New Privacy-Preserving Ascending Auction for Assignment Problems

De Liu
University of Minnesota
deliu@umn.edu *

Adib Bagh
University of Kentucky
adib.bagh@uky.edu†

March 7, 2016

Abstract

We introduce a new ascending auction that allocates heterogeneous objects among bidders with purely private unit demands. Our auction design differs from existing dynamic auctions in a number of ways: it economizes on information solicited from bidders by requiring marginal bidders to reveal a single new bid at a time; it uses a transparent price adjustment process; and it allows the seller to set starting prices above his reservation valuations. Despite these new features, (i) the auction stops in a finite time, (ii) sincere bidding is an ex-post Nash equilibrium, (iii) the auction ending prices and revenue depend only on bidders valuations and starting prices, and (iv) the auction is efficient if it starts with the seller's valuations. To test the new auction's privacy preservation feature, we propose a novel entropy-based measure of bidder privacy, and a hybrid quasi-Monte Carlo procedure that allows the measure to be numerically computed and compared across any number of mechanisms. Our numerical simulation illustrates that our auction reveals about 10% less information about bidders compared to an improved version of the Demange, Gale and Shapley (1986)'s dynamic auction over a good range of parameters.

Keywords: Assignment Problem, Ascending Auctions, Privacy Preservation, Entropy, Quasi-Monte Carlo.

*Corresponding author. 321 19th Ave S, Minneapolis, MN 55455, USA. Phone:+011(612)626-4480. This research received support from the National Science Foundation of China under Grant No.71171052. The authors wishes to thank Xinyi Zhao for research assistance on the initial computation of entropy, and discussion and feedback by seminar participants at Workshop on Information Technology and Systems, Theory in Economics of Information Systems, INFORMS annual meetings, University of Minnesota, University of Kentucky, and the University of International Business and Economics.

†335Y Gatton College of Business and Economics, University of Kentucky, Lexington KY, 40506, USA.

1 Introduction

We introduce a dynamic ascending auction that assigns heterogeneous items among several bidders. We focus on the case where bidders have unit demand and their valuations are private. Such auctions may be used for task allocation in online crowdsourcing labor markets, allocation of rental properties in AirBnB type markets, allocation of Electronic Vehicle (EV) charging slots across multiple locations, and so on. Recent literature on multi-item dynamic auctions considers bidders with more general demand functions (Gul and Stacchetti, 2000; Bikhchandani and Ostroy, 2002; Ausubel and Milgrom, 2002; Ausubel, 2006; de Vries et al., 2007; Perry and Reny, 2005; Mishra and Parkes, 2009). However, problems with unit demand, due to their well-known applications in matching markets and job assignments (Gale and Shapley, 1962; Leonard, 1983; Crawford and Knoer, 1981), remain an important subclass and warrant special attention (Mishra and Parkes, 2009; Sankaran, 1994; Andersson et al., 2013).

Our contribution is twofold. First, we propose a new ascending auction design that economizes on information solicited from bidders. Bidders are generally reluctant to reveal information regarding their preferences, especially when they expect to participate in subsequent activities and negotiations in which information about their preferences over the auctioned items is relevant (Ausubel, 2004; Rothkopf et al., 1990). For example, knowing bidders true willingness to pay can lead to price discriminations in subsequent sales. As a result of better privacy protection and informational efficiency, we can expect our proposed auction to promote greater bidder participation than designs that require bidders to reveal their entire demand sets at every given price.

Our second contribution is a novel entropy-based method for measuring the amount of information revealed about bid valuations, as a bidder privacy measure, and a procedure for calculating such a measure. To our knowledge, no prior work has conducted formal analysis on the information revealed about bidder values. Our measurement, based on Shannon's information entropy, is general enough to be used by any number of mechanisms, including both dynamic and one-shot auctions. We further propose a procedure for extracting the informational events from auction paths and estimating changes in entropy as a result of such events. The calculation of bidder information entropy amounts to a high-dimensional integration problem. To ensure a good convergence rate, we use a mixture of analytical methods and quasi-Monte Carlo simulation methods.

Nearly all existing dynamic auctions require all bidders to report, at any iteration, their *entire* demand set, that is, what they desire the most given the current prices (Demange et al., 1986; Gul and Stacchetti, 2000). The reported demand sets are then used to determine a set of items whose prices should be increased - a key step for dynamic auctions. Such a clearing-house style reporting is reminiscent of simultaneous auction designs where all bidders must simultaneously submit bids and winners are calculated at once. Such centralized report is a quick and centralized way of gathering demand information. However, it relies on the success and complete reporting of demand sets at any time. Failure to report or report fully could disrupt the auction process, causing misallocations, lost revenues, and chaos. A few authors have noted that the demand-set reporting is excessive, and cause a number of problems due to inaccurate bids(Ausubel, 2004; Perry and Reny, 2005; Rothkopf et al., 1990). To overcome these issues, at any time, we inquire only a selected subset of bidders on their intent to submit a new bid, and only one needs to report a single new bid from his demand set each time. In other words, we solicit bidder demand on an "as-needed" basis, using *iterative* methods within each auction iteration to gather enough demand information that allows use to choose a set of over-demanded items for price increase.

Despite our departure from demand-set reporting, our auction maintains several desirable properties: (i) the auction stops in a finite time, (ii) sincere bidding at every stage of the auction is an ex-post Nash equilibrium, and (iii) for a given set of valuations, the auction revenue and ending prices depend only on starting prices (i.e., path independence). Moreover, if our auction starts with the auctioneer's reservation value, then (iv) it ends with an efficient allocation and (v) Vickery-Clarke-Groves (VCG) payments. To our knowledge, this is the first multi-item ascending auction that does not require full demand set reporting while maintaining a sincere bidding equilibrium.

Demange et al. (1986) (henceforth DGS)'s "exact auction", a pioneering work in the domain of dynamic auctions, can also allocate heterogenous items while satisfying properties (i), (ii), (iii) and (v). However, it requires a) each bidder to report his entire demand set at every posted price, and b) the auctioneer to compute a "minimal over-demanded set" of items, which requires not only the set is an *over-demanded set*, i.e., the number of bidders who *only* demand from the set should exceed the number of items in the set, but also "minimal" in the sense that no subset of the over-demand set should be an over-demanded set. The same reporting requirement is used in Gul and Stacchetti (2000)'s design and the more recent auctions of Ausubel (2006) and de Vries et al. (2007), which extend beyond unit demands. Though convenient for mathematical proofs, the calculation of minimal over-demand set requires complete demand information, and is rather opaque and complex. Sankaran (1994) relaxes the minimal over-demand set requirement and suggests that the same results can be achieved if over-demand sets are used in each iteration instead. But, the reporting requirements remain the same. Krishnappa and Plaxton (2011) examines a dynamic auction where every round consists of a sealed-bid auction such as sealed-bid unit demand auction with put options. Again, it employs the same clearing-house style simultaneous bidding within each auction iteration.

