

Impact of Information Feedback in Continuous Combinatorial Auctions: An Experimental Study of Economic Performance

Gediminas Adomavicius, Shawn P. Curley, Alok Gupta

Carlson School of Management, University of Minnesota, Minneapolis, MN 55455
{gedas@umn.edu, curley@umn.edu, alok@umn.edu}

Pallab Sanyal

School of Management, George Mason University, Fairfax, VA 22030
psanyal@gmu.edu

Abstract

Advancements in information technology offer opportunities for designing and deploying innovative market mechanisms that can improve the allocation and procurement processes of businesses. For example, combinatorial auctions – in which bidders can bid on combinations of goods – have been shown to increase the economic efficiency of a trade when goods have complementarities. However, the lack of real-time decision support tools for bidders has prevented this mechanism from reaching its full potential. With the objective of facilitating bidder participation in combinatorial auctions, this study, using recent research in real-time bidder support metrics, discusses several novel feedback schemes that can aid bidders in formulating combinatorial bids in real-time. The feedback schemes allow us to conduct *continuous* combinatorial auctions, where bidders can submit bids at any time. Using laboratory experiments with two different setups, we compare the economic performance of the continuous mechanism under three progressively advanced levels of feedback. Our findings indicate that information feedback plays a major role in influencing the economic outcomes of combinatorial auctions. We compare several important bid characteristics to explain the observed differences in aggregate measures. This study advances the ongoing research on combinatorial auctions by developing continuous auctions that differentiate themselves from earlier combinatorial auction mechanisms by facilitating free-flowing participation of bidders and providing exact prices of bundles on demand in real time. For practitioners, the study provides insights on how the nature of feedback can influence the economic outcomes of a complex trading mechanism.

Key words: online auctions, continuous combinatorial auctions, experimental economics, information feedback.

Conditionally Accepted at MIS Quarterly

Impact of Information Feedback in Continuous Combinatorial Auctions: An Experimental Study of Economic Performance

1. INTRODUCTION

Rapid advances in computing and information technologies present the potential for deploying complex market mechanisms to achieve gains for all stakeholders involved. Combinatorial auctions, in which bidders can bid on combinations of goods (commonly referred to as *bundles* or *packages*), represent an innovative market mechanism that can enhance the ability to allocate multiple assets efficiently. While only a limited number of applications have been found so far, due to the complexity of this mechanism, combinatorial auctions have the potential to become more pervasive as theory develops and as information technology (IT) reduces the costs of computation and information processing that these auctions require.

To understand the benefits of selling multiple items together, consider an auction for the sale of a number of equally-sized property lots surrounding a lake. A potential buyer may be more interested in acquiring two adjoining lots rather than two separate lots, since having adjacent lots can provide more options for development. Consequently, she may be willing to pay more for the two lots together than separately. However, the sale of the lots in separate single-item auctions limits her ability to express her super-additive valuation, and also may lead to reduced revenue for the seller. Furthermore, the assets may not be won by persons who value them the most, resulting in loss of efficiency.¹

Thus, combinatorial auctions have received a considerable amount of attention in the research literature in recent years, and have been used or proposed to be used in some prominent

¹ Efficiency (or allocative efficiency) of auctions is maximized when goods are acquired by persons who value them the most. Higher efficiency is a desirable goal in auctions because it leads to greater social welfare.

applications, including the auctioning of the rights to use railroad tracks (Brewer and Plott 1996), delivery routes (Caplice 1996), spectrum rights (McAfee and McMillan 1996), airport time slots (Rassenti et al. 1982), and the procurement of school meals (Epstein et al. 2002). The compelling motivation for the use of combinatorial auctions in these cases was the presence of item complementarities that differed across bidders (Cramton et al. 2005). In such cases, a combinatorial auction can mitigate the so-called *exposure problem*² (Bykowsky et al. 2000) by allowing bidders to bid on an entire bundle of interest. Thus, these auctions have been found to increase the efficiency, seller's revenue, and bidders' willingness to participate when bidder valuations of multiple items are super-additive (Ledyard et al. 2002).

In spite of their benefits, one barrier to practical implementations of combinatorial auctions has been the complexity of bid evaluation. Adomavicius and Gupta (2005) proved that, while there can never be more than N provisionally winning bids at any stage in a combinatorial auction of N distinct items, it is possible to have up to $2^N - 1$ bids at any time that can potentially win the auction. As a result, determining winners in combinatorial auctions in general is computationally intractable, i.e., an *NP-hard* problem³ (de Vries and Vohra 2003).

² The exposure problem may occur when procuring multiple items through multiple single-item auctions if bidders have complementary (super-additive) valuations for the items. For example, if there are two items and a bidder has individual valuations of \$100 for each item but a joint valuation of \$250, then two single-item auctions may allocate goods to other bidders who have individual valuations greater than \$100 but combined valuation of less than \$250. If the bidder with complementary valuation bids too high on one of the items in the hope of winning the other item too, she exposes herself to a potential loss in the event that she is unable to win the second lot.

³ In computational complexity, NP-hard (non-deterministic polynomial-time hard) problems denote a class of problems that do not have an efficient computational solution (i.e., a solution with running time that is polynomial with respect to the problem size). See Garey and Johnson (1979) for more details.

Other complications in bid formulation also arise. In single-item ascending auctions (e.g., English auctions), if a bidder has been outbid, she needs to bid an amount higher than the current highest bid to stand a chance of winning the auction. However, in continuous combinatorial auctions this is not necessarily the case. A bid that is not among the current winners can be among the future winners based on the combinations of later bids on other items. For example, if we have a 2-item combinatorial auction for selling items A and B, and if the current bids are: (i) \$10 for {AB}; and (ii) \$5 for {A}, the second bid is not a current winner, assuming the auctioneer is maximizing her revenue (a standard assumption in the combinatorial auction literature). However, if a new bid of \$6 for {B} arrives, then bid (ii) becomes part of the currently winning bids.

Owing to these complexities, *continuous* versions of combinatorial auctions, similar to the classic English auction, are rare. However, several *iterative* solutions aimed at reducing the computational complexity have been introduced recently (e.g., Ausubel et al. 2005, Goeree and Holt 2008, Kwasnica et al. 2005, Parkes 1999, Porter et al. 2003). These approaches primarily focus on creating rules and restrictions to allow several well-defined rounds of bidding, with the auctioneer declaring the intermediate results after each round. In contrast, we are interested in facilitating a version of the combinatorial auction that is simple, transparent, optimal, efficient, and also *continuous*. A goal is to make the rules of the auction as simple as possible, so that the mechanism can be more easily deployed in online marketplaces for business-to-consumer (B2C) and consumer-to-consumer (C2C) exchanges. Our mechanism eliminates the need for an auctioneer's intervention and provides a simple, asynchronous, ascending-bid environment that is similar to traditional online ascending auctions, where bidders can potentially join an ongoing auction at any time and are able to find up-to-date information to facilitate their bidding. Our

primary objective in this paper is to test the feasibility of conducting such continuous auctions. Second, we study how the nature of feedback influences the economic performance of the combinatorial auctions. This is a novel approach since, to the best of our knowledge, no one has tested the effects of different levels of feedback within similar combinatorial auction designs, implicitly assuming that providing more information is better from an efficiency perspective. However, high efficiency can increase both the seller's revenue and bidders' surplus or any one of the two. Through the use of controlled experiments, we demonstrate how varying quality and quantity of information feedback affects various metrics of interest, such as efficiency, auction revenue, and bidder's surplus. We also analyze the bids to explain the differences in aggregate economic outcome as a function of feedback.

2. DESIGNING CONTINUOUS COMBINATORIAL AUCTIONS

Kwasnica et al. (2005) noted that winner determination no longer remains a major challenge for combinatorial auctions, since the availability of increased computing power permits computational solutions in reasonable time for reasonably-sized problems using commercial software such as CPLEX (see, for example, Andersson et al. 2000). Instead they identified bidder support as the major obstacle to combinatorial auctions reaching their potential. Pekeč and Rothkoph (2003) similarly noted that "bidtakers should take particular care in providing tools that help bidders in bid preparation" (p. 1501). To date, only limited bidder support techniques have been developed for combinatorial auctions. For example, Banks et al. (1989) created an Adaptive User Selection Mechanism (AUSM) in which it is the responsibility of the bidder to look at the existing bids in a bulletin board and submit her new bid in such a way that it becomes a part of the bid combination that provides maximal revenue for the auctioneer. Our research is focused on facilitating a *continuous* combinatorial auction environment by designing

a useful bidder support system that not only provides bidders with information regarding provisional allocation, but also guides bidders through appropriate and *exact* price feedback, without the intervention of the auctioneer.

The development of such an environment requires *real-time* bid evaluation. A major step towards that direction was taken by Adomavicius and Gupta (2005), who classified bids into categories and identified theoretical relationships among these categories. Based on the identified relationships, it became possible to define several novel constructs, such as *live* (non-losing) and *dead* (losing) bids, which help in developing efficient approaches towards providing real-time responses to various bid-related queries, such as: “*Is my bid currently winning?*” or “*Does my bid stand a chance of winning in the future even if it is not currently winning?*” We use these constructs to develop several feedback schemes to support a continuous combinatorial auction environment, where bids can be submitted at any time. If continuous auctions can be made practical, the use of combinatorial auctions can be greatly expanded beyond their current uses in iterative settings with restrictive rules and the need for an overseeing auctioneer. Bidder interactions with continuous auctions would be more similar to the commonly available and widely used ascending auctions for single items and, consequently, such auctions would become easier to adopt in the broader online marketplace for consumer-centric commerce.

The construction of the information infrastructure for real-time bidder support raises the issue of feedback and its impact upon the economic properties of the auction as well as the behavior of the auction participants. *Will bidders be able to properly interpret and synthesize all the information available to them in order to formulate efficient bids? Is more feedback regarding the state of the auction necessarily better for the bidder?* While the issue of the impact of information revelation policies on the outcome of auctions has been studied both

theoretically (e.g., Arora et al. 2007) as well as experimentally (e.g., Koppius and van Heck 2003) in the context of procurement, similar studies have not been conducted for continuous combinatorial auctions. The increased information revelation can be expected to decrease bidders' uncertainty, thereby increasing the efficiency of auctions by lowering the probability that assets go to the bidders with relatively lower valuations. However, how the gains of the trade from higher efficiency will be distributed between the seller and the buyers is not apparent from the auctions literature. *Will higher efficiency lead to higher seller's revenue, higher bidders' surplus, or both?* In a complex combinatorial bidding environment, the answer will likely depend on how bidders utilize the information provided to them via different forms of feedback. While a better understanding of the state of the auction can benefit the bidders in making better bids, greater transparency can also result in increased competition, thus benefitting auctioneers (Krishna 2002). To develop testable hypotheses concerning the effect of feedback on the auction outcomes, we turn to behavioral theory on how feedback affects performance. A primary goal of this study is to analyze the effects of progressively advanced levels of feedback on the economic performance of continuous combinatorial auctions in terms of efficiency, seller's revenue, and bidder's surplus. Such studies are absent in the literature, primarily because the capabilities have only now been developed to carry out such auctions in real-time and to provide potentially useful real-time feedback.

