Multi-Object Auctions with Package Bidding: An Experimental Comparison of Vickrey and *i*BEA *

Yan Chen Kan Takeuchi

December 27, 2005

Abstract

We study two package auction mechanisms in a laboratory setting, a sealed bid Vickrey auction and an ascending version of Vickrey, the *i*BEA auction. Unlike the single-unit Vickrey auction, where bidders tend to overbid in the laboratory, most of our bidders either underbid or bid their true values. Furthermore, at the aggregate level, Vickrey generates significantly higher revenue and efficiency than iBEA. We also find that human bidders learn from their robot opponents when the robot strategies are (myopic) best responses.

Keywords: combinatorial auctions, procurement auctions, experimental economics
JEL Classifications: C91, D44

*We would like to thank Ted Bergstrom, Vincent Crawford, Tom Finholt, Ted Groves, Dan Houser, Elena Katok, Tony Kwasnica, John Ledyard, Dan Levin, Kevin McCabe, Lydia Mechtenberg, David Parkes, David Reiley, Tatsuyoshi Saijo, Tim Salmon and seminar participants at Meinheim, Michigan, Penn State, Shanghai Jiaotong, UC Santa Barbara, UC San Diego, the 2004 Economic Science Association Meetings (Amsterdam and Tucson), and the Eighth Experimental Economics Conference at Osaka University (November 2004) for helpful discussions. We thank Alexandra Achen for excellent research assistance. Chen gratefully acknowledges financial support from the National Science Foundation through grant no. SES-0339587.
Chen: School of Information, University of Michigan, 1075 Beal Avenue, Ann Arbor, MI 48109-2112. Email: yanchen@umich.edu. Takeuchi: Department of Economics, University of Michigan, 611 Tapan Street, Ann Arbor, MI 48109-1220. Email: ktakeuch@umich.edu.

1 Introduction

In recent years, the use of various multi-object auction mechanisms to determine resource allocation has been rapidly increasing. In particular, the Federal Communications Commission (FCC) spectrum auctions, characterized by synergies across licenses, stimulated tremendous research interests in complex auction design for multiple objects and with synergies. Since 1994, the FCC has used the Simultaneous Multiple Round (SMR) auction to allocate spectrum licenses and raised over $41 billion in revenue (Milgrom 2004). However, this auction format does not allow package bidding. Ledyard, Porter and Rangel (1997) demonstrate that the performance of the FCC design degrades in the presence of complementarities. More generally, when bidder valuations for multiple objects are super-additive, package bidding is necessary to increase efficiency, seller revenue and bidder willingness to participate (Bykowsky, Cull and Ledyard 2000). Economic research on package auction mechanisms helps policy makers such as the FCC as well as others who must allocate complex resources. For example, FCC auction No. 31 for selling spectrum licenses in the 700 MHz band, was designed to permit bids for any of the 4095 possible packages of the twelve licences on offer.[1]

The earliest research on package auction design is a proposal for a sealed bid combinatorial auction sale of paired airport takeoff and landing slots proposed by Rassenti, Smith and Bulfin (1982). In another study, Banks, Ledyard and Porter (1989) present two kinds of iterative package auctions to allocate uncertain and unresponsive resources, AUSM (Adaptive User Selection Mechanism) and the iterative Vickrey-Clarke-Groves mechanism (Vickrey (1961), Clarke (1971), Groves (1973)). In laboratory experiments, these package auctions significantly outperform markets and administrative procedures. More recently, Kwasnica, Ledyard, Porter and DeMartini (2005) create and test a new design for multi-object iterative auctions by merging the better features of the AUSM and the FCC SMR designs. The resulting new Resource Allocation Design (RAD) is shown to perform better than either parent.

Another important application of package auction design is the Business-to-Business (B2B) auction, which is predominantly multi-object and often involves synergies. Package auctions have the potential to provide value to both buyers and sellers of goods and services. The descending Dutch auction used to sell flowers in Aalsmeer, Holland is a simple version of a package auction for homogeneous goods (Katok and Roth 2004). More recently, package auctions have been successfully applied to transportation procurement. For example, Sears Logistics Services is the first procurer of trucking services to use a package auction to reduce its costs. Since 1993, it consistently saved 13 percent over past procurement practices, mostly through bilateral negotiations (Ledyard, Olson, Porter, Swanson and Torma 2002). Finally, the London bus routes provide an example of the use of a package auction format in public procurement. In this instance, the local transportation authority has adopted a form of package auction because of expected economic synergies among routes located in the same area of London. The London bus routes auction has led to increased quality of service and lower costs, and thus is considered a success (Cantillon and Pesendorfer 2006).