Noting the complex reporting requirements of existing dynamic auctions, Demange et al. (1986) proposed an "approximate" auction design as an alternative in the same seminal paper. In the approximate auction, Each item can accept at most one bidder and a new bidder displaces the incumbent bidder and increases the price by a small increment. DGS showed that, with a small enough increment, this auction can be made arbitrarily close to the efficient DGS "exact" auction that we have mentioned earlier. Our auction design is different from DGS's approximate auction in the key dimension of how items are selected for price increases. More importantly, the approximate auction assumes sincere bidding while it is guaranteed by the ex post equilibrium in our design.

To demonstrate our auction can indeed protect bidder privacy. We propose a new entropy-based method for measuring the amount of bidder privacy. Our work lays the foundation for analyzing bidder information revelation in a broad spectrum of dynamic and one-shot mechanisms. Our specific work towards measuring bidder information revelation include, (a) introducing a new bidder privacy measure based on Shannon's information entropy, (b) identifying the specific events in the course of a dynamic auction that is sufficient to capture information revelations, and (c) coming up a hybrid method that uses analytical and quai-Monte Carlo simulation methods to estimate the entropy reduction. We illustrates the approach by comparing our auction design with the DGS exact auction and VCG mechanism. Our results demonstrate benefits of our auction for bidder privacy, and confirm the overall advantages (and disadvantages) of dynamic auctions.

In a broader level, our research is related to work in applying dynamic auctions to grid-resources market (Bapna et al., 2011), and in supporting bidders in a complex combinational dynamic auctions (Adomavicius and Gupta, 2005; Petrakis et al., 2013). Though these studies focus on issues other than bidder privacy.

The rest of the paper proceeds as follows. We next describe our auction design and establish its theoretical prosperities. We then proceed to describe our information revelation measure and procedures for calculation the amount of information revealed for each bidder. Finally, we present and discuss results of several numerical experiments comparing our auction, DGS auction and the VCG mechanism.

2 Auction Design

Consider a set of indivisible items $\mathcal{J} = \{0, 1, 2, .., m\}$, where 0 is a dummy item representing outside options, and a set of bidders $\mathcal{I} = \{0, 1, 2, .., n\}$, where 0 denotes a special dummy bidder, the auctioneer. Unless otherwise noted, we use terms "bidder" and "item" to refer to regular (non-dummy) bidders and items. Each bidder has a unit demand for items (including dummy item). We use $v_{ij} \in [0, \infty)$ to denote bidders' (including the dummy bidder's) private valuation of items (including the dummy item). We interpret v_{0j} as the auctioneer's reservation valuations and fix $v_{00} = 0$.

The auction consists of iterations of price-adjustment and reassignment phases. In a price-adjustment phase, the prices of a chosen set of items, called the *active set*, will rise until a bidder signals his intent to switch to a new item outside of the active set (this bidder is called a *marginal bidder*). In a reassignment phase, bidders submit new bids and are reassigned but the prices of all items remain still.

To facilitate the description of the auction procedure, we can imagine a clock-like device publicly displayed for each non-dummy item. The clock displays the current price of the item, and has a switch that turns the clock on and off. In the "on" state, the display price will rise, in discrete steps or continuously.[1] In the "off" state, the price will stay at its current level. Several clocks can also be linked together in which case, they will rise and stop synchronously.

We denote $p^t = (p^t_0, p^t_1, p^t_2, ..., p^t_m)$ as a set of prices at iteration $t = 0, 1, ...$ The price for the outside option p^t_0 is fixed at zero at all times. We often omit the superscript t for simplicity. The details of the auction design are explained below.

1. (Initialization) The auctioneer sets starting prices at p^0 (i.e. starting value of clocks) and set all clocks to the "off" state. Bidders simultaneously name a single item as their most preferred item and are assigned to it. The iteration counter t is set to zero.

2. While the clocks are off (thus in a reassignment phase), the auctioneer decides between terminating the auction, entering a price adjustment phase, or continuing the reassignment phase using the following rules:

   (a) (Termination rule) If no item has more than one bidder, the auction ends. Each bidder receives the item he is assigned to and pays the current price of the item.

[1] our auction design works for both discrete and continuous values, though we assume discrete values in our study of bidder privacy.

(b) (Choose the initial active set) Among items with more than one bidders, choose the one with the largest number of bidders, say j, as the active set $\mathcal{A} = \{j\}$. Break a tie using an arbitrary tie-breaking rule.

(c) (Solicit intent to bid) Active bidders (i.e., bidders assigned to one of the active items) are inquired on their *intent to bid*, that is, given the current prices, whether they are indifferent between their current assignment and an inactive item (including dummy). If no bidder signals an intent to bid, enter a price adjustment phase (step 3). Otherwise, the bidder who signals his intent to bid is declared the *marginal bidder* and must submit an inactive item (including dummy) as his "new bid." If there are multiple signals, only one is let through (by arbitrary tie-breaking rule) and the auctioneer does not know the existence of other signals.[2].

(d) (Link clocks) If the marginal bidder has been assigned to other items (his indifferent items) in the current reassignment phase, expand the active set $\mathcal{A}$ to include these indifferent items. If the marginal bidder is the sole bidder of the item, find a bidder who has previously bid on the item in the current phase ("replacement" bidder), and include the item the bidder is currently at in the active set. If the bidder is also the sole bidder, search recursively until the replacement bidder is not the sole bidder of the item.[3]

If this step results in an expansion of the active set, then re-solicit intent to bid (step 2c). Otherwise, reassign the marginal bidder (step 4).

3. (Price adjustment) Turn on all the clocks in the active set and let the prices of all active items rise synchronously. The synchronous rise of prices continues until an active bidder signals an intent to revise his bid ("the marginal bidder"), which ends the current price adjustment phase. Again, if there are multiple signals, only one is let through and the auctioneer does not know the existence of other signals. As soon as the marginal bidder appears, proceed to Step 4 to reassign him.

4. (Reassignment) The marginal bidder is asked to submit an inactive item (including dummy item) as his new bid and is reassigned to it.[4] If this is an end of a price increase phase, let

[2]One may think of each bidder being in front of a silent bid button. A press of any button will register the bidder as the marginal bidder in this round, and disable the buttons in front of all other bidders.

[3]We can always find a finite chain of replacement bidder because the only way a single-bidder item can enter the active set is when a previous marginal bidder of the item has linked this item in from somewhere else.

[4] If the marginal bidder is the sole bidder at his current item, we will reassign a replacement bidder, among the active bidders, to this item, before reassigning the marginal bidder. If the replacement bidder is also a sole bidder at his item, apply the same rule recursively until we find a replacement bidder who is active and not the sole bidder

	Valuation for item			
bidders	1	2	3	0 (dummy item)
1	16	0	15	5
2	20	23	25	7
3	0	25	25	8
4	20	30	25	6
5	15	28	13	6
0 (auctioneer)	10	20	15	0

Table 1: Valuations

$t \leftarrow t+1$ and start a new iteration from Step 2a. Otherwise, continue the current reassignment phase at Step 2a.