In order to test the feasibility of a continuous mechanism, and also examine the effects of various quality and quantity of feedback on the performance of combinatorial auctions, we conduct laboratory experiments where bidders participate in one of three possible treatments that differ only in the type of feedback provided. We consider a hypothetical auction environment wherein individuals bid on real-estate properties surrounding a lake. Our experimental

environment is scalable, and it provides a plausible scenario in which certain sets of items, e.g., adjacent properties, might have greater value as a set than the sum of their values individually. This feature provides the opportunity for combinatorial bidding to offer advantages over traditional, non-combinatorial, single-item bidding. To construct the experimental setup, we relied on theoretical and empirical advances in experimental economics.

Laboratory data forms an important means of analyzing and comparing complex auction mechanisms before they are implemented in the field. As Kwasnica et al. (2005) argue: "... test bed environments in the laboratory [...] exhibit as much complexity or simplicity as one wishes. In these environments, one can test any auction" (p. 421). Several recent studies that have used an experimental methodology to test various combinatorial auction designs, primarily in evaluating the design for FCC spectrum auctions,⁴ provide a starting point for our research. For example, Ledyard et al. (1997), during the evaluation of proposals for the FCC spectrum auctions, explored whether multiple items should be auctioned sequentially or simultaneously; they also examined the question of whether bundle bidding should be allowed. They found that simultaneous auctions were a better choice for heterogeneous items, and that bundle bidding is only preferable when there are significant complementarities among items. Issac and James (2000) successfully operationalized the Vickrey combinatorial auction in the laboratory.

Banks et al. (2003) compared simultaneous multi-round auctions (SMA) with combinatorial multi-round auctions (CMA) and found that CMA outperforms SMA from an

⁴ In the United States, since 1994, the Federal Communications Commission (FCC) has been conducting competitive auctions of licenses for electromagnetic spectra rather than assigning spectra through comparative hearings or through lotteries. FCC spectrum auctions are open to any company or individual that is determined by the Commission to be a qualified bidder.

efficiency perspective when individuals have super-additive preferences for bundles; however, SMA performs better from a revenue perspective, while CMA takes more rounds to complete. Kwasnica et al. (2005) proposed a new design called the Resource Allocation Design (RAD) and compared it to Banks et al.'s (1989) AUSM, and SMA. Kwasnica et al. reported that their design, which provides bidders with approximate price guidance at the end of each round, was more efficient than SMA, and at least as efficient as AUSM. Note that, although AUSM is a continuous mechanism, it provides very limited feedback to bidders for composing and revising bids. Recently, Goeree and Holt (2008) introduced the Hierarchical Package Bidding (HPB) mechanism that performed better than RAD and SMA in the laboratory; however, the mechanism is tailored for a significantly constrained structure of hierarchical bundles proposed by Rothkopf et al. (1998).

The most obvious conclusion from these studies is that allowing combinatorial bidding improves both efficiency and revenue of the auctions. However, any rules that limit bidders' ability to express the complete extent of their willingness to pay for all bundles will likely hurt the efficiency. Further, iterative auctions, where bids can be placed only in well-defined rounds are hard to implement in online markets. In this study, we extend the research on combinatorial auctions by examining the feasibility of conducting continuous auctions using the bidder support schemes that we have developed. In these auctions, bidders can place bids at any time, and on any bundle they want. In the following sections we describe the design of our continuous auction environment, performance measures, the research hypotheses, and the experimental design and procedures. The results and their discussion complete the paper.

3. BIDDER SUPPORT FOR CONTINUOUS AUCTIONS

To develop the bidder support schemes for running continuous combinatorial auctions, we build

upon the real-time bid evaluation metrics developed by Adomavicius and Gupta (2005) that can present bidders with exact price information whenever a bidder wants to explore her alternatives, thereby providing a continuous environment. The technical details of the computational infrastructure are provided in Appendix A. In the following sub-section, we explain the basis of the advanced feedback schemes by describing some characteristics of combinatorial auctions that are different from those of single-item auctions.

3.1 Characteristics of Bids in Combinatorial Auctions

In single-item ascending auctions, if a bidder is not the highest bidder, she must bid an amount higher than the current highest bid to have a chance of winning the auction. However, in combinatorial auctions, even if a bid is not currently winning, it can still be among the future winners depending on the later bids. At any given stage of a combinatorial auction, a bid can be in one of the following three possible states: (i) currently winning (we call this a *winning* state), (ii) currently non-winning but with a possibility of winning in the future (we call this a *live* state), and (iii) currently non-winning with no chance of winning in the future (we call this a *dead* state). This is in contrast to traditional single-item auctions where a bid can only be in either of two possible states (winning/losing). Furthermore, due to the existence of three bid states, combinatorial auctions have two price levels at which a bidder can place non-losing bids: (i) Winning Level (WL), above which a bid would be in a winning state, and (ii) Deadness Level (DL), above which a bid would be in a live state.

As an example of these states and levels, consider an auction with three items: A, B, and C. Table 1 shows the status of all the placed bids at an auction state $k = 5$. Bids 1 and 4 form the *winning* bid combination at the current state of the auction, with total revenue of \$25, which is greater than that of any other bid combination.

Table 1. Example of bid status after five bids have been placed

Bid sequence	[span; value]	Status after 5 bids
1.	[A; \$10]	<i>Winning</i>
2.	[AC; \$20]	<i>Dead</i>
3.	[C; \$11]	<i>Live</i>
4.	[BC; \$15]	<i>Winning</i>
5.	[AB; \$13]	<i>Live</i>

Bid 2 is *dead* (after five bids are submitted) because the combined revenue from the bids on {A} (Bid 1) and {C} (Bid 3) exceeds that of the bundle {A, C}, so Bid 2 cannot win at any subsequent state of the auction. Bids 3 and 5 still have a chance to win depending on subsequent bids (e.g., a new bid of \$13 on {C} would make Bid 5 *winning*) and hence are still *live*. Now, suppose a bidder wants to bid on item {B}. Since no bid on {B} has been placed yet, the DL on this single-item bundle is \$0, i.e., any bid above \$0 would be live. The WL, on the other hand, is \$4 because the bidder has to bid more than that in order to make it a winning bid in combination with the existing bids on {A} and {C}. Thus, if the bid increment is \$1, the bidder will have to bid at least \$1 on item {B} to have a chance of winning the item in the future and at least \$5 in order to be winning at the current state of the auction.

Additional theoretical results about bid states and price levels in combinatorial auctions can be found in Adomavicius and Gupta (2005).

The complexity of combinatorial auctions emphasizes the need to offer real-time support capabilities to help bidders evaluate the existing bids as well as to provide price guidance in formulating new bids. The metrics we discussed in this section provide the opportunity for conducting continuous combinatorial auctions similarly to ascending English auctions and their counterparts by forming the basis for real-time feedback to bidders.

3.2 Feedback Schemes

From the discussion above it is clear that, unlike in traditional single-item auctions, in combinatorial auctions simply displaying all the bids is not sufficient for bidders to formulate smart bids. The computational infrastructure derived from the concepts presented above allows us to provide advanced information feedback to bidders that could potentially be more useful in planning and executing their bidding strategies than just displaying all the bids. The possible feedback items are as follows:

- a) Identifying the *winning* bids at all stages of the auction;
- b) Removing the *dead* bids at all stages of the auction;
- c) Specifying the *DL* for any chosen bundle at all stages of the auction;
- d) Specifying the *WL* for any chosen bundle at all stages of the auction.

The first two forms of feedback address the computational complexity of calculating winning and losing bids in combinatorial auctions and are expected to provide the bidders better awareness regarding the state of all the submitted bids. However, as explained earlier, in combinatorial auctions, even if the bidders realize that their bids are not currently included in the provisional allocation, it is a cognitively challenging task to revise their bids so that they become winning. Specification of DL and WL provides the knowledge regarding the *exact* prices (unlike the approximate prices provided in most of the iterative auctions summarized earlier) for any bundle: DL for knowing the minimum potential non-losing bid, and WL for knowing the minimum potential winning bid, at any stage of the auction.

Furthermore, the ability to place live bids along with the provision of the DL potentially allows the bidders to cooperatively overcome the *threshold problem* (Banks et al. 1989). In combinatorial auctions, the threshold problem refers to the difficulty that bidders seeking smaller

bundles face in outbidding a single bidder who places a high bid on a larger bundle of which the smaller bundles are subsets. For example, consider three bidders competing for two assets: A and B. Suppose that {A} is worth \$50 to Bidder 1; {B} is worth \$50 to Bidder 2; and the bundle {AB} is worth \$75 to Bidder 3. Now, if Bidder 3 places a bid of \$60 on {AB}, neither of the other two bidders acting unilaterally can afford the \$60 necessary to match the bid by Bidder 3. However, in our setting, due to the availability of DL, Bidders 1 and 2 can each place smaller *live* bids, which they can gradually increase to unseat the bigger bid. For instance, in the above example, Bidders 1 and 2 could each incrementally place live bids of \$5, \$10, \$15, \$20, \$25, \$30, \$35 on {A} and {B}, respectively. As a result, when Bidder 1 reaches \$35 on {A} and Bidder 2 reaches \$35 on {B}, their combined bids will become winning (without any of the bidders having to singlehandedly outbid Bidder 3). Banks et al. (1989) allowed two live bids per bidder to be placed on a bulletin board to minimize the threshold problem. However, by allowing continuous bidding and making DL explicitly available, our auction environment gives bidders the ability to place minimum live bids for bundles whose winning levels at any stage of the auction are higher than the bidders' valuations, thus mitigating the threshold problem.

If the bidders are able to utilize all the feedback to construct appropriate bundles and place bids with appropriate amounts, we would expect more efficient bidding. Since higher efficiency means that the total value of the trading system is higher, it has the potential to result in higher revenues for the seller, better surpluses for the bidders, or both. A caveat arises from the bidders' abilities to incorporate the feedback in their decision making: *Can bidders exploit the information provided to them in formulating their bids in a way that maximizes their utility?* To be able to study any differences on the outcome of auctions with differing quality and quantity of feedback, we employ three distinct treatments, i.e., three different levels of feedback,

each cumulatively providing more comprehensive information to the bidders:

1. Level 1: Baseline feedback (control) – this represents the continuous combinatorial auction setup where all submitted bids are visible to all bidders, but no other feedback is provided. The baseline provides a basis for identifying any benefits provided by greater feedback. Note that in most iterative versions of combinatorial auctions this is the only information that is available to the bidders.
2. Level 2: Outcome feedback – this level includes all feedback provided in Level 1 plus identification of the winning bids at all stages of the auction, i.e., the bidders will always be aware of which bids would win if the auction ended right then. The identification of the winning bids represents a non-trivial feedback, since winner determination in combinatorial auctions is a computationally complex problem, as mentioned earlier. However, as mentioned earlier, since winner determination problem can be solved quite fast, it is conceivable that this information can be provided to bidders in iterative combinatorial auctions after each round of bidding. In our setup, this information is updated after every submitted bid.
3. Level 3: Price feedback – this level includes all feedback provided in Level 2 plus price feedback. As mentioned above, this includes the ability to specify the DL (i.e., maximum price at which bids can never be part of a winning bid combination) and the WL (i.e., minimum price above which bids become part of the set of currently winning bids) on demand for every possible bundle. Furthermore, the dead bids are removed in this feedback level. This feedback is expected to direct the bidder toward formulating smarter, more informed bids.

