.jstor.org/stable/1879431
2. Azoulay, R., David, E.: Truthful and efficient mechanisms for site oriented advertizing auctions, multiagent and grid systems. vol. 10(2), pp. 67–94. An International Journal Press (2014)
3. Bagnall, A., Toft, I.: Autonomous adaptive agents for single seller sealed bid auctions. Autonomous Agents and Multi-Agent Systems 12(3), 259–292 (2006)
4. Board, S.: Revealing information in auctions: the allocation effect. Economic Theory 38(1), 125–135 (2009)
5. Bredin, J., Parkes, D., Duong, Q.: Chain: A dynamic double auction framework for matching patient agents. Journal of Artificial Intelligence Research (JAIR) 30, 133–179 (2007)
6. David, E., Azoulay-Schwartz, R., Kraus, S.: Bidding in sealed-bid and english multi-attribute auctions. Decision Support Systems 42(2), 527–556 (2006)
7. David, E., Manisterski, E.: Strategy proof mechanism for complex task allocations in prior consent for subtasks completion environment. In: IAT. pp. 209–215 (2013)
8. Dobzinski, S., Nisan, N.: Mechanisms for multi-unit auctions. Journal of Artificial Intelligence Research (JAIR) 37, 85–98 (2010)
9. Dufwenberg, M., Gneezy, U.: Information disclosure in auctions: an experiment. Journal of Economic Behavior & Organization 48(4), 431–444 (2002)
10. Dughmi, S., Immorlica, N., Roth, A.: Constrained signaling for welfare and revenue maximization. ACM SIGecom Exchanges 12(1), 53–56 (2013)

11. Emek, Y., Feldman, M., Gamzu, I., Leme, R., Tennenholtz, M.: Signaling schemes for revenue maximization. In: Proceedings of the 12th ACM Conference on Electronic Commerce (EC-12). pp. 514–531 (2012)
12. Eső, P., Szentes, B.: Optimal information disclosure in auctions and the handicap auction. The Review of Economic Studies 74(3), 705–731 (2007)
13. Ganuza, J.J., Penalva, J.S.: Signal orderings based on dispersion and the supply of private information in auctions. Econometrica 78(3), 1007–1030 (2010)
14. Ganuza Fernández, J.J.: Ignorance promotes competition: An auction model with endogenous private valuations (2005)
15. Gerkey, B., Mataric, M.: Sold!: auction methods for multirobot coordination. IEEE Transactions on Robotics 18(5), 758–768 (2002)
16. Goeree, J., Offerman, T.: Efficiency in auctions with private and common values: An experimental study. American Economic Review 92(3), 625–643 (2002)
17. Goeree, J., Offerman, T.: Competitive bidding in auctions with private and common values. The Economic Journal 113, 598–613 (2003)
18. Hajaj, C., Hazon, N., Sarne, D., Elmalech, A.: Search more, disclose less. In: Proceedings of the Twenty-Seventh AAAI Conference on Artificial Intelligence (2013)
19. Hajiaghayi, M., Kleinberg, R., Mahdian, M., Parkes, D.: Online auctions with re-usable goods. In: Proceedings of the Sixth ACM Conference on Electronic Commerce (EC-2005). pp. 165–174 (2005)
20. Jewitt, I., Li, D.: Cheap-talk information disclosure in auctions. Tech. rep., Working paper (2012)
21. Johnson, J.P., Myatt, D.P.: On the simple economics of advertising, marketing, and product design. The American Economic Review pp. 756–784 (2006)
22. Juda, A., Parkes, D.: An options-based solution to the sequential auction problem. Artificial Intelligence 173(7-8), 876–899 (2009)
23. Klemperer, P.: Auction theory: A guide to the literature. Journal of Economic Surveys 13(3), 227–286 (1999)
24. Klemperer, P.: Auctions: Theory and Practice. Princeton University Press (2004)
25. Krishna, V.: Auction Theory. Academic Press (2002)
26. Lagoudakis, M., Markakis, E., Kempe, D., Keskinocak, P., Kleywegt, A., Koenig, S., Tovey, C., Meyerson, A., Jain, S.: Auction-based multi-robot routing. In: Proceedings of Robotics: Science and Systems. pp. 343–350 (2005)
27. Lavi, R., Nisan, N.: Competitive analysis of incentive compatible on-line auctions. In: Proceedings of the Second ACM conference on Electronic commerce (EC '00). pp. 233–241 (2000)
28. Milgrom, P., Weber, R.: A theory of auctions and competitive bidding. Econometrica 50(5), 1089–1122 (1982)
29. Miltersen, P., Sheffet, O.: Send mixed signals: earn more, work less. In: Proceedings of the 13th ACM Conference on Electronic Commerce(EC '12). pp. 234–247 (2012)
30. Ng, C., Parkes, D., Seltzer, M.: Virtual worlds: fast and strategyproof auctions for dynamic resource allocation. In: Proceedings Fourth ACM Conference on Electronic Commerce (EC-2003). pp. 238–239 (2003)
31. Nisan, N.: Algorithms for selfish agents. In: Proceedings of the Sixteenth Annual Symposium on Theoretical Aspects of Computer Science (STACS 99). pp. 1–15 (1999)
32. Parkes, D., Shneidman, J.: Distributed implementations of vickrey-clarke-groves mechanisms. In: Proceedings of the Third International Conference on Autonomous Agents and Multiagent Systems (AAMAS-04). vol. 1, pp. 261–268 (2004)
33. Perry, M., Reny, P.: On the failure of the linkage principle in multi-unit auctions. Econometrica 67(4), 895–900 (1999)

34. Sarne, D., Alkoby, S., David, E.: On the choice of obtaining and disclosing the common value in auctions. Artificial Intelligence 215, 24–54 (2014)
35. Sarne, D., Grosz, B.J.: Sharing experiences to learn user characteristics in dynamic environments with sparse data. In: Proceedings of the Sixth International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2007). pp. 202–209 (2007)
36. Sarne, D., Hadad, M., Kraus, S.: Auction equilibrium strategies for task allocation in uncertain environments. In: Proceedings of the Eight International Workshop on Cooperative Information Agents (CIA). pp. 271–285 (2004)
37. Sarne, D., Kraus, S.: Solving the auction-based task allocation problem in an open environment. In: Proceedings of the Twentieth National Conference on Artificial Intelligence and the Seventeenth Innovative Applications of Artificial Intelligence Conference (AAAI 2005). pp. 164–169 (2005)
38. Stone, P., Schapire, R., Littman, M., Csirik, J., McAllester, D.: Decision-theoretic bidding based on learned density models in simultaneous, interacting auctions. Journal of Artificial Intelligence Research 19(1), 209–242 (Sep 2003)
39. Tennenholtz, M.: Tractable combinatorial auctions and b-matching. Artificial Intelligence 140(1/2), 231–243 (2002)
40. Vetsikas, I., Jennings, N.: Bidding strategies for realistic multi-unit sealed-bid auctions. Autonomous Agents and Multi-Agent Systems 21(2), 265–291 (2010)
41. Yokoo, M., Sakurai, Y., Matsubara, S.: Robust combinatorial auction protocol against false-name bids. Artificial Intelligence 130(2), 167–181 (2001)