As an example, consider an auction with five bidders and three items with valuations given by Table 1. For the purpose of illustration, we assume all bidders bid sincerely - that is, they bid on the item that provides the highest surplus at the current prices. Table 2 illustrates the auction process. At time 0, the auctioneer sets the starting prices to his reservation valuations and bidders bid on their preferred items. The auctioneer chooses item 1 (has bidders 1 & 2) to be "active". Bidder 2 signals bid intent (for item 3) and is reassigned to item 3. Next, the chosen active item is 3 (has bidders 2, 3, & 4) and bidder 2 again signals intent to bid (on item 1). Because bidder 2 has previously been assigned to item 1 at $t = 0$, the auctioneer links items 1 and 3 as the new active set (active bidders 1,2,3,& 4). Bidder 4 bids on item 2 and is reassigned to it. The reassignments continue until bidder 4 links all three items. At this moment, there is no intent to bid so the clock starts and prices of active items $\{1, 2, 3\}$ rise continuously until bidder 1 signals an intent to bid at $t = 1$ and drops out. After some additional reassignments of bidders 2 and 4, the clock starts again with all three items linked. The auction ends at $t = 2$ after bidder 5 drops out. In the end, items 1, 2, and 3 are assigned to bidders 2, 4, and 3 for prices 12, 22, and 17 respectively.

As noted in the auction procedure and the example, only active bidders are inquired on their intent to bid, and only one may become a marginal bidder who will be required to submit a single new bid. Rather than asking for the complete demand information from all the bidders to build a (minimal) over-demanded set at each iteration (as in DGS's exact auction and GS's auction), we iteratively build the active set (with excessive demand) by soliciting bidder information on an "as-needed" basis, thus economizing on information solicitation. Unlike DGS's approximate auction,

of his current item. We can always find a path of replacement because our way of constructing the active set. The "chain" replacement always ends in finite steps because at least one item in the active set has two or more bidders. The replacement step ensures that the number of bidders at an item will never drop to zero.

		Bidders for item				Active	Marginal	
t	Prices	1	2	3	0	items	bidder	What happens next
0	$10, 20, 15$	1,2	5	3,4	-	1	2	2 reassigned to item 3
0	$10, 20, 15$	1	5	2,3,4	-	3	2	2 links items 1 & 3
0	$10, 20, 15$	1	5	2,3,4	-	1,3	4	4 reassigned to item 2
0	$10, 20, 15$	1	4,5	2,3	-	2	4	4 links items 2 & 3
0	$10, 20, 15$	1	4,5	2,3	-	2,3	4	4 reassigned to item 1
0	$10, 20, 15$	1,4	5	2,3	-	1	4	4 links items 1, 2, & 3
0	$10, 20, 15$	1,4	5	2,3	-	1,2,3	-	Clocks 1, 2, & 3 are on
1	$11, 21, 16$	1,4	5	2,3	-	1,2,3	1	1 drops out
1	$11, 21, 16$	4	5	2,3	1	3	2	2 reassigned to item 1
1	$11, 21, 16$	2,4	5	3	1	1	2	2 links items 1 & 3
1	$11, 21, 16$	2,4	5	3	1	1,3	4	4 reassigned to item 2
1	$11, 21, 16$	2	4,5	3	1	2	4	4 links items 1 & 2
1	$11, 21, 16$	2	4,5	3	1	1,2	2	2 links item 1, 2, & 3
1	$11, 21, 16$	2	4,5	3	1	1,2,3	-	Clocks 1, 2, & 3 are on
2	$12, 22, 17$	2	4,5	3	1	1,2,3	5	Cocks are off; 5 drops out
2	$12, 22, 17$	2	4	3	1,5	-	-	Auction ends

Table 2: An Illustration of the Auction Process

we may pivot or expand the active set before raising prices so that sincere bidding is maintained.

Our design prevents cycling: an active bidder can move multiple times during a reassignment phase but cannot make a round trip because once the bidder becomes marginal, his indifferent items are linked so that he cannot go back to them (refer to step 2d). Similarly, bidders who drop out can never re-enter because the dummy item is never "active" in our design.[5]

3 Auction Properties

We represent an allocation using a partition of items $x^t = \left(x_0^t, x_2^t, ..., x_n^t\right)$, where x_i^t represents the item(s) assigned to bidder i at the end of iteration t. Similarly, we may also represent an allocation using the partition of bidders $\delta^t = \left(\delta_1^t, \delta_2^t, ..., \delta_m^t\right)$, where δ_j^t denotes the set of bidders assigned to item j at t.

Given a price vector p, we define the demand set for a bidder i as

$$D_i(p) = \arg\max{}_{j \in \mathcal{J}}(v_{ij} - p_j), \forall i = 1, 2, ...n.$$

[5]We note that the dummy cannot enter the active set via linking. The argument is as follows. We suppose that it has not entered the active set so far. The only way for it to enter the active set is via linking, which would require that the current margin bidder has previously been assigned to the dummy item and left it - a contradiction.

3.1 Sincere Bidding

The auction requires two kinds of actions from bidders: signaling an intent to submit a new bid and actual submission of a new bid (including the initial bid).

Definition 1 ' ***(Sincere Bidding)*** *A bidder bids sincerely if his actions are consistent with his demand set, that is, he signals an intent to bid if and only if there is an inactive item in his demand set, and his new bid is always an inactive item from his demand set. A bidder is sincere, if he always bids sincerely.*

An auction's final outcome is an allocation-price pair (x, p). Whether bidders are sincere or not, an auction outcome (x, p) must satisfy the following three conditions:

1. x is feasible (i.e., each bidder is assigned to a single item (include dummy) and no two bidders are assigned to the same non-dummy item).

2. $\forall j \in \mathcal{J}$, $p_j \geq p_j^0$ (i.e., final prices are no lower than starting prices).

3. $\forall j \in \mathcal{J}, \delta_j = \emptyset \Rightarrow p_j = p_j^0$ (i.e., unassigned items must be at their starting prices).

The first two conditions are intuitively true. The third condition is supported by the following lemma.