To our knowledge, at present, no other mechanism exists that can provide such information. Note that, the designs such as RAD (Kwasnica et al., 2005) provide a linear combination of individual item prices after each round of bidding – however, whether this price is an upper bound or a lower bound on true WL or DL is an open question. Again, our computational mechanism provides the exact WL and DL in real time.

Note that the auctions with any of the above feedback schemes are theoretically equivalent because, even with baseline feedback, the provisional winners and even the prices for all bundles can be computed by a fully rational bidder. However, the requirements of bidders' cognitive abilities are very different in the three treatments. Therefore, we expect to see differences in bidding behavior in the three treatments leading to differences in auction outcomes that auction theory does not capture. Such differences in outcomes among theoretically isomorphic auction mechanisms due to behavioral elements have also been observed in single-item auctions—both in the single-unit case (Cox et al. 1982; Kagel et al. 1987) as well as the multi-unit case (Kagel and Levin 2009). Thus, Kagel and Levin (2009) assert that, when bidders cannot be expected to be fully rational, “they may benefit from the additional information and transparency embedded in a dynamic format” (p. 222) motivating the above feedback schemes.

4. PERFORMANCE MEASURES OF AUCTIONS

When choosing an auction design, a variety of criteria and measures may be used: allocative efficiency, seller's revenue, and bidders' profit (or surplus) (Kwasnica et al. 2005; Ledyard et al. 1997). In general, as noted earlier, there are tradeoffs among these measures. We considered the following criteria and measures in comparing the economic outcomes of the auctions under various feedback regimes:

1. *Allocative efficiency.* The allocative efficiency of a mechanism measures the social welfare from the allocation using the mechanism as compared to the maximum social benefits that could have been achieved. An auction is said to be 100% efficient when it assigns the set of offered items so that the total value that society obtains from the items is maximized. The optimal allocation occurs when: (i) the winning combination represents the itemsets that have the largest possible aggregate valuation among the bidders, i.e., no other combination of itemsets can produce higher combined valuation; and (ii) the itemsets are sold to individuals that have the highest valuations for those itemsets (regardless of the actual auction price).

The allocative efficiency (η) can be calculated as follows:

$$\eta = \frac{(\text{Total Auction Revenue} + \text{Total Consumer Surplus})}{\text{Max Possible Total Valuation}} = \frac{\text{Valuation of Auction Winners}}{\text{Max Possible Total Valuation}}.$$

Allocative efficiency has been widely used as a performance measure in the combinatorial auctions literature.

2. *Seller's Revenue.* The seller's revenue is the sum of the prices that the bidders pay for the items they win. The amount of revenue generated from a particular auction mechanism partially depends on the distribution of the asset valuations across bidders. This distribution changes when the number of bidders participating in the auction changes. However, in our experiments the number of bidders and the distribution of the valuations are held constant across treatments. Hence the revenue generated from all the treatments can be easily compared. While a more competitive environment may yield higher revenues, a less information-rich environment also may yield higher revenues depending upon bidders' risk preferences.
3. *Bidder's Profit.* Each bidder's profit (or surplus) reflects the increase in net worth of the

bidder after the trade. Together, the balance of seller's revenue and bidders' surplus impact the perceived fairness of the auction mechanism. Bidders may be unwilling to participate in auctions where the seller extracts relatively high revenue from the trade, i.e., bidders' net worth after the trade does not increase significantly. Bidders' surplus is also important from the perspective of comparisons across feedback levels. For example, higher efficiency implies that the total value of the trade is higher (as compared to a lower efficiency trade); however, if a seller is able to extract all the increases in value, bidders may prefer a mechanism with lower efficiency that may provide them similar or higher surpluses. We measure and report individual bidder surpluses in each of the auctions we conducted.

5. RESEARCH HYPOTHESES

As summarized earlier, recent studies on iterative auctions have argued that, if individuals or entities with high valuations can be provided with information on competitors' bids, they are more likely to post competitive bids themselves (Kwasnica et al. 2005; Porter et al. 2003). This feedback increases the likelihood of more efficient outcomes, as the bidders with relatively small valuations are outbid. Therefore, most of the studies in the literature have focused on some form of *outcome* feedback within or across mechanisms. For example, while Porter et al. (2003) found that their Combinatorial Clock (CC) auction was more efficient as compared to SMA (due to complementarity of items), they also observed that auction systems that provide feedback and allow bidders to update their bids produce more efficient outcomes than their sealed-bid counterparts. Similarly, Kwasnica et al. (2005) found that their RAD design, which provides winning bid information after every round, produces higher efficiency than SMA. Since we provide winning bid information in both outcome and price feedback cases, we hypothesize that these two cases will have a significantly higher efficiency than the baseline case, where we do

not provide winning bid information.

Hypothesis 1a: The efficiency in the case of *outcome* feedback will be greater than the efficiency in the case of *baseline* feedback.

Hypothesis 1b: The efficiency in the case of *price* feedback will be greater than the efficiency in the case of *baseline* feedback.

As for differentiating between outcome and price feedback, less is known. The literature on single-item auctions is not relevant since price and outcome feedback are not distinct for these auctions. Would price feedback help direct the goods to the bidders who most highly value them? Price feedback (which, in addition to pointing out winning bids, provides *winning* and *deadness* levels for any item or bundle of interest) provides information that is implicitly oriented toward avoiding too high a bid at a given point in the auction. This can lead to different bidding behavior; but, it is not clear whether there should be an impact on efficiency as compared to outcome feedback. Since the winning bids are displayed, bidders can iteratively bid and discover any bundles that are available below their valuations; therefore, theoretically, any additional benefit with respect to efficiency is questionable with price feedback. On the other hand, practically, bidders are less likely to evaluate all possible bids through the bidding process with outcome feedback; hence, we would expect that efficiency with price feedback will at least be as high as with outcome feedback. Therefore, we have the following hypothesis:

Hypothesis 1c: The efficiency in the case of *price* feedback will not be less than the efficiency in the case of *outcome* feedback.

Next, we consider the impact of feedback on seller's revenue and bidders' surplus. First, we use behavioral theories concerning the effect of feedback on performance to develop

hypotheses on the effect of feedback on bidders' surplus. Then, we derive hypotheses on seller's revenue based on our hypotheses of auction efficiency and bidder's surplus.

Researchers in different fields have studied the role of feedback as well as the effects of different types of feedback on performance and learning (Annett 1969; Balzer et al. 1989; Kluger and DeNisi 1996; Sengupta and Abdel-Hamid 1993). Although most of the early research on feedback suggested that availability of feedback improved learning and performance, more recent research has shown that the effects depend on, among other things, the characteristics of the feedback itself (Balzer et al. 1989; Hogarth et al. 1991; Kluger and DeNisi 1996). In particular, simply providing outcome feedback generally has been insufficient for decision makers to make optimal decisions in complex environments. Although individuals have benefitted from outcome feedback in simple tasks (Doherty and Balzer 1988), in complex tasks, outcome feedback often does not have a positive influence on performance (Brehmer 1980; Hammond et al. 1975; Hoffman et al. 1981). When tasks are complex, what is needed is cognitive feedback on the processes leading to the outcomes (Kluger and DeNisi 1996). The most effective component of cognitive feedback capable of affecting individual performance is task information (Balzer et al. 1989), which enables a decision maker to learn more about the environment. According to Sengupta and Abdel-Hamid (1993), cognitive information enables an individual to obtain greater insight into her strategy. In the absence of task information, a decision maker may have very little idea regarding what constitutes a better strategy (Brehmer 1990), which may lead to perpetuation of flawed strategies. As described earlier, bidding in combinatorial auctions constitutes a complex task in a dynamic environment. Since the information provided to bidders in the outcome feedback treatment lacks strategic information regarding the itemsets and amounts to bid, the feedback is not expected to improve bidder

performance beyond what is achieved using baseline feedback. The price feedback condition, on the other hand, provides a higher level of cognitive feedback, offering task-related (i.e., bid-centric) information that is effective in improving individual's economic performance.

Therefore, we have the following hypotheses:

Hypothesis 2a: The bidders' profit in the case of *outcome* feedback will not be greater than the bidders' profit in the case of *baseline* feedback.

Hypothesis 2b: The bidders' profit in the case of *price* feedback will be greater than the bidders' profit in the case of *baseline* feedback.

Hypothesis 2c: The bidders' profit in the case of *price* feedback will be greater than the bidders' profit in the case of *outcome* feedback.

Finally, since the experimental treatment conditions do not affect seller's behavior, the hypotheses concerning seller's revenue are based on Hypotheses 1 and 2, regarding efficiency and bidders' surplus. Higher efficiency has the potential to increase either the revenue of the seller or the surplus of bidders, or both. Comparing the effect of outcome feedback and baseline feedback, we hypothesized that the efficiency with outcome feedback will be greater (H1a) but the bidders' profit will not be greater (H2a). Since we do not expect the efficiency gains obtained with outcome feedback to increase bidders' profit (relative to baseline feedback), we can expect them to accrue to the seller. A similar argument can be used to compare the impacts of outcome feedback and price feedback on seller's revenue. Since we do not expect the efficiency with price feedback to be any better than what is achieved with outcome feedback (H1c), but the bidders' profits to be higher with price feedback (H2c), we propose that the revenue generated with price feedback will be lower than that with outcome feedback.

Lastly, price feedback is expected to increase both the efficiency and bidders' profit

compared to baseline feedback (H1b and H2b). The increase in efficiency was argued based on the increased information provided by the outcome information of winning bids. The increase in bidders' surplus was argued based on the added information about the task; specifically, price feedback provides information that allows for more strategic formulation of bids. As stated above, increased efficiency can simultaneously increase seller's revenue and bidders' surplus or either one of them. Intuitively, we expect the task information with price feedback to be more effective than the winning bid information with outcome feedback, allowing the bidders to place more precise bids and extract higher surplus. Thus, if the bidders are able to effectively use price feedback, we expect most of the gains from increased efficiency to go to bidders and, consequently, the seller's revenue with price feedback will not be greater than that with baseline feedback. Hence, our final set of hypotheses is as follows:

Hypothesis 3a: The seller's revenue in the case of *outcome* feedback will be greater than the revenue in the case of *baseline* feedback.

Hypothesis 3b: The seller's revenue in the case of *price* feedback will not be greater than the revenue in the case of *baseline* feedback.

Hypothesis 3c: The seller's revenue in the case of *outcome* feedback will be greater than the revenue in the case of *price* feedback.

6. METHODOLOGY

To construct the experimental setup, we rely on Smith's (1976) induced-value theory that identifies sufficient conditions for experimental control. The key idea is that the proper use of a reward mechanism allows an experimenter to induce pre-specified characteristics in experimental subjects. Proper use is further defined to consist of a monotonic non-satiable utility

for the reward and that the incremental reward a person receives depends on her actions (and those of other agents) as defined by the institutional rules that she understands. The use of real currency is known to satisfy these important conditions. Furthermore, Jamal and Sunder (1991) find that such salient rewards tend to increase the reliability of results. Smith and Walker (1993) provide a summary of evidence that further supports the use of real monetary rewards in experimental economics. We borrow significant design aspects from researchers who have conducted experiments in the domain of combinatorial auctions summarized earlier (e.g., Banks et al. 2003; Kwasnica et al. 2005; Ledyard et al. 1997). To test the robustness of the results, we conducted two sets of experimental auctions with different setups. Setup 1 used a symmetric, systematic valuation scheme across bidders. Setup 2 used an asymmetric, random valuation scheme. In the following subsections, we describe our experimental environment.