Lemma 1 *Unassigned items must be at their starting prices.*

Proof. We first argue that once an item gets its first bidder, it will maintain at least one bidder in subsequent bidding. By our auction rule, only an active item may lose a bidder. If an active item has two or more bidders, the argument holds trivially. If an active item has only one bidder, which can happen only when it becomes active via linking, our auction procedure as described in Footnote 4 ensures that it maintains at least one bidder. Thus, an unassigned item must have no prior bids. Because an item with no prior bids has never been active and the auctioneer only increases prices of active items, an unassigned item must be at its starting price. ■

To examine the incentives for bidders to bid sincerely, we need to evaluate bidders' payoffs in the final auction outcomes. In order to compare auction outcomes, we first need to know that the auction ends in a finite number of iterations t. Clearly, if there are two or more insincere bidders, the auction can go on indefinitely (though not all insincere bidders can profit from the situation). However, when there is at most one insincere bidder, we have the following:

Lemma 2 *If there is at most one insincere bidder, the auction must end in a finite number of iterations.*

Proof. We first argue that when all bidders are sincere, then the auction must stop at a finite t. Suppose it does not. Because at any t at least one price should increase, as $t \to \infty$, there must exist one item j whose price is high enough such that no sincere bidder would bid on this item. But the price of an unassigned item must equal its starting price (Lemma 1), a contradiction. Now we consider all but one bidder bid sincerely. Again, for a large enough t, at least one item has a price high enough and it exceeds all bidders' valuations. Since the item must be assigned and any sincere bidder must have left the item already, the only remaining bidder on this item must be the insincere one. But with just one bidder on this item, the price of this item will never rise again. So the auction of the rest of the items, which are participated by only the sincere bidders, must end in finite iterations. ■

We additionally note that not only the auction must end in a finite number of iterations, there can be only a finite number of reassignments during any reassignment phase. This is because our activity rule (the linking mechanism) prevents the marginal bidder from bid on items he has previously bid on.

For a given set of valuations and starting prices, let Ω denote the set of all possible outcomes when all bidders are sincere, and Ω_{-1} denote those when all but bidder 1 are sincere. Because the insincere bidder can also choose to bid sincerely, Ω is a subset of Ω_{-1}. We further denote x_I as the allocation of a subset of bidders I. For example, $x_{(1,2,3)} = (3, 2, 5)$ means that bidders 1, 2, and 3 are assigned to items 3, 2, and 5 respectively.

Lemma 3 *Consider two auction outcomes $(x, p) \in \Omega$ and $(x', p') \in \Omega_{-1}$. If bidder 1 strictly prefers (x', p') to (x, p), there must be a set of bidders I such that 1 is part of I, x_I is a permutation of x'_I, and*

$$p'_j < p_j, \forall j \in x_I \tag{1}$$

Proof. We discuss two cases: (a) 1 is assigned to the same item, say j, in both outcomes and (b) 1 is not. In case (a), we let $I = 1$. Because 1 strictly prefers (x', p') to (x, p), we clearly have $p'_j < p_j$. We now consider case (b). Suppose 1 is assigned to j_1 in x and j_2 in x' ($j_1 \neq j_2$). Because bidder 1 strictly prefers (x', p') to (x, p) and j_1 is a sincere bid at p, we have $v_{1j_1} - p_{1j_1} < v_{1j_2} - p'_{j_2}$ and $v_{1j_2} - p_{j_2} \leq v_{1j_1} - p_{1j_1}$, which implies

$$p'_{j_2} < p_{j_2}. \tag{2}$$

By Lemma 1, j_2 must be assigned in (x,p) because $p_{j_2} > p'_{j_2} \geq p^0_{j_2}$. Let C be the winner of item j_2 in (x,p). Because the price of j_2 is lower in (x',p'), C must not drop out and be a winner of some item j_3 in (x',p'). The fact that C bids sincerely at p and p' (C is not 1) implies $v_{C,j_2} - p_{j_2} \geq v_{C,j_3} - p_{j_3}$ and $v_{C,j_3} - p'_{j_3} \geq v_{C,j_2} - p'_{j_2}$. So

$$p'_{j_2} - p'_{j_3} \geq v_{C,j_2} - v_{C,j_3} \geq p_{j_2} - p_{j_3} \Rightarrow p'_{j_3} - p'_{j_2} \leq p_{j_3} - p_{j_2}.$$

Combining this with (2), we have

$$p'_{j_3} < p_{j_3} \tag{3}$$

Similarly, we can find another bidder who is winner of j_3 in (x,p) and a winner of another item at (x',p'). Since we have a finite number of items, we must end up with a bidder who is a winner of item j in (x,p) - thus we have a permutation. By construction, every item in this permutation has a lower final price in (x',p') than in (x,p). ■

Lemma 3 says that for an insincere bidder to profit from manipulating the auction's outcome, he must be part of a group that obtains the same items (possibly with permutation) at strictly lower prices (compared to an unmanipulated outcome). Such a group of bidders cannot exist as per the next lemma.

Lemma 4 *Consider two auction outcomes $(x,p) \in \Omega$ and $(x',p') \in \Omega_{-1}$, bidder 1 weakly prefers (x,p) to (x',p').*

Proof. We prove it by contradiction. Suppose the two outcomes result from auctions A and B respectively and 1 strictly prefers (x',p') to (x,p). By Lemma 3, there must be a subset of bidders I, such that $1 \in I$, x'_I is a permutation of x_I, and $p'_j < p_j, \forall j \in x_I$.

Case 1: Suppose I includes all the assigned bidders in x. Consider the last reassignment involving items x_I in auction A. There are only two possible subcases.

a) A bidder i chooses dummy over item $j \in x_I$ at price p_j and all other items at their respective prices. Since bidder i bids sincerely, i is indifferent between j and 0 at this price, which means i strictly prefers j to 0 when j's price is $p'_j < p_j$. Since i is not a winner of any item in x'_I in auction B and he prefers j to 0 under p', he must be a winner of another item, say $k \notin x'_I$, and strictly prefer k to 0 under p'. This contradicts the fact that i chooses 0 over k in auction A when $p_k = p^0_k \leq p'_k$.

b) The last reassignment involves a bidder i moving into an unassigned item $j \in x_I$. By Lemma 1, item j must still have a price of $p_j = p^0_j$, contradicting $p_j > p'_j \geq p^0_j$.

Case 2: Suppose I does not include all the assigned bidders in (x,p). Consider the last reassignment involving items x_I in auction A. By b), the last reassignment cannot be some bidder moving into an unassigned item in x_I. The only other possibility is a bidder leaving x_I. Suppose bidder i is the last bidder leaving x_I. Suppose further that he leaves item $j_1 \in x_I$ at time t of auction when j's price is p_{j_1}. Because $p'_{j_1} < p_{j_1}$, i must be a winner of some item under p', say item j_2. By sincere bidding of i (recall that $1 \in x_I$ and $i \notin x_I$) at two auctions, we have $v_{i,j_2} - p'_{j_2} \leq v_{i,j} - p'_{j_1}$ and $v_{i,j_1} - p_{j_1} \geq v_{i,j_2} - p^t_{j_2}$ respectively. Because $p'_{j_1} < p_{j_1}$ and $p^t_{j_2} \leq p_{j_2}$ (no price-drop), we can infer:

$$p'_{j_2} < p_{j_2} \tag{4}$$

From this point on we can repeat the argument of Lemma 3 and obtain a new set of bidders $I_1 = I \cup \Delta$, $i \in \Delta$, such that x'_{I_1} is a permutation of x_{I_1} and $p'_j < p_j, \forall j \in x_{I_1}$. By induction and the fact that there are a limited number of bidders, we must end with case 1 - a contradiction. ■

Intuitively, the group described in the Lemma 3 cannot exist because there is always a "last person" problem: in order to depress prices for the items assigned to the group, eventually one needs an outsider who would prematurely exit the competition, which a sincere bidder would not do. Now Lemma 4 immediately yields the following proposition.

Proposition 1 *Bidding sincerely is an ex-post Nash equilibrium.*

We note that our proof of sincere bidding does not depend on the actual path of our auction process despite the fact that the final allocation can be path dependent.

3.2 Efficiency

Starting the auction with arbitrary prices might lead to an inefficient outcome. However, our auction can achieve efficient outcomes for a specific set of starting prices:

Proposition 2 *When the starting prices are the auctioneer's reservation valuations, the final outcome is efficient.*

Proof. We reformulate the assignment problem by replacing the dummy bidder 0 with m dummy bidders $\{n+1, ..., n+m\}$ and the dummy item with n dummy items $\{m+1, ..., m+n\}$. With the reformulated assignment problem, any feasible allocation can be represented by a permutation of bidders.

Consider a sincere auction outcome $(x, p) \in \Omega$. We can easily verify that, when the starting prices are set to the auctioneer's reservation valuations, the condition

$$x_i \in D_i(p), i \in \{1, 2, .., m, m+1, m+2, .., m+n\} \tag{5}$$

holds for both true and dummy bidders. In other words, every bidder, including dummy ones, weakly prefers his assignment.

We now claim that any cyclic permutation of x weakly decreases efficiency. We suppose that $i_1, i_2, ..., i_l$ are originally assigned to items $j_1, j_2, ..., j_l$ but a cyclic permutation reassigns them to $j_2, j_3, ..., j_l, j_1$. By condition (5), we know

$$\begin{aligned} v_{i_1 j_1} - p_{j_1} &\geq v_{i_1 j_2} - p_{j_2} \\ v_{i_2 j_2} - p_{j_2} &\geq v_{i_1 j_3} - p_{j_3} \\ &... \\ v_{i_l j_l} - p_{j_l} &\geq v_{i_l j_1} - p_{j_1} \end{aligned}$$

Adding two sides of inequations, we have

$$\sum_{k=1..l} v_{i_k j_k} \geq \sum_{k=1..l-1} v_{i_k j_{k+1}} + v_{i_l j_1}$$

The left (right) hand side is the efficiency of the original (permuted) allocation, implying that the cyclic permutation weakly decreases efficiency. Since any permutation can be decomposed into several disjoint cyclic permutations, we conclude that any feasible allocation other than x weakly decreases efficiency. ∎

3.3 Efficiency, Revenue, and Ending Prices

We now show the connection between our auction and the VCG mechanism.

Proposition 3 *If an auction starts with the auctioneer's reservation values, the bidders' final payments coincide with their VCG payments for the same final allocation.*

Proof. Consider an auction outcome (x, p). Without loss of generality, we assume bidder 1 is assigned to $j_1 \in \mathcal{J}$ under x. We consider an alternative economy where bidder 1 is excluded from participation. The VCG payment for bidder 1 can thus be calculated as the difference in social welfare of all other players in the original and alternative economies. We compare auction payments and VCG payments for two cases: (a) $p_{j_1} > p^0_{j_1}$ and (b) $p_{j_1} = p^0_{j_1}$.

Case (a): We construct a new efficient allocation $\hat{x}$ in the alternative economy and use it to calculate VCG payment for bidder 1. There must be a last bidder who leaves j_1 at price p_1. Suppose this bidder is i_2 and assigned to $j_2 \in \mathcal{J}$ under x. If $p_{j_2} > p^0_{j_2}$, we continue to search for the last bidder who leaves j_2, and so on. Eventually we can find a series of bidders $\{i_2, i_3, ..., i_k\}$ assigned to $\{j_2, j_3, ..., j_k\}$ such that i_l is the last bidder to leave j_{l-1}, for $l = 2..k$ and the price of item j_l is strictly above its starting price except $p_{j_k} = p^0_{j_k}$. We now construct a new allocation $\hat{x}$ under which every bidder in the series $\{i_2, i_3, ..., i_k\}$ takes the place of the preceding bidder (e.g., bidder i_2 is assigned to item j_1), i_1 is unassigned, and all other bidders keep their assignments. For bidder i_2, we know by construction that (t is the moment of the last bidder departure at j_1)

$$v_{i_2,j_1} - p_{j_1} \geq v_{i_2,j_2} - p^t_{j_2} \geq v_{i_2,j_2} - p_{j_2} \tag{6}$$

where the first inequality is by sincere bidding and the second inequality is because of nondecreasing prices. Hence, given p, i_2 would prefer his assignment j_1. The same argument can be made for other bidders in the series $\{i_2, i_3, ..., i_k\}$. Overall, we conclude that $\hat{x}_i \in D_i(p)$ holds for all bidders and by Proposition 2, $\hat{x}$ is efficient in the alternative economy.

By sincere bidding in the original economy, we also have $v_{i_2,j_1} - p_{j_1} \leq v_{i_2,j_2} - p_{j_2}$. Combining this with (6), we have

$$v_{i_2,j_1} - p_{j_1} = v_{i_2,j_2} - p_{j_2} \tag{7}$$

The equal (7) is also true for other bidders in the series. Summing up all equations and rearranging terms, we have

$$\left(v_{i_2,j_1} + v_{i_3,j_2} + ... + v_{i_k,j_{k-1}} + p_{j_k}\right) - \left(v_{i_2,j_2} + v_{i_3,j_3} + ... + v_{i_k,j_k}\right) = p_{j_1}$$

Note that $p_{j_k} = p^0_{j_k} = v_{0,j_k}$, so the first parenthesis represents the social welfare of bidders $\{i_2, i_3, ..., i_k\}$ and the auctioneer in the alternative economy and the second parenthesis, the original economy. So the left hand side is exactly the VCG payment of bidder 1 for allocation x, suggesting that the auction payment p_{j_1} for bidder 1 coincides with his VCG payment for the same allocation.
■

The above proposition shows that when starting prices are equal to the auctioneer's reservation valuations, the auction implements the VCG mechanism. This special case of our auction can be therefore viewed as a decentralized implementation of the VCG mechanism.

We now turn to the properties of auction revenue and ending prices. Our numerical experiments show that the auctioneer may obtain higher revenues by setting starting prices above his true reservation valuations, though this may cause allocative inefficiency. If the auctioneer has some knowledge of the distribution of bidders valuation, he may use that information to choose optimal starting prices to maximize his expected revenue. In this sense, our auction generalizes the classic English auctions with reserve prices.

Now, given the starting prices, can the auction end with different ending prices or revenues? We note that, because of tie-breaking by bidders and by the auctioneer, the auction may end with different sets of winners and/or the same bidder winning different items. So it is not automatically clear whether the ending prices and the auction revenue are path-dependent. The following result establishes that auction ending prices and revenue are path independent.