6.1 Experimental Design

In our auctions, m bidders compete to acquire N property lots surrounding a lake. For the specific experiments described in this paper we used $m = 3$ and $N = 6$. Let \mathbf{J} be the set of bidders and \mathbf{I} the set of distinct items. The property lots are adjoining and successively labeled A through F as shown in Figure 1a. The first and the last lots are also adjoining. Note that, while the *spatial environment* considered by Kwasnica et al. (2005), and Ledyard et al. (1997) also had six items, the set of items each bidder could bid was restricted. However, we allow bids on all $2^6 - 1$ (i.e., 63) possible combinations (i.e., all possible item bundles).

Setup 1 (Symmetric, systematic valuations). In this setup, each bidder has a *preferred lot* – the lot that she values most. The identity of the preferred lot is private information to each bidder. The value of this lot is represented by v_p . The preferred lots make it a hard environment for the outcome or price feedback to have comparative advantages over the baseline case, since

good efficiency can be obtained even without feedback if bidders focus their attention around their preferred lot. This provides a strong test of the feedback hypotheses. The value of all other lots can be set to decrease by a certain percentage r_d as the lot is farther from the preferred position. We call this the *rate of decay*. Let the *distance* of a lot imply how far it is from the preferred lot in terms of the number of lots in between. Thus, all adjoining lots of the preferred lot for a bidder are at distance 1. If d_{ij} denotes the shortest distance of the i^{th} lot for the j^{th} bidder (i.e., minimum of the clockwise and counterclockwise distances), then the value v_{ij} of the lot i to the bidder j is given by:

$$v_{ij} = v_p (1 - r_d)^{d_{ij}}, \text{ for } i \in \mathbf{I}, j \in \mathbf{J} \quad (1)$$

For example, if Lot A is the preferred lot with $v_p = \$100$, and if the decay rate r_d is 50%, then the values of the individual lots will be as shown in Figure 1b. This setup creates a distribution of valuations for each bidder that is very easy to scale with more bidders and more assets.

Furthermore, any significant differences in efficiency we see with the provision of feedback in the design can be attributed to feedback.




Figure 1a. Property lots surrounding a lake




Figure 1b. Valuation of lots with preferred lot A, $v_p = \$100$, and decay rate $r_d = 50\%$

Setup 2 (Asymmetric, random valuations). In this setup, values are generated and assigned in two phases. In Phase 1, six unique values are randomly drawn from a uniform distribution $U(\alpha, \beta)$. In Phase 2, for each bidder, the six values picked in Phase 1 are randomly

assigned as the bidder's valuations for the six lots. We refer to each assignment set as an *instance*. Each instance was used in all three treatments to allow comparison across treatments.

Phases 1 and 2 were repeated to generate instances to run multiple auctions.

Complementarity. A key characteristic that contributes to the efficiency of combinatorial auctions (compared to selling multiple items in separate single-item auctions) is the non-linear valuation of items. That is, the valuation of a bundle is not simply additive but super-additive due to the existence of complementarities among the items in the bundle. In both setups, we introduce this possibility in our design through the use of variable r_s , called the *super-additivity rate*. Owning adjoining lots increases the combined value of the lots by r_s for every additional adjoining lot, thereby creating super-additive valuations that depend on the number of adjoining lots in the bundle. In our case, if a is the number of consecutively adjoining lots in the bundle, then the super-additive adjustment for the valuation of this bundle is $(a - 1)r_s$. For example, if $r_s = 10\%$ then the valuations of some bundles, for the case depicted in Figure 1b, are shown in Table 2.

Table 2. Examples of super-additive valuation calculation

	Bundles	Number of adjoining lots	Valuations of bundles
Example 1	{AF}	2 adjoining	$(\$100 + \$50) * 1.1 = \$165.00$
Example 2	{CDE}	3 adjoining	$(\$25 + \$12.5 + \$25) * 1.2 = \75.00
Example 3	{BEF}	1 separate; 2 adjoining	$\$50 + (\$25 + \$50)*1.1 = \132.50

The aforementioned example can be generalized with a set of parameters to the experimental environment of Banks et al. (2003), where the valuation of an itemset can be defined as:

$$V_{\Omega_j} = \sum_{i \in \Omega} v_{ij} + A_j \sum_{i \in \Omega} \sum_{\substack{k \in \Omega \\ k > i}} l_{ik} \quad (2)$$

where $V_{\Omega j}$ is the value of the bundle Ω for bidder j . Here bundle Ω is represented by an ordered set of items; v_{ij} is the valuation of item $i \in \Omega$ by bidder j ; A_j is the super-additivity factor for bidder j ; and l_{ik} the minimum of clockwise and counter-clockwise distances between two items. For example, if $l_{ik} = 1$ (say, for adjacent items i and k), the total value of a bundle just consisting of these two items would be $(v_{ij} + v_{kj} + A_j)$. The model in Equation (2) allows both for a compact description of the scenario that can be provided to bidders participating in the auction as well as for building a simulation of the auction environment.

6.2 Valuations

In Setup 1, the preferred lots for the three bidders were A, C, and E respectively. The values of the other variables were set as: $v_p = \$100$, $r_d = 50\%$, and $r_s = 10\%$. The distribution of the valuation of the lots is shown in Figure 2. Since the bidder valuations are identically distributed, any consistent differences in bidders' behaviors among auctions with differing levels of feedback can be reasonably attributed to varying bidder responses to feedback. We conducted several simulation runs with computerized bidding agents as well as several pilot tests with human bidders to refine values of the parameters in our model before carrying out our main experiments. For instance, a variation of r_s (the super-additivity rate) can have a significant impact on the outcome of the auction. For example, in our setup, if $r_s = 25\%$, the optimal allocation is for one bidder to buy all the lots. In that case, if a bidder places a high bid on the bundle that includes all the lots, no other combination of bids can outbid her. After several simulations we chose to set $r_s = 10\%$; this leaves ample room for combinations of bundle bids to outbid a single bid on all the lots, providing a competitive bidding environment.

In Setup 2, the six unique values were randomly drawn from the uniform distribution with support $[\$5, \$100]$ as part of Phase 1. In Phase 2, for each bidder, the six values picked in

Phase 1 were randomly assigned to the six lots. As an illustrative example, suppose the set of six values randomly drawn in Phase 1 is $\{\$84, \$17, \$58, \$25, \$21, \$97\}$, then a possible assignment from Phase 2 can be as shown in Figure 3. Finally, the same super-additivity rate as in Setup 1 (i.e., $r_s = 10\%$) was used in Setup 2 as well.




Figure 2. Valuation of lots for the bidders in Setup 1




Figure 3. A sample assignment of values for Setup 2

6.3 Auction Procedure and Rules

In each session, prior to the beginning of the auctions, instructions explaining the rules of the auction were read out loud so that everyone could hear. The instructions were followed by short tests to ensure that the participants understood the rules of the auction and were familiar with the bidding environment. The participants in each session were randomly assigned to a particular auction. They were not told how many other participants were in their specific auction, because

in online auctions bidders are usually unaware of the number of people interested in the commodity. Furthermore, while the rules to generate the valuations of the lots were common knowledge and each bidder in an auction knew the distribution of their own values, participants had no explicit knowledge of the valuations of other bidders.

Bidders were allowed to place any number of bids on individual lots as well as on any combination of lots. They could place as many bids as they wanted and at any time they wanted during the course of the auction. In keeping with practical online auctions, bid withdrawal was not allowed. At the end of the auction, winners were determined, and the participants were notified of the result along with their individual profits based on the final allocation. In addition to these general rules, we adopted the following specific rules:

- (i) *Bid increments*: The bid increment was set at \$1. The auction interface ensured that only integer (dollar amount) bids could be submitted;
- (ii) *Stopping rule*: A *soft* stopping rule was used, i.e., after an initial time period, the auction ended if no new bids were placed for x minutes. This rule of extending the auction was followed in order to eliminate *sniping*, i.e., placing bids in the last few seconds of the auction. The initial time period was chosen as 13 minutes, with $x = 2$ minutes in Setup 1 and $x = 1$ minute (again, varied to test robustness of results) in Setup 2. So, each auction lasted at least 15 minutes in Setup 1 and 14 minutes in Setup 2. The average duration of the auctions was 27 minutes for Setup 1 and 22 minutes in Setup 2.

Bidders were instructed as to how the valuations were generated. The auction interface allowed them to see their individual valuations for any bundle. The bidders were not given any fixed budget; but, in the final compensation scheme, bidders were paid for the lots they won in

proportion to their profits from the auction. Profit was calculated as the difference between their valuation of the item(s) and their winning bid(s). Consequently, their profits were positive, zero, or negative depending on whether their winning bid was less than, equal to, or greater than their valuation, respectively. If they did not win any lot, their profit was zero. Each subject was paid an up-front sum of \$10 for participation. Friedman and Sunder (1994) recommend this practice for three reasons: (a) to reduce tardiness, (b) to establish ex ante credibility with the subjects that the rewards being promised to them will be paid to them promptly, and (c) to provide an initial cushion of wealth they can afford to lose in the actual experiment without dipping into their own wallets. At the end of the auction, they were paid 20 cents for every experimental dollar of their profit. Similarly, they were charged 20 cents from their participation fee for every dollar of loss. The maximum amount that could be taken off was their participation fee. Participants were paid privately in sealed envelopes.

6.4 Auction Interface

The screen snapshots of the auction interfaces are presented in Appendix B, and the descriptions of various fields are provided in Appendix C. The auction interface for both setups and all three treatments (i.e., feedback levels) was the same except for the type of feedback provided. With baseline feedback, all bids were displayed, with the bidder's own bids identified on her own screen. The bidders' assigned valuations for the individual lots were displayed on their screens at all times. The bidders could find their private valuations for any possible combination of the lots (bundle) by clicking on checkboxes provided against each lot. Bids could be placed by selecting the lots, entering a bid amount, and then pressing the <Submit Bid> button. The total elapsed time of the auction and the time since the last bid was placed were also displayed.

With outcome feedback, in addition to displaying all the bids, the winning bids at any given stage of the auction were identified. The set of winning bids was updated after every new winning bid.

Price feedback consisted of a specification of the following prices for any given bundle at any point during the auction: (1) DL^* : $DL + \epsilon$ (where ϵ is the bid increment = \$1 in our experiments) – the minimum price at which their bid stands a chance of winning in the future given all the other bids at that state of the auction, and (2) WL^* : $WL + \epsilon$ – the minimum price at which their bid would be winning at that state of the auction. These two prices provided a range of values for utility-maximizing bidders to bid on a selected bundle. A bidder could find these two bounds for any possible bundle by simply clicking on the lots constituting the bundle of interest.