Proposition 4 *If (x, p) is a final outcome of an auction, then p is the smallest price vector in Ω.*

Proof. Because $(x, p) \in \Omega$, by Lemma 4, bidder i weakly prefers (x, p) to any outcome in Ω_{-1}. Since $\Omega \subseteq \Omega_{-i}$, the bidder must also weakly prefer (x, p) to any other $(x', p') \in \Omega$. To see that p is indeed the smallest price vector in Ω, we discuss two cases. (a) If bidder i is assigned to the same item in x and x', $p_{x_i} \leq p'_{x_i}$ holds trivially. Otherwise, say, (b) i is assigned to item 1 under x but 2 under x'. We have

$$\begin{aligned} v_{i,2} - p'_2 &\geq v_{i,1} - p'_1 \\ v_{i,1} - p_1 &\geq v_{i,2} - p'_2 \end{aligned}$$

where the first inequality is because item 2 must be weakly preferred by i under (x', p') and the second inequality is because i weakly prefers (x, p) to (x', p'). These two inequalities imply $p_1 \leq p'_1$. Hence, any assigned item must have the lowest price in Ω. If the item is unassigned, its price is the starting price, which is the lowest possible price in Ω. Overall, we conclude that p is the lowest price in Ω. ∎

Proposition 4 implies that our auction must end with a unique set of prices (since there can only be one smallest price vector Ω). Even though the auction may end with different allocations, the final price vector is always the same (path independence). Therefore, the auction revenue is also unique and path-independent.

A further question is whether identical items (i.e., all bidders have the same valuations for these items) fetch the same price in our auction (e.g., for fairness reasons). The next result shows that identical items with the same starting price will in fact have the same ending price.

Lemma 5 *If two items are identical and have the same starting price, they must have the same ending price.*

Proof. Suppose two items 1 and 2 have different final prices. Without loss of generality, we assume $p_1 > p_2$. Because their starting prices are the same, so 1 must be assigned to a bidder, say, bidder 1. By sincerity, we have $v_{11} - p_1 \geq v_{12} - p_2$. But since $v_{11} = v_{12}$, we have $p_1 \leq p_2$, a contradiction. ∎

4 Measuring Bidder Privacy

By bidder privacy, we specifically mean the amount of information gets revealed about a bidder's private valuation of items. So far we have argued that because our auction only collects demand information as needed, therefore it should better preserve bidder privacy. But in a dynamic auction, it would as well be the case that, due to additional rounds of reporting, that overall information revealed is the same or even more than other auctions. Additionally, due to the randomness in the auction path, the amount of information revealed in the same auction could vary depending on the auction path. So it is not practical to analytically draw conclusions about the relative amount of two auctions when both can reveal variable amount of bidder information. Hence, we resort to numerical simulation methods for comparing amount of information revealed.

Despite the importance of bidder privacy preservation in practical auction design, there is no formal measure of bidder privacy in the auction literature. So our first step towards studying the bidder privacy aspect of auction design is to formalize the notion of bidder privacy. A key idea is to compare posterior value distributions. For example, in sealed bid first-price auctions, the auctioneer knows exactly all bidders' values after the auction, but he only knows a value range for a winner of an English auction. By measuring how much more precise the value distribution has become after the auction, we can quantify the information revealed by the auction. For this purpose, we choose Shannon's information entropy, which measures the amount of information contained in a probability distribution as information entropy. The entropy for the random variable $x \in X = \{x_1, x_2, ..., x_k\}$, with probability function $P(x)$ is given by

$$-\sum_{i=1}^{k} P(x_i)\, log_2 P(x_i)$$

Definition 2 ***(Amount of Information Revealed)****: Let $P(\mathbf{v})$ and $P'(\mathbf{v})$ be the probability functions of a bidder's valuation vector before and after a mechanism. We define the amount of information revealed about this bidder as the entropy reduction, i.e.,*

$$\sum_{\mathbf{v}_i \in \mathbf{V}} \left[P'(\mathbf{v}_i)\, log_2 P'(\mathbf{v}_i) - P(\mathbf{v}_i)\, log_2 P(\mathbf{v}_i)\right]$$

In the special case where the k dimensions of bidder valuation are independently distributed, the above can be rewritten as:

$$\sum_{i=1}^{k} \sum_{v_i \in V_i} \left[P'_i(v_i)\, log_2 P'_i(v_i) - P_i(v_i)\, log_2 P_i(v_i)\right]$$

where V_i and $P_i(\cdot)$ are the support and the probability function of the i'th dimension respectively.

4.1 Determine the Posterior Value Distribution

In ascending auctions, at any time, a bidder i's bid for item j reveals that

$$v_{ij} - p_j \geq v_{il} - p_l, \forall l \in I_1,\ l \neq j \tag{8}$$

If the bidder is indifferent between two items j and k, 8 should hold for both j and k.

The amount of information revealed is generally path dependent. For example, bidding on item 1 all the way contains different information than first bidding on item 2, then switching to item 1. The following Lemma specifies all the events that need to be tracked for the purpose of the capturing all the information revealed.

Lemma 6 *In our auction, the total information revealed is the sum of information revealed at the moments of the initial bid, subsequent bid revisions, and the final bid, at their corresponding prices.*

Proof. Given the private valuation assumption, if there is no price increase and the bidder does not change his bid, there is no new information about the bidder. If the price increases (and the bid remains the same), information revealed at a higher price overrides that at a lower price. A combination of initial bids, intermediate bid revisions, and final bid capture all the bidder states at their highest price points. ■

With this Lemma and condition 8, we only need to record a series of inequations (such as 8) at informational events. Note that at event of a bid revision, we need to record a set of inequations for both items before and after the revision.

Let $C \equiv \{(j_k, p_k)\}_{k=1}^{K}$ denote a set of item-price-vector pairs captured by all the informative events. The posterior probability is given by the conditional probability function:

$$P'(\mathbf{v}) = P\left(\mathbf{v} \middle| \left\{ \left\{ v_{ij_k} - p_j^k \geq v_{il} - p_l^k, \forall l \in I_1,\ l \neq j \right\}, (j_k, p_k) \in C \right\} \ and\ \mathbf{v} \in \mathbf{V} \right)$$

4.2 Calculating Entropy

Analytically solving the updated probability distribution, given a set of linear constraints, maybe feasible in some cases, but not for general cases. This problem amounts to a high-dimension integration over a convex feasible region (as defined by the linear constraints), a problem known to be difficult even for numerical methods of integration due to the curse of dimensionality. We resort to a Monte Carlo simulation approach instead.

The basic idea of the Monte Carlo simulation method is as follows. We draw N random samples according to the prior distribution of random vector. We test each sample against all the constraints. If all constraints are satisfied, we record 1 ("good sample"), otherwise 0. The proportion of 1's approximates the probability P of the random vector falls into the feasible region. For the uniform distribution, the entropy reduction due to the added constraints is simply $log_2\left(P\right)$, which is approximated by:

$$log_2\left(N\right) - log_2\left(\#good\ samples\right)$$

The idea seems simple, but in order to obtain a reliable estimate of entropy, lots of draws (sometimes more than 20,000) are often needed, largely because the MC method is slow in convergence. To accelerate convergence, we adopt two strategies. First, whenever there are equal constraints such as $v_1 - v_2 = 3$, a degree of freedom is lost, and a lot of draws are required to meet the equality constraint, which results in low convergence. To speed it up, we decompose the value vector into free and derived variables. By construction, the derived variables can be mathematically derived using a system of linear equations. We can then simulate the free variables, and analytically evaluate the derived variables. The entropy reduction among derived variables can be directly calculated numerical without simulation, and combined with the simulated entropy reduction for the free variables.

The second strategy is to use quasi-MC method instead of regular MC method. The former is popular in computational finance for calculating high-dimensional integration (Boyle et al., 1997). quasi-MC method works the same way as regular Monte Carlo simulation, but the points of integration are given by a deterministic low-discrepancy sequence rather than the usual pseudo-random sequence. We use the popular Sobol sequence, which is known to achieve high convergence rate at high dimensions compared with the pseudo-random sequence in regular MC(Morokoff and Caflisch, 1995).

Now we describe an algorithm for calculating the entropy for our auction and other auctions. Let C_i be a set of constraints gathered across all relevant informational events for bidder i.

1. For each bidder i, given C_i, extract equality constraints, and use them to split a set of items into F and D such that any value $v_{ij},\ j \in D$, can be derived from values $\{v_{bj}, j \in F\}$ via a system of linear equations. Let $n' \leftarrow \|F\|$ denote the number of free variables.

2. Draw N random size-n' vectors $\mathbf{v}^F$ according to the prior distribution.

3. Total $\leftarrow$0

4. For each random vector $\mathbf{v}^F$, derive $\mathbf{v}^D = \{v_{bj}, j \in D\}$ by solving a system of linear equations. Tests $\mathbf{v}^F$ against constraints in C_i and the lower/upper bounds for elements in $\mathbf{v}^F$. Also test the lower/upper bounds for values in $\mathbf{v}^D$. If all conditions are met, $Total \leftarrow Total + 1$.

5. $EntropyReduction_i \leftarrow log_2(N) - log_2(Total) + Entropy\ Reduction\ for\ \mathbf{v}^D$

We note that we analytically evaluate the probability of drawing a specific $\mathbf{v}^D$ (e.g. 1/N for a single uniformly distributed discrete variable), and use it to calculate the corresponding entropy reduction. Hence, our method of calculating entropy is a hybrid of analytical and simulation methods.

4.3 Numerical Simulations and Results

In this section, we compare our auction with two benchmark mechanisms, i.e. multi-item VCG and the DGS auction. For the DGS exact auction, Sankaran (1994) points out it is unnecessary to calculate the minimal over-demanded set in each iteration, which is a computational expensive. He shows that calculating an over-demanded set is sufficient to ensure the same outcome. We therefore have implemented Sankaran (1994)'s version of improved DGS exact auction. We note that Lemma 6 can be applied to the DGS auction by considering each reported demand set as a set of bids. For the VCG auction, since all the information is revealed via sealed truthful bid, we derive the entropy reduction analytically instead of using simulation (a simple experiment shows that simulation results are close to the analytical result).

We run a series of numerical experiments on the three mechanisms, our auction, DGS auction, and VCG. In all of these experiments, we assume the bidder values are independently and uniformly distributed. We also set the value of outside options and auctioneer reservation values to zero for simplicity. Unless otherwise noticed, we draw integer bidder values from a uniform distribution between 0 and 20. For each parameter set, we simulate 100 random scenarios. For each scenario, we run three mechanisms, and record total entropy reduction across all bidders, record the number of iterations, running time. We use 5000 draws for all Monte Carlo simulation as our experiments show that this is a reasonable number. The experiments are run on MacBook Pro 2014 with 16Gb RAM.

In the first set of analysis, we let the number of items be one less than the number of bidders, representing a high item/bidder ratio (low competition) scenario. We then vary the total number of items. As we can see from Table 3, our auction consistently results in about 10% less entropy reduction than DGS, suggesting its privacy preservation advantage. Both dynamic auctions show significant less entropy reduction than VCG, and the gain is more significant as the number of items

		Entropy Reduction					Iterations				Running Time (sec)				
m	n	our	DGS	t	p	VCG	our	DGS	t	p	our	DGS	t	p	VCG
3	2	12.24	13.4	-2.96	**	26.35	3.08	3.08	0		0.002	0.002	4.02	***	0.001
5	4	26.32	29.09	-3.91	***	87.85	5.3	5.18	0.80		0.007	0.003	18.46	***	0.001
7	6	41.52	45.55	-3.40	***	184.5	7.44	7.08	1.43		0.019	0.005	12.42	***	0.001
9	8	55.68	61.12	-4.03	***	316.2	8.61	8.3	1.17		0.035	0.008	26.29	***	0.002
11	10	71.73	80.21	-5.54	***	483.2	10.41	9.66	2.28	*	0.058	0.011	23.91	***	0.002

***: $p < 0.001$, **: $p<0.01$, *: $p<0.05$

Table 3: Numerical studies - high item/bidder ratio (low competition) scenarios

		Entropy Reduction					Iterations				Running Time (sec)				
m	n	our	DGS	t	p	VCG	our	DGS	t	p	our	DGS	t	p	VCG
3	2	12.1	13.61	-3.50	***	26.35	3.05	3.05	0		0.002	0.001	7.04	***	0.001
5	3	29.86	31.51	-2.36	*	65.88	6.55	5.75	3.992	***	0.007	0.003	14.62	***	0.001
7	4	48.06	52.33	-3.75	***	123	11.75	8.92	7.115	***	0.018	0.007	9.344	***	0.001
9	5	68.49	76.47	-5.27	***	197.7	18.13	11.85	10.76	***	0.037	0.009	17.64	***	0.001
11	6	91.32	99.35	-4.75	***	289.9	25.42	14.17	14.66	***	0.064	0.013	18.72	***	0.003

Table 4: Numerical studies - low item/bidder ratio (high competition) scenarios

increase (85% less at 10 items). This confirms the benefits of dynamic auctions in preserving bidder privacy. Our auction uses comparable number of iterations as DGS auctions, but its running time is longer than DGS, which makes sense because it iteratively requests for bids within an iteration to find over-demanded items. Unsurprisingly, both dynamic auctions are significantly slower than VCG.

In the next set of experiments, we keep the number of items approximately one-half of the number of bidders, simulating a low item/bidder ratio (high competition) situation. As shown in Table 4, the number of iterations is noticeably higher, but our auction maintains a roughly 10% less entropy reduction than DGS. Both dynamic auctions show significant less entropy reduction compared with VCG, but the gaps are noticeably smaller, presumably because competition reveals more information about bidder values.