6.5 Participants

With Setup 1, we conducted a total of 51 auctions over 15 experimental sessions. Three to four auctions were simultaneously conducted in each session. Bidders did not know who were the competing bidders or how many bidders they were competing against. For example, in a given session, there might be 4 auctions conducted with 12 bidders, however, the bidders did not know how many auctions were running simultaneously. Further, bidders were not told that the valuations were symmetrically distributed in Setup 1 and had no knowledge of other bidders' preferred lots. No bidder was allowed to participate in more than one auction. The 153 unique participants in the 51 auctions were all undergraduate business students who responded to volunteer solicitation announcements throughout the university campus. The average age of the subject pool was 20 years; 54% were male. We excluded seven auctions from our analysis, because in these at least one bidder mistakenly placed a bid significantly above her valuation.

They immediately notified us of the mistake; however, since our design disallowed bid withdrawal, rectification of the user error was not possible. Therefore, we removed these auctions from further analysis, attributing the irrational bids to bidding errors.

With Setup 2, we conducted a total of 43 auctions over 12 experimental sessions. Again, three to four auctions were simultaneously conducted in each session with no bidder being allowed to participate in more than one auction. The 129 unique participants in these auctions were also undergraduate business students but at a different university. The composition of the subjects was fairly similar with the average age of the subject pool being 22 years, and 51% male. We excluded four auctions from our analysis, because in these at least one bidder mistakenly placed a bid significantly above her valuation.

7. RESULTS

In Sections 7.1 and 7.2, we present the descriptive statistics and the hypotheses test results for Setups 1 and 2, respectively.

7.1 Setup 1: Symmetric and Systematic Valuations

Descriptive statistics for efficiency, seller's revenue, and bidder's profits for the continuous combinatorial auctions are provided in Table 3. To test the hypotheses, we conducted Mann-Whitney Rank Sum tests, the non-parametric counterpart of the unpaired t -test. Under the conditions of the auction setting, it is not necessarily reasonable to assume that the data meet the normality assumptions required for a t -test (see, for example, Kwasnica et al. 2005; Porter 1999). The rank sum test only requires the two samples to be independent (Siegel and Castellan 1988).

The results of the hypotheses tests, displayed in Table 4, show that either form of feedback – outcome or price – results in a significant improvement in auction efficiency over the baseline case. Thus, the hypotheses regarding the positive impact of outcome and price feedback

on efficiency (H1a and H1b) are supported. Efficiency under price feedback was not less than that with outcome feedback, though the greater efficiency was not statistically significant: The two advanced levels of feedback were comparable and the null hypothesis of equality was not rejected (H1c supported).

Table 3. Descriptive statistics of the auctions with Setup 1

Treatments	Number of Auctions Conducted (number of participants)	Mean (SE) Efficiency	Mean (SE) Bidder's Profit	Mean (SE) Seller's Revenue as a percentage of maximum possible revenue
Baseline Feedback	14 (42)	86.23% (3.0%)	\$33.60 (\$5.02)	66.14% (4.3%)
Outcome Feedback	15 (45)	90.64% (3.5%)	\$21.57 (\$3.06)	80.66% (1.8%)
Price Feedback	15 (45)	93.48% (1.6%)	\$47.54 (\$6.49)	65.27% (5.3%)

Table 4. Results of hypotheses tests with Setup 1

Comparison	Efficiency	Bidder's Profit	Seller's Revenue
Outcome feedback vs. Baseline feedback	Outcome > Baseline (H1a)* $z = 1.433; p = 0.075$	Outcome \leq Baseline (H2a)* \dagger $z = 1.384; p = 0.083$	Outcome > Baseline (H3a)*** $z = 2.620; p = 0.004$
Price feedback vs. Baseline feedback	Price > Baseline (H1b)** $z = 1.934; p = 0.026$	Price > Baseline (H2b)* $z = 1.357; p = 0.087$	Price \leq Baseline (H3b) $z = 0.131; p = 0.448$
Price feedback vs. Outcome feedback	Price \geq Outcome (H1c) \dagger $z = 0.064; p = 0.474$	Price > Outcome (H2c)*** $z = 2.581; p = 0.004$	Price < Outcome (H3c)** $z = 2.096; p = 0.018$

Note: All tests are one-tailed. Inequalities in the table are in the hypothesized direction.

\dagger Test is of the directional inequality as the alternative hypothesis.

* $p < 0.1$ ** $p < 0.05$ *** $p < 0.01$

Turning to the bidder side, as we hypothesized, price feedback led to higher bidders' profit compared to baseline feedback (H2b supported). Auctions with just outcome feedback significantly decreased bidders' profit compared to the auctions with baseline feedback.

Although we did not anticipate this decrease in bidders' profits with outcome feedback, the result

is consistent with the expectation that the outcome feedback would not increase bidders' profits (H2a supported). The sister hypotheses regarding the seller's revenue are also supported. Since outcome feedback made the auction environment more efficient compared to the baseline case but did not lead to higher bidder surplus, the benefits of the increased efficiency went to the sellers, resulting in higher seller's revenue, thus supporting H3a. In contrast and as hypothesized, the increased efficiency of price feedback resulted in higher bidders' profits. The seller did not realize the benefits of the increased efficiency; in this case, the lower seller's revenue compared to baseline feedback was not statistically significant (H3b supported).

The contrast between price and outcome feedback is highlighted by their direct comparison. As hypothesized, price feedback allowed bidders to make better winning bids and thus to retain higher surpluses as compared to outcome feedback (supporting H2c); correspondingly, price feedback led to comparatively lower seller's revenues (H3c supported). Returning to the data in Table 3 and comparing outcome feedback to the baseline case, we observe that, while the efficiency increases, all the gains (and more) go to the seller, as indicated by substantial increase in seller's revenue and a decrease in bidders' profits. However, when comparing price feedback to the baseline case, we observe that the efficiency gains result in relatively more equitable distribution of benefits, as observed by no decrease in revenue but a significant increase in bidders' profits. Overall, in terms of efficiency, both outcome as well as price feedback clearly improve the performance of the auctions. However, our results reveal that the distribution of the gains (resulting from higher efficiency) between the bidders and the seller is influenced by the nature of the feedback.

7.2 Setup 2: Asymmetric and Random Valuations

To test the robustness of the results presented above, we conducted a second set of experiments with a different setup and a separate subject pool from a different university as described in Section 6. The descriptive statistics of the auctions and the hypotheses test results for Setup 2 are shown in Tables 5 and 6, respectively.

Table 5. Descriptive statistics of the auctions with Setup 2

Treatments	Number of Auctions Conducted (number of participants)	Mean (SE) Efficiency	Mean (SE) Bidder's Profit	Mean (SE) Seller's Revenue as a percentage of maximum possible revenue
Baseline Feedback	13 (39)	78.78% (2.4%)	\$22.27 (\$6.93)	64.76% (4.7%)
Outcome Feedback	13 (39)	86.74% (2.4%)	\$20.67 (\$7.31)	74.52% (3.9%)
Price Feedback	13 (39)	89.39% (2.6%)	\$39.89 (\$7.92)	65.99% (4.8%)

Table 6. Results of hypotheses tests with Setup 2

Feedback	Efficiency	Bidder's Profit	Seller's Revenue
Outcome feedback vs. Baseline feedback	Outcome > Baseline (<i>H1a</i>)** $z = 2.098; p = 0.018$	Outcome \leq Baseline (<i>H2a</i>) † $z = 0.175; p = 0.431$	Outcome > Baseline (<i>H3a</i>)** $z = 2.132; p = 0.017$
Price feedback vs. Baseline feedback	Price > Baseline (<i>H1b</i>)*** $z = 2.481; p = 0.007$	Price > Baseline (<i>H2b</i>)* $z = 1.503; p = 0.067$	Price \leq Baseline (<i>H3b</i>) † $z = 0.314; p = 0.377$
Price feedback vs. Outcome feedback	Price \geq Outcome (<i>H1c</i>)† $z = 0.943; p = 0.173$	Price > Outcome (<i>H2c</i>)* $z = 1.572; p = 0.058$	Price < Outcome (<i>H3c</i>) $z = 1.153; p = 0.125$

Note: All tests are one-tailed. Inequalities in the table are in the hypothesized direction.

† Test is of the directional inequality as the alternative hypothesis.

* $p < 0.1$ ** $p < 0.05$ *** $p < 0.01$

Unlike the first setup, the assignment of values for each bidder was different in different auction instances of the experiment in this setup; however, the instances were matched across treatment conditions. Since each instance had a different set of values, we conducted Wilcoxon

matched-pairs test (Conover 1980) to test our hypotheses. This test is the non-parametric equivalent of the paired *t*-test. We compared the outcomes of each instance pair-wise across the three treatments.

Table 7. Summary of hypotheses test results

Hypothesis	Supported in Setup 1?	Supported in Setup 2?
<i>Hypothesis 1a</i> : The efficiency in the case of <i>outcome</i> feedback will be greater than the efficiency in the case of <i>baseline</i> feedback.	Yes	Yes
<i>Hypothesis 1b</i> : The efficiency in the case of <i>price</i> feedback will be greater than the efficiency in the case of <i>baseline</i> feedback.	Yes	Yes
<i>Hypothesis 1c</i> : The efficiency in the case of <i>price</i> feedback will not be less than the efficiency in the case of <i>outcome</i> feedback.	Yes	Yes
<i>Hypothesis 2a</i> : The bidders' profit in the case of <i>outcome</i> feedback will not be greater than the bidders' profit in the case of <i>baseline</i> feedback.	Yes (less than)	Yes
<i>Hypothesis 2b</i> : The bidders' profit in the case of <i>price</i> feedback will be greater than the bidders' profit in the case of <i>baseline</i> feedback.	Yes	Yes
<i>Hypothesis 2c</i> : The bidders' profit in the case of <i>price</i> feedback will be greater than the bidders' profit in the case of <i>outcome</i> feedback.	Yes	Yes
<i>Hypothesis 3a</i> : The seller's revenue in the case of <i>outcome</i> feedback will be greater than the revenue in the case of <i>baseline</i> feedback.	Yes	Yes
<i>Hypothesis 3b</i> : The seller's revenue in the case of <i>price</i> feedback will not be greater than the revenue in the case of <i>baseline</i> feedback.	Yes	Yes
<i>Hypothesis 3c</i> : The seller's revenue in the case of <i>outcome</i> feedback will be greater than the revenue in the case of <i>price</i> feedback.	Yes	No

The test results obtained with Setup 2 are very similar to those obtained with Setup 1 except for Hypotheses 2a and 3c. Hypothesis 2a is still supported; however, in this case, the bidders' profits are not significantly less in the outcome feedback condition compared to the baseline condition. For Setup 2, unlike Setup 1, Hypothesis 3c was not supported. In Setup 2, although the average revenue with outcome feedback was greater than that with price feedback as in Setup 1 (i.e., in the hypothesized direction), the difference was not statistically significant.

Overall, the similarity of the results provides confidence in our analysis of how feedback influences the economic outcomes of continuous combinatorial auctions. The hypotheses test results from the two setups are summarized in Table 7.

8. DISCUSSION: BIDDING BEHAVIOR

In order to better understand the rationale for the observed aggregate outcomes as a function of feedback, we examine several auction- and bid-level characteristics. Overall, the expectations for these characteristics across the different feedback levels arise from the discussion of economic and behavioral theories in Section 5. Outcome feedback, compared to baseline feedback, leads to more competitive bidding with greater efficiency (Kwasnica et al. 2005; Porter et al. 2003). This should be reflected by more frequent bidding. However, outcome feedback alone is not expected to lead to better strategic performance. For complex tasks like the continuous combinatorial auction, task information is needed (Balzer et al. 1989). In our setting, price feedback provides this information. Thus, with price feedback, we should see not only more frequent bidding (compared to the baseline condition), but also better bidding in a strategic sense. Thus, characteristics that connect to more strategic bidding behavior in response to the task-based price feedback should show improvement in these analyses.

A comparison of the average duration of each auction, the number of bids placed in each, and the number of unique bundles that were bid upon (maximum = 63) are presented in Figure 4. For each treatment, the average duration of the auctions is lower for Setup 2. This is expected, because the minimum duration for Setup 2 was 14 minutes whereas it was 15 minutes for Setup 1. However, in spite of the shorter duration, more bids were placed in each treatment in Setup 2, most likely because of its asymmetric design. The auctions with baseline feedback required less time to complete (though the difference in duration is not statistically significant for Setup 2) on

average and were less competitive in terms of the number of bids placed. Without explicit knowledge of whether they were winning, the bidders in the baseline treatment stopped bidding earlier than their counterparts in the other two treatments. Furthermore, bidders used fewer distinct bundles in each auction, on average, with baseline feedback than outcome and price feedback cases (though not statistically significant for Setup 2). Both of these tendencies are as expected, and are indicative of the effectiveness of the feedback in promoting the increased efficiency observed compared to the baseline condition.




Figure 4. Comparison of average auction duration, number of bids, and number of distinct bundles used

A comparison of the percentages of the three types of bids (winning, live, and dead) submitted in each treatment (shown in Figure 5) demonstrates the ineffectiveness of outcome feedback (over and above baseline feedback) in helping bidders improve their performance in complex dynamic tasks. With outcome feedback, as with baseline feedback, a large percentage of bids were dead (losing): over 30% in Setup 1 and over 40% in Setup 2.⁵ However, with the availability of price feedback, the percentage dropped significantly ($\alpha = 0.05$): less than 8%⁶ of

⁵ Setup 1: 30% with baseline feedback and 31% with outcome feedback. Setup 2: 42% with baseline feedback and 43% with outcome feedback.

⁶ 6% in Setup 1 and 7% in Setup 2.

the bids⁷ in each setup. Thus, outcome feedback led to increased bidding frequency that increased the efficiency compared to the baseline case; however, a good percentage of this extra bidding was not economically beneficial to bidders as compared to the price feedback condition.




Figure 5. Comparison of bids based on bid types

The key to the effectiveness of cognitive feedback in dynamic decision making tasks lies in presenting task information that enables decision makers to devise suitable decision strategies and to calibrate them with system changes (Sengupta and Abdel-Hamid 1993). While outcome feedback only enabled the bidders to realize whether they were winning, it lacked the task

⁷ The existence of any such bids with price feedback might seem surprising, but one likely cause of this behavior is that the DL could have increased (by virtue of other bids) between the times a bidder observed the level and placed a bid. This phenomenon of price increase while a bidder is in the process of placing a bid is not limited to our setup or just combinatorial auctions but is possible in any bidding environment. For example, in single-item online auctions, a bidder may see the current highest bid and start placing a marginal bid; but, while she is placing the bid, another bidder can place a higher bid, so that when the first bidder finishes placing the bid, it has been already outbid.

information that is required for high performance in a dynamic environment. Through the provision of exact bid prices, price feedback enabled bidders to avoid inconsequential bids and accrue higher profits for themselves.

A second noteworthy observation is that when both the DL^* and WL^* were available (i.e., with price feedback) bidders more often chose to bid above WL^* . In Setup 1, 67% of the bids placed with price feedback were winning compared to 27% live, and in Setup 2, 59% winning compared to 34% live, with the differences significant at $\alpha = 0.05$. While placing bids at or above WL^* ensured that the bids were immediately provisionally winning, the trade-off is foregone surplus. That is, if the bidders placed a live bid (at or above DL^*) instead of a winning bid (at or above WL^*) on a bundle, then they had to wait for suitable complementary bids to arrive from other bidders (that may or may not come) in order to win. But, had such a complementary bid arrived, it could have generated greater surplus. This phenomenon of placing bids bigger than the minimum non-losing value can be explained using behavioral theory that has been used to explain jump bids (bid increments that are larger than the minimum necessary) in other ascending auctions, including FCC spectrum auctions (Banks et al. 2003; Issac et al. 2007). Through an empirical analysis of FCC auctions, Issac et al. (2007) found that jump bidding exists even in auctions with large bid increments. They assert that jump bidding is not necessarily an indication of irrationality, but rather a sign of impatience or strategic bidding (to preemptively outbid others). Banks et al. (2003) observed that “jump bidding is encouraged by impatient bidders who desire to speed up the pace of the auction” (p. 312). In our context, the bidders were frequently willing to forego surplus rather than wait for a complementary bid to arrive from another bidder. However, as we explain later, the winning bids with price feedback were closer to WL^* as compared to those with outcome feedback, generating higher surplus for

the bidders with price feedback. Between the two types of feedback where WL^* was not available (i.e., baseline and outcome feedback), the percentages of live and winning bids were not significantly different indicating that, as expected, bidders were doing a lot of guesswork even when they were generally placing more aggressive bids (as evidenced by the winning bids being much larger than WL^* as compared to price feedback).

To further understand the basis for the observed differences in aggregate outcome across feedback, we study three important characteristics of each bid: (i) **SIZE**, defined as the number of items in the bundle being bid; (ii) **VALUE**, defined as the percentage of the total value of the bundle that the bidder bid, e.g., if the bundle {ABC} was worth \$210 to a bidder, and she bid \$70 on it, the **VALUE** of the bid is 0.33; and (iii) **LEVEL**, defined as the percentage by which the amount of the bid was above the minimum required to be live, calculated as $(\text{Bid} - DL^*) / (WL^* - DL^*)$. Thus, by definition, the **LEVEL** of all dead bids is negative; the **LEVEL** of all live bids is greater than 0 but less than 1; and the **LEVEL** of all winning bids is greater than or equal to 1. Since the dead bids were inconsequential, and played no role in influencing the economic outcome of the auctions, we compare the metrics for live and winning bids only.

Figure 6 shows a comparison of the **SIZE** of the bundles. Since the number of items in each auction was six, the maximum size of any bundle was six. Firstly, across all treatments and for both setups, over 70% of the bids were on bundles of size greater than 1, pointing to the complementarity effect of the assets. That is, the bidders saw the value of combinatorial bidding and tried to capitalize on the offered capability across all auctions. Secondly, as can be observed, no significant differences in bundle size were noticed across the three treatment levels. That is the complementarity effect was not influenced by the nature of feedback available to the bidders.




Figure 6. Comparison of bids based on bundle sizes

The other two bidding metrics did show significant differences, as can be observed from the averages presented in Table 8. Comparing price and outcome feedback, in both setups we find that VALUE and LEVEL were higher ($\alpha = 0.05$) with outcome feedback. With outcome feedback, bidders placed bids closer to their individual valuations of the bundles and, without explicit knowledge of the DL^* and WL^* , farther up from the required minimum. As a result, having similar efficiency as with price feedback, the auctions with outcome feedback generated higher revenue for the seller but lower surplus for the bidders.

Comparing price and baseline feedback, we find that in Setup 1 both the VALUE and LEVEL were higher (significant at $\alpha = 0.05$) with baseline feedback; however, in Setup 2 only the difference in LEVEL was significant ($\alpha = 0.05$). Relative to price feedback, the high VALUE bids resulted in lower surplus for the bidders with baseline feedback, but the high LEVEL bids did not result in higher revenue due to the lower efficiency in the case of baseline feedback. Conversely, relative to baseline feedback, even though a far greater number of bids

were placed with price feedback (Figure 4) and with most of them winning (Figure 5), LEVEL with price feedback was still lower because with the provision of DL* and WL* bidders were able to place more accurate marginal bids.

Comparing outcome and baseline feedback, we find that, although LEVEL is higher with baseline feedback (though not significantly in Setup 2), VALUE is significantly lower, resulting in lower revenue in the baseline case. Moreover, with fewer bids and shorter duration of the auctions, the efficiency of the auctions with just baseline feedback was lower than the other two cases.

Table 8. Average VALUE and LEVEL for non-losing bids

	VALUE		LEVEL	
	Setup 1	Setup 2	Setup1	Setup 2
Baseline Feedback	0.63	0.52	7.41	4.97
Outcome Feedback	0.72	0.70	5.42	4.47
Price Feedback	0.50	0.53	3.47	2.24

Overall, we observe significant differences in bidding behavior across experimental conditions. The cognitive feedback provided in the form of price feedback allowed bidders to choose bundles strategically with more room for profits. Furthermore, with explicit task-related information on the minimum price for placing provisionally winning bids, they avoided inconsequential bids and did not significantly overbid. These results provide evidence that the high efficiency of the auctions as well as high surplus for the bidders achieved with price feedback is a consequence of the cognitive feedback. With outcome feedback, since the bidders were aware of whether they were winning, they placed a large number of bids. The greater number of bids increased the efficiency of the auctions compared to the baseline case; however, without the aid of exact prices, the bidders tended to overbid, thus, benefitting the sellers at the

expense of their own profit. In the baseline case, without the aid of either outcome or price feedback, the auctions ended early and with fewer bids, resulting in lower efficiency and lower revenue compared to the outcomes with feedback.

9. CONCLUSIONS

IT is increasingly being used to automate existing market processes, but it also presents opportunities to design and deploy new, innovative market mechanisms. While the benefits of combinatorial auctions in the presence of super-additive valuations of assets have been theoretically and empirically demonstrated, the inability to provide meaningful feedback in real time has resulted in limited application of such auctions. Continuous mechanisms, where bidders update their bids in an ongoing fashion and not just in discrete rounds, are primary characteristics of most online auction institutions. A basic requirement for generating bidder participation in such auctions is the availability of information regarding the current state of the auction, e.g., identification of the currently winning bid(s). Online auctioneers have implemented many different variations of classical single-item auctions (including auctions with multiple units – see, for example, Bapna et al. 2003; Kagel and Levin 2009), but there have been no widespread implementations of continuous combinatorial auctions to sell multiple heterogeneous items to multiple bidders. This has led researchers (Adomavicius and Gupta 2005; Kwasnica et al. 2005) to assert that supporting bidders in combinatorial auctions is the next big challenge in facilitating wider use of combinatorial auctions.

In contrast to prior research on combinatorial auctions, our research focused on exploring the feasibility of *continuous* combinatorial auctions without limiting the scope to a specific application (e.g., the FCC spectrum auctions) and without imposing restrictive bidding rules, such as pre-specification of biddable bundles. The contributions of this paper (expanded in

Table 9) are twofold: (1) we implement novel feedback schemes that can facilitate bidder participation in continuous combinatorial auctions; and (2) we analyze the impact of different types of feedback on auction outcomes and bidding behavior using laboratory experiments.