In the third set of experiments, we fix the number of items and bidders to 5 and 9 respectively, but vary the support of the uniform distribution, thereby varying the degree of uncertainty about bidder valuation. As the support increases, both dynamic auctions use more iterations and longer running time. At a narrow support (0-5), our auction enjoys a 14% less entropy reduction than DGS but that advantage decreases as the support gets wider: at [0, 200], there is no significant difference. This is because the privacy advantage of our auction is most pronounced when the bidder has more than two items in the demand set, which is less likely when the support is very wide.

	Entropy Reduction					Iterations				Running Time				
Support	our	DGS	t	p	VCG	our	DGS	t	p	our	DGS	t	p	VCG
[0,5]	59.15	68.92	-9.09	***	116.3	7.34	6.27	4.04	***	0.019	0.005	24.9	***	0.001
[0,10]	69.62	79.24	-7.71	***	155.7	11.76	9.14	6.83	***	0.023	0.007	13.2	***	0.001
[0,20]	84.46	90.26	-4.01	***	197.7	16.43	11.53	8.22	***	0.029	0.009	14.1	***	0.001
[0,40]	98.34	102.2	-2.69	**	241.1	21.64	13.63	9.95	***	0.034	0.010	19.2	***	0.001
[0,100]	114.9	120.0	-2.68	**	299.6	23.51	15.15	9.62	***	0.035	0.011	18.4	***	0.001
[0,200]	130.3	128.4	1.01		344.3	26.60	14.80	13.6	***	0.040	0.011	16.4	***	0.001

Table 5: Numerical studies - degree of uncertainty

Interestingly, as the support grows from 5 to 200, the bidder privacy benefits of the two dynamic auctions grow stronger, from 50% less entropy reduction to 62% less than the VCG. This may be caused by greater amount of uncertainty in the bidder valuation, thus more room for bidders to hide their valuations through dynamic auctions.

5 Conclusion

Privacy and transparency issues affect bidders' participation, which in turn affects the probability of sale and revenues. In this research, we have presented a new ascending auction design that can mitigate privacy concerns of bidders and improve the process transparency by economizing on information solicited from bidders, making price adjustment processes transparent, and accommodating explicit starting prices. In this auction, sincere bidding is an ex-post Nash equilibrium, and given a set of valuations, ending prices depend only on starting prices. The starting prices of our auction can be set strictly above the reservation values of the auctioneer, which is important to accommodate auctioneers' revenue goals. However, if the starting prices are set equal to the reservation values, our auction implements the VCG mechanism. Our numerical study using a newly defined bidder privacy measurement shows that our auction routinely reduces amount of information revealed than DGS auctions and both dynamic auctions significantly reduce information revelation than the VCG auction.

There are several directions for extending the current work. First, it would be useful to conduct complexity analysis on the proposed auction design. Second, we have not studied the heuristics of choosing auction paths, which can have a sizable impact on both bidder privacy and speed of convergence. Third, our numerical studies could be further improved both in terms of systematic exploration of parameter space and in terms of the simulation techniques for better performance. Finally, it will be interesting to see how our approach to dynamic auctions could be extended to more complex settings, such as non unit-demand settings and team-based assignments, and for applications in online environments such as online labor markets.

References

Adomavicius, G. and Gupta, A. (2005). Toward Comprehensive Real-Time Bidder Support in Iterative Combinatorial Auctions. *Information Systems Research*, 16(2):169–185.

Andersson, T., Andersson, C., and Talman, A. J. J. (2013). Sets in excess demand in simple ascending auctions with unit-demand bidders. *Annals of Operations Research*, 211(1):27–36.

Ausubel, L. (2004). An efficient ascending-bid auction for multiple objects. *The American Economic Review*, 94(5):1452–1475.

Ausubel, L. M. (2006). An efficient dynamic auction for heterogeneous commodities. *The American Economic Review*, 96(3):602–629.

Ausubel, L. M. and Milgrom, P. R. (2002). Ascending Auctions with Package Bidding. *Frontiers of Theoretical Economics*, 1(1).

Bapna, R., Das, S., Day, R., Garfinkel, R., and Stallaert, J. (2011). A Clock-and-Offer Auction Market for Grid Resources When Bidders Face Stochastic Computational Needs. *INFORMS Journal on Computing*, 23(4):630–647.

Bikhchandani, S. and Ostroy, J. M. (2002). The Package Assignment Model. *Journal of Economic Theory*, 107(2):377–406.

Boyle, P., Broadie, M., and Glasserman, P. (1997). Monte Carlo methods for security pricing. *Journal of Economic Dynamics and Control*, 21(8-9):1267–1321.

Crawford, V. P. and Knoer, E. M. (1981). Job matching with heterogeneous firms and workers. *Econometrica*, 49(2):437–450.

de Vries, S., Schummer, J., and Vohra, R. V. (2007). On ascending Vickrey auctions for heterogeneous objects. *Journal of Economic Theory*, 132(1):95–118.

Demange, G., Gale, D., and Sotomayor, M. (1986). Multi-item auctions. *The Journal of Political Economy*, 94(4):863–872.

Gale, D. and Shapley, L. S. (1962). College admissions and the stability of marriage. *The American Mathematical Monthly*, 69(1):9–15.

Gul, F. and Stacchetti, E. (2000). The English Auction with Differentiated Commodities. *Journal of Economic theory*, 92(1):66–95.

Krishnappa, C. and Plaxton, C. G. (2011). A Dynamic Unit-Demand Auction Supporting Bid Revision. In *Proceedings of the 13th International Conference on Electronic Commerce.*

Leonard, H. (1983). Elicitation of honest preferences for the assignment of individuals to positions. *Journal of Political Economy*, 91(3):461–479.

Mishra, D. and Parkes, D. C. (2009). Multi-item Vickrey-Dutch auctions. *Games and Economic Behavior*, 66(1):326–347.

Morokoff, W. J. and Caflisch, R. E. (1995). Quasi-Monte Carlo Integration. *Journal of Computational Physics*, 122(2):218–230.

Perry, M. and Reny, P. J. (2005). An Efficient Multi-unit Ascending Auction. *Review of economic studies*, 72(2):567–592.

Petrakis, I., Ziegler, G., and Bichler, M. (2013). Ascending Combinatorial Auctions with Allocation Constraints: On Game Theoretical and Computational Properties of Generic Pricing Rules. *Information Systems Research*, 24(3):768–786.

Rothkopf, M. H., Teisberg, T. J., and Kahn, E. P. (1990). Why Are Vickrey Auctions Rare? *Journal of Political Economy*, 98(1):94–109.

Sankaran, K. (1994). On a dynamic auction mechanism assignment problem. *Mathematical Social Sciences*, 28(1994):143–150.