Table 9. Summary of contributions

Contributions	Research contributions	Managerial contributions
<p>1. Implement a continuous combinatorial auction mechanism with novel feedback schemes that is designed to mimic the success of the open ascending auction in the single-item case.</p>	<p>Due to the complexity of a combinatorial bidding environment, most of the existing research on combinatorial auctions focuses on designing iterative auction mechanisms with numerous rules and restrictions. Through a series of controlled experiments, we demonstrate that, with appropriate feedback, efficient continuous combinatorial auctions can be conducted. We answer the previous researchers' call for designing effective bidder support schemes for a combinatorial bidding environment. Thus, we contribute to the active area of research in finding an efficient allocation mechanism for assets that exhibit value synergies – one that can be easily deployed in the online marketplace (being continuous).</p>	<p>While the benefits of conducting combinatorial auctions when assets exhibit complementarities have been recognized in prior literature, the mechanism is not generally available in the online marketplace. The primary reason for the absence is the non-availability of a continuous mechanism that could be easily implemented and where the cognitive costs of participation are not prohibitive. The rules of the continuous mechanism that we have designed are very similar to those of single-item auctions prevalent on the Web. Thus, the mechanism could be easily deployed in the online marketplace. It also has the potential to generate bidder participation, since the cognitive complexity of participation is lowered through the provision of relevant non-trivial feedback.</p>
<p>2. Analysis of the impact of varying levels of feedback on auction outcomes.</p>	<p>Most of the existing research on combinatorial auctions focuses on comparing different auction mechanisms based on economic performance measures without regard to the fact that, even within a given mechanism, the provision of different types of information can significantly influence the economic outcomes. Using two novel feedback schemes along with a baseline condition, we demonstrate that information feedback plays a key role in affecting bidding behavior resulting in differences in the economic performance of combinatorial auction mechanisms. We also highlight the importance of fairness considerations in assessing efficiency gains. Depending on the type of feedback, the efficiency gains in our studies were distributed differently between the seller and bidders.</p>	<p>The analyses provide guidelines for managers to choose the appropriate form of feedback depending on the goals of the auction. If revenue maximization is the goal, then outcome feedback is preferred. However, outcome feedback leads bidders to place high bids, thus, squeezing them out of surplus and might alienate them in the long run. Therefore in the long run, price feedback is likely to be a better information regime.</p>

While most of the current research on combinatorial auctions focuses on either comparing novel combinatorial auction designs with multiple single-item auctions (e.g., Banks et al. 1989; Banks et al. 2003; Ledyard et al. 1997) or comparing several combinatorial auction designs (e.g., Brunner et al. 2007; Goeree and Holt 2008; Kwasnica et al. 2005), we studied how different types of information feedback influence the traditional metrics of interest, such as efficiency, seller's revenue, and bidder profits. In order to study these features, we first developed a simulation testbed that facilitated the creation of a robust experimental environment, including the appropriate choice of parameter values, so that the relevant effects can be isolated with minimal noise due to experimental instrument bias. To test the robustness of the results, we conducted the experiments with two different setups and achieved essentially the same results.

Our study adds to the active experimental exploration and design of new combinatorial auction mechanisms. The study showed that, with appropriate feedback, efficient continuous auctions can be conducted and, further, that information feedback plays a crucial role in the distribution of gains between the buyers and the sellers. One limitation of our study is that we only test the impact of feedback in a small auction setting with six items and three bidders. Since one of our goals was to study how different types of feedback influence bidding behavior, we maintained a simple bidding environment in order to avoid various confounding factors. However, even in this simple setting, we see significant differences in bidding behavior across different feedback regimes, leading to significant differences in outcome. Future research can test how bidding behavior changes with different designs, e.g., with more items and bidders, varying the parameters of the auction (e.g., super-additivity rate), modifying the distribution of assigned valuations of items. Furthermore, a thorough analysis of how bidder behavior is influenced by the nature of feedback might throw more light on the observed differences in

outcome of the auctions. Finally, the analysis of potential strategic interactions among the bidders represents yet another interesting direction for future research.

REFERENCES

- Adomavicius, G., and Gupta, A. 2005. Towards Comprehensive Real-Time Bidder Support in Iterative Combinatorial Auctions. *Information Systems Research* 16(2) 169-185.
- Andersson, A., Tenhunen, and M., Ygge, F. 2000. Integer Programming for Combinatorial Auction Winner Determination. *Proceedings of the 4th International Conference on Multiagent Systems, Boston, MA*.
- Annet, J. 1969. *Feedback and Human Behavior*. Harmondsworth, Middlesex, England: Penguin Books.
- Arora, A., Greenwald, A., Kannan, K. N., and Krishnan R. 2007. Effect of Information Revelation Policies Under Market Structure Uncertainty. *Management Science* 53(8) 1234-1248.
- Ausubel, L., Cramton, P., and Milgrom P. 2005. The clock-proxy auction: A practical combinatorial auction design. P. Cramton, Y. Shoham, R. Steinberg, eds. *Combinatorial Auctions*, Chapter 5. MIT Press, Cambridge, MA.
- Balzer, W., Doherty, M., and O'Connor, R. J. R. 1989. Effects of Cognitive Feedback on Performance. *Psychological Bulletin* 106(3) 410-433.
- Banks, J. S., Ledyard, J. O., and Porter, D. P. 1989. Allocating Uncertain and Unresponsive Resources: An Experimental Approach. *The Rand Journal of Economics* 20(1) 1-25.
- Banks, J., Olson, M., Porter, D. P., Rassenti, S. J., and Smith V. L. 2003. Theory, Experiment and the Federal Communications Commission Spectrum Auctions. *J of Economic Behavior and Organizations* 51(3) 303-350.
- Bapna, R., Goes, P., and Gupta, A. 2003. Analysis and Design of Business-to-Consumer Online Auctions. *Management Science* 49(1) 85-101.
- Brehmer, B. 1980. In One Word: Not from Experience. *Acta Psychologica* 45 223-241.
- Brehmer, B. 1990. Strategies in Real-Time, Dynamic Decision Making. R. Hogarth (Ed.) *Insights in Decision Making: A Tribute to Hillel J. Einhorn*. University of Chicago Press.
- Brewer, P. J., and Plott, C. R. 1996. A Binary Conflict Ascending Price (BICAP) Mechanism for the Decentralized Allocation of the Right. *International Journal of Industrial Organization* 14(6) 857-886.

- Brunner, C., Goeree, J. K., Holt, C. A., and Ledyard, J. O. 2007. An Experimental Test of Combinatorial FCC Spectrum Auctions. *Forthcoming in American Economic Journal: Microeconomics*.
- Budescu, D. V., and Maciejovsky, B. 2005. The Effect of Payoff Feedback and Information Pooling on Reasoning Errors: Evidence from Experimental Markets. *Management Science* 51(12) 1829-1843.
- Bykowsky, M. M., Cull, R. J., and Ledyard, J. O. 2000. Mutually Destructive Bidding: The FCC Auction Design Problem. *Journal of Regulatory Economics* (17) 205–228.
- Caplice, C. G. 1996. An Optimization Based Bidding Process: A New Framework for Shipper-Carrier Relationships. *Thesis, Dept. of Civil and Environmental Engineering, MIT, Cambridge, MA*.
- Chewning, J. E., and Harrell, A. M. 1990. The Effect of Information Load on Decision Makers' Cue Utilization Levels and Decision Quality in a Financial Distress Decision Task. *Accounting, Organizations and Society* 15(6) 527-542.
- Conover W.J. 1980. *Practical Nonparametric Statistics*. John Wiley and Sons, New York.
- Cox, C.R., Roberson, B., and Smith, V. L. 1982. Theory and Behavior of Single Object Auctions. V. L. Smith (Ed.) *Res. Exper. Econ.*, (2) JAI Press, Greenwich, CT.
- Cramton, P., Shoham, Y., and Steinberg, R. 2005. Introduction to Combinatorial Auctions. P. Cramton, Y. Shoham, R. Steinberg (Eds.) *Combinatorial Auction*. Introduction, MIT Press, Cambridge, MA.
- de Vries, S., and Vohra, R. V. 2003. Combinatorial Auctions: A survey. *INFORMS Journal on Computing* 15(3) 284-309.
- Doherty, M. E, and Balzer, W. K. 1988. Cognitive Feedback. B. Brehmer, C.R.B. Joyce (Eds.) *Human Judgment: The SJT Approach*. Amsterdam: North-Holland.
- Epstein, R., Henriquez, L., Catalan, J., Weintraub, G. Y., and Martinez, C. 2002. A Combinational Auction Improves School Meals in Chile. *Interfaces* 32(6) 1-14.
- Freidman, D., and Sunder, S. 1994. *Experimental Methods: A Primer for Economists*. Cambridge University Press.

- Garey, M. R., and Johnson, D.S. 1979. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. New York: W.H. Freeman.
- Goeree, J. K., and Holt, C. A. 2008. Hierarchical Package Bidding: A Paper & Pencil Combinatorial Auction. *Games and Economic Behavior* forthcoming.
- Grisé, M. L., and Gallupe, R. B. 1999. Information Overload: Addressing the Productivity Paradox in Face-to-Face Electronic Meetings. *Journal of Management Information Systems* 16(3) 157-185.
- Hammond, K. R, Stewart, T. R., Brehmer, B., and Steinmann, D. O. 1975. Social Judgment Theory. M. Kaplan, S. Schwartz (Eds.) *Human Judgment and Decision Processes*. New York: Academic Press.
- Hausch, D. B. 1986. Multi-Object Auctions: Sequential vs. Simultaneous Sales. *Management Science* 32(12) 1599-1610.
- Hoffman, P. J., Earle, T. C., and Slovic, P. 1981. Multidimensional Functional Learning (MFL) and Some New Conceptions of Feedback. *Organizational Behavior and Human Performance* 27, 75 - 102.
- Hogarth, R. M., Gibbs, B. J., McKinzie, C. R. M., and Marquis, M. A. 1991. Learning from Feedback: Exactingness and Incentives. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 17, 734 - 752.
- Isaac, R. M., and James, D. 2000. Robustness of the Incentive Compatible Combinatorial Auction. *Experimental Economics* (3) 31-53.
- Isaac, R. M., Salmon, T. C., and Zillante, A. 2007. A Theory of Jump Bidding in Ascending Auctions. *Journal of Economic Behavior* (62) 144-164.
- Jacoby, J. 1984. Perspectives on Information Overload. *The Journal of Consumer Research* 10(4) 432-435.
- Jamal, K., and Sunder, S. 1991. Money vs. Gaming: Effects of Salient Monetary Payments in Double Oral Auctions. *Organizational Behavior and Human Decision Processes* (49) 151-166.
- Kagel, J. H., Harstad, R. M., and Levin, D. 1987. Information Impact and Allocation Rules in Auctions with Affiliated Private Values: A Laboratory Study. *Econometrica* (55) 1275-1304.

- Kagel, J. H., and Levin, D. 2009. Implementing Efficient Multi-object Auction Institutions: An Experimental Study of the Performance of Bounded Rational Agents. *Games and Economic Behavior* (66) 221-237.
- Kluger, A. N., and DeNisi, A. 1996. The Effects of Feedback Interventions on Performance: A Historical Review, a Meta-Analysis, and a Preliminary Feedback Intervention Theory. *Psychological Bulletin* 119(2) 254-284.
- Koppius, O. R., and van Heck, E. 2003. Information architecture and electronic market performance in multidimensional auctions. *Working paper*, Erasmus University, Rotterdam, The Netherlands.
- Krishna, V. 2002. *Auction Theory*, Academic Press.
- Kwasnica, A. M., Ledyard, J. O., Porter, D. P., and DeMartini, C. 2005. A New and Improved Design for Multiobject Iterative auctions. *Management Science* 51(3) 419-434.
- Ledyard, J. O., Olson, M., Porter, D. P., Swanson, J. A., and Torma, D. P. 2002. The First Use of a Combined Value Auction for Transportation Services. *Interfaces* 32(5) 4-12.
- Ledyard, J. O., Porter, D. P., and Rangel, A. 1997. Experiments Testing Multiobject Allocation Mechanisms. *Journal of Economics and Management Strategy* 6(3) 639-675.
- Malhotra, N. K. 1982. Information Load and Consumer Decision Making. *The Journal of Consumer Research* 8(4) 419-430.
- McAfee, R. P., and McMillian, J. 1996. Analyzing the Airwaves Auction. *Journal of Economic Perspectives* 10(1) 159-175.
- O'Reilly III, C. A. 1980. Individuals and Information Overload in Organizations: Is More Necessarily Better? *The Academy of Management Journal* 23(4) 684-696.
- Parkes, D. C. 1999. iBundle: An Efficient Ascending Price Bundle Auction. *Proceedings of the First ACM Conference on Electronic Commerce* 148-157.
- Pekec, A., and Rothkopf, M. H. 2003. Combinatorial Auction Design. *Management Science* 49(11) 1485-1503.
- Porter, D. P. 1999. The Effect of Bid Withdrawal in a Multi-Object auction. *Review of Economic*

Design 4(1) 73-97.

Porter, D., Rassenti, S. J., Roopnarine, A., and Smith, V. L. 2003. Combinatorial Auction Design. *Proc. of the National Academy of Sciences of the USA* 100(19) 11153-11157.

Rassenti, S. J., Smith, V. L., and Bulfin, R. L. 1982. A Combinatorial Auction Mechanism for Airport Time Slot Allocation. *Bell Journal of Economics* 13(2) 402-417.

Rothkopf, M. H., Pekec, A., and Harstad, R. M. 1998. Computationally Manageable Combinatorial Auctions. *Management Science* 44(8) 1131-1147.

Sengupta, K., and Abdel-Hamid, T. K. 1993. Alternative Conceptions of Feedback in Dynamic Decision Environments: An Experimental Investigation. *Management Science* 39(4) 411-428.

Siegel, S., and Castellan, N. J. 1988. *Nonparametric Statistics for Behavioural Science*. McGraw Hill, New York.

Smith, V. L. 1976. Experimental Economics: Induced value theory. *American Economic Review* 66(2) 274-279.

Smith, V. L., and Walker, J. M. 1993. Rewards, Experience and Decision Costs in First Price Auctions. *Economic Inquiry* 31(2) 237-244.

Tversky, A., and Kahneman, D. 1974. Judgment under uncertainty: Heuristics and biases. *Science* 185 1124-1130.

Witten, I., and Frank, E. 2005. *Data Mining: Practical Machine Learning Tools and Techniques* (2nd ed.), Morgan Kaufmann.

Appendix A: Design of Real-Time Bidder Support

We build upon the real-time bid evaluation metrics developed by Adomavicius and Gupta (2005) that can present bidders with exact price information whenever a bidder wants to explore her alternatives, thereby providing a continuous environment. In this appendix, we provide technical details of the computational real-time bidder support capabilities.

Let I be the set of distinct items to be sold in a combinatorial auction, and let $N = |I|$. The terms *auction set* and *auction size* refer to I and N , respectively. Bidders can place bids on any *itemset*, which refers to any non-empty subset of I . A bid b can be represented by the tuple $b = (S, v, id)$. Here S denotes the itemset the bid was placed on ($\emptyset \subset S \subseteq I$), also called the *span* of the bid; v denotes the *value* of the bid ($v > 0$), e.g., the monetary amount specified in the bid; and id denotes the *bidder* who submitted this particular bid. Given bid b , $S(b)$, $v(b)$, and $id(b)$ are used to refer to the span, value, and bidder of the bid, respectively. We also use the notion of *auction states*. In particular, auction state k (where $k = 0, 1, 2, \dots$) refers to the auction after the first k bids are submitted. The bid set is denoted as B_k , i.e., $B_k = \{b_1, \dots, b_k\}$. Auction state 0 refers to the auction before any bids are made, i.e., $B_0 = \emptyset$. Obviously, $B_k \subseteq B_l$, for any k and l such that $k \leq l$.

Given an arbitrary set of bids B in a combinatorial auction, a bid set C (where $C \subseteq B$) is called a *bid combination in B* if all bids in C have non-overlapping spans, i.e., for every $b_x, b_y \in C$ such that $b_x \neq b_y$, we have $S(b_x) \cap S(b_y) = \emptyset$. Let C_k denote the set of all bid combinations possible at auction state k , or, more formally, $C_k = \{ C \subseteq B_k \mid b_x, b_y \in C, b_x \neq b_y \Rightarrow S(b_x) \cap S(b_y) = \emptyset \}$. We assume that the winners of the auction are determined by maximizing the seller's revenue, i.e., $\max_{C \in C_k} \sum_{b \in C} v(b)$, which is a standard assumption in the combinatorial auction research literature. The bid combination that maximizes this expression is called a *winning bid*

combination and is denoted as WIN_k (for auction state k). Moreover, given auction state k , bid $b \in B_k$ is called a *winning bid* in B_k if $b \in WIN_k$. Furthermore, if bid $b \in B_k$ is not a winning bid in B_k and cannot possibly be a winning bid in *any* subsequent auction state, then b is called a *dead bid* in B_k . Formally, bid $b \in B_k$ is dead if $b \notin WIN_k$ and $(\forall B_l \supseteq B_k)(b \notin WIN_l)$. The set of all dead bids in B_k is denoted as $DEAD_k$. On the other hand, if $b \notin DEAD_k$ then bid $b \in B_k$ is called a *live bid* in B_k . The set of all live bids in B_k is denoted as $LIVE_k$. Based on the definitions of WIN_k , $DEAD_k$ and $LIVE_k$, it is easy to see that: (i) $DEAD_k \cap LIVE_k = \emptyset$ and $DEAD_k \cup LIVE_k = B_k$, i.e., at any auction state k any bid $b \in B_k$ can either be live or dead, but not both; (ii) $WIN_k \subseteq LIVE_k$, i.e., every winning bid is obviously live; and (iii) $DEAD_k \subseteq DEAD_{k+1}$, i.e., once a bid becomes dead, it can never become live again.

Now, assume that an auction participant is interested in bidding on a bundle, say $X \subseteq I$. It is important for a bidder to know how much she should bid on X at a given time (i.e., at any auction state k), in order to guarantee that her bid is either winning or at least stands a chance of winning in the future (i.e., it is not dead). For this purpose the following bid evaluation metrics are used (Adomavicius and Gupta 2005):

1. Bid *winning level* (WL): for itemset X at auction state k , $WL_k(X)$ denotes the minimal value that auction participants have to bid on itemset X in order for this bid to be winning. In other words, after k bids have already been submitted, any bid b_{k+1} on itemset X that has value above $WL_k(X)$ will be winning, i.e., $v(b_{k+1}) > WL_k(X)$ where $X = S(b_{k+1})$ implies $b_{k+1} \in WIN_{k+1}$.
2. Bid *deadness level* (DL): for itemset X at auction state k , $DL_k(X)$ denotes the minimal value that auction participants have to bid on itemset X in order for this bid to be *live*. Similar to above, after k bids have already been submitted, any bid

b_{k+1} on itemset X that has value above $DL_k(X)$ will be *live*, i.e., $v(b_{k+1}) > DL_k(X)$

where $X = S(b_{k+1})$ implies $b_{k+1} \in LIVE_{k+1}$.

Appendix B: Screenshots of the Auction Interface

Elapsed Time: 00:03:01 (hh:mm:ss) Time since last bid: 00:00:47 (hh:mm:ss) Refresh

Select lots:

Specify your bid (\$): (no decimals) Submit Bid

Your valuation: \$165.00

Submit Bid

Bid History (Reverse chronological)

Bid No.	Bid Set	Bid Amount	Bid Time
6.	A	\$75.00	00:02:14
5.	BC	\$80.00	00:02:05
4.	ABC	\$140.00	00:01:58
3.	BC	\$80.00	00:01:42
2.	C	\$50.00	00:01:28
1.	AB	\$100.00	00:01:06

Figure B1. Auction interface for baseline feedback.

Elapsed Time: 00:01:11 (hh:mm:ss) Time since last bid: 00:00:13 (hh:mm:ss) Refresh

Select lots:

Specify your bid (\$): (no decimals) Submit Bid

Your valuation: \$165.00

Submit Bid

The winning bids are in bold red

Bid History (Reverse chronological)

Bid No.	Bid Set	Bid Amount	Bid Time
6.	A	\$75.00	00:00:57
5.	BC	\$80.00	00:00:48
4.	ABC	\$140.00	00:00:42
3.	BC	\$80.00	00:00:32
2.	C	\$50.00	00:00:19
1.	AB	\$100.00	00:00:10

Figure B2. Auction interface for outcome feedback.

Elapsed Time: 00:01:44 (hh:mm:ss) Time since last bid: 00:00:09 (hh:mm:ss) Refresh

Select lots:

Specify your bid (\$): (no decimals) Submit Bid

Your valuation: \$50.00

Submit Bid

The winning bids are in bold red

Bid History (Reverse chronological)

Bid No.	Bid Set	Bid Amount	Bid Time
6.	A	\$75.00	00:01:35
3.	BC	\$80.00	00:00:42
2.	C	\$50.00	00:00:29
1.	AB	\$100.00	00:00:14

To at least stand a chance of winning, bid: \$51.00

To be currently winning, bid: \$56.00

Figure B3. Auction interface for price feedback.

Appendix C: Descriptions of experimental interface elements marked in Appendix B.

Interface Element	Description
1	Any individual lot or a combination of lots could be selected by simply clicking on the checkboxes beside each lot. The amounts next to the checkboxes denote the valuations of the individual property lots. These amounts were displayed during the entire course of the auction.
2	This table displayed all the placed bids in the baseline feedback case. All the bids placed by a particular bidder were highlighted on his/her screen.
3	This label displayed the valuation of the selected individual lot or bundle. The valuation of the bundle {B,C} in this example is \$165.00.
4	This table displayed all the placed bids in the outcome feedback case. All the bids of placed by a particular bidder were highlighted on his/her screen as in the baseline case. Furthermore, all the provisionally winning bids were identified in bold red at all stages of the auction.
5	This table displayed all the non-losing bids in the price feedback case. All losing bids were removed from display (e.g., as can be seen from this table, bids 4 and 5 are not displayed). All the bids placed by a particular bidder were highlighted on his/her screen as in the other two cases. Further, all the provisionally winning bids were identified in bold red at all stages of the auction as in outcome feedback case.
6	This label displayed the minimum price for placing a non-losing bid on a chosen bundle. In this example, the current minimum price for placing a non-losing bid on Lot C is \$51.00.
7	This label displayed the minimum price for placing a winning bid on a chosen bundle. In this example, the current minimum price for placing a winning bid on Lot C is \$56.00.