Empirical investigations of package auctions have important implications not only in procurement and privatization, but also potentially in the allocation of scarce equipment time

[1] See http://wireless.fcc.gov/auctions/31/ for more information.

in large scientific collaboratories.[2] The National Science Foundation and other federal agencies are poised to make significant investments in expanding the ability of geographically-distributed groups of scientists to conduct research via the Internet. There are currently over eighty collaboratories in use within multiple scientific communities, including space physics, HIV/AIDS, software engineering and neuroimaging (Finholt 2002). In many collaboratories, a critical feature of the equipment time allocation problem is that contiguous time slots are more valuable than the sum of separate slots, i.e., user valuation for multiple slots exhibits synergy. Therefore, package auctions might be an important mechanism in achieving efficient allocation of equipment time.

Despite successful results for package auctions, the theoretical properties of these auctions are not often understood. Therefore, it is unclear whether a given choice of design option is the most appropriate one.

An important standard for nearly all mechanism design work, and for auctions in particular, is the Vickrey-Clarke-Groves (VCG) mechanism. The VCG mechanism is dominant strategy incentive compatible, i.e., bidding one's true valuation is always optimal regardless of others' strategies. Furthermore, it implements the efficient outcome. For the single object case, VCG mechanism becomes the familiar second-price auction. In the auction context, we follow convention and call the VCG mechanism Vickrey auctions.

However, despite its attractive theoretical properties, the Vickrey auction has some disadvantages. In the package auction context, there are three main concerns.[3] First, the Vickrey auction might be vulnerable to collusion. For example, bidders have the incentive to use shill bidders through which they could manupulate the allocation and prices in their favor. Second, it might suffer from the monotonicity problem,[4] i.e., adding bidders might reduce equilibrium revenues. Third, previous laboratory experiments show that, in the single object case, the dominant strategy in second-price auctions is not transparent. Many experimental subjects consistently overbid in second-price auctions and do not seem to learn from prior experience (see, e.g., Kagel (1995)).

Contrary to the economist belief that Vickrey auctions are rarely used in practice, Lucking-Reiley (2000b) presents evidence that Vickrey auctions have long been the predominant auction format for mail sales of collectible postage stamps, at least 65 years earlier than the publication of Vickrey's seminal paper. Vickrey-like auctions have also been appearing in auctions on the Internet, sometimes with an additional feature of "proxy bidding." In these auctions, a bidder tells his proxy his maximum willingness to pay. The proxy keeps this information secret and bids on the bidder's behalf in an ascending auction in a pre-announced increment. If every bidder uses a proxy, then the bidder with the highest maximum price wins and pays (approximately) the second highest price. The most prominent examples of such auctions include Amazon and eBay.[5]

To retain the advantages of the sealed-bid Vickrey auctions while reducing their dis-

[2] First proposed in the late eighties, a collaboratory is a center without walls, in which researchers can perform their research without regard to physical location - interacting with colleagues, accessing instrumentation, sharing data and computational resources, and accessing information in digital libraries (Wulf 1993).

[3] Note the first two problems do not appear when all goods are substitutes for all bidders.

[4] See Milgrom (2004) Chapter 8 for examples.

[5] Ockenfels and Roth (2002) analyze the closing rules of eBay and Amazon. They show that eBay auctions, with a hard closing rule, give bidders incentives for sniping, while Amazon auctions, with a soft closing rule, do not give such incentives and hence are a more faithful dynamic version of Vickrey.

advantages, researchers have searched for an ascending bid package auction with comparable theoretical properties. One example is the Ausubel-Milgrom Ascending Proxy auction (Ausubel and Milgrom 2002), which replicates the performance of the Vickrey auction when all goods are substitutes, and has full-information equilibrium outcomes that are bidder-optimal points in the core. Another example is the iBundle Extend & Adjust (*i*BEA) auction (Parkes and Ungar 2002), an ascending bid auction with package bidding which implements efficient allocation and Vickrey payments in *ex post* Nash equilibrium, with only a free-disposal requirement on agent preferences. Thus, theoretically, *i*BEA represents a major advance in modeling package auctions.

In this paper, we investigate two package auction mechanisms in the laboratory, Vickrey and *i*BEA, to evaluate their performance among boundedly rational individuals. Our study advances research on auctions in two ways. First, compared to previous experimental studies of package auctions, the theoretical properties for these two auctions are well understood. Second, this is the first experimental study of *i*BEA, a new promising ascending bid package auction mechanism.

We study these two mechanisms in a simple environment where three bidders compete for four items (and thus 15 packages), with synergies across subsets of the items. Each human bidder competes against two automated bidders. In half of the treatments, the automated bidders are programmed to follow the dominant strategy in Vickrey, and Myopic Best Response in *i*BEA, while in the other half of the treatments, automated bidders are programmed to follow random, or zero-intelligence strategies. The use of automated agents serve two purposes. First, it allows the experimenter to compare the performance of the two mechanisms in an environment free from the strategic uncertainties inherent in interactions between human bidders.[6] Second, the use of automated agents is becoming increasingly widespread in Internet auctions, which allow for conveniently asynchronous bidding (Lucking-Reiley 2000a). Therefore, it is important to study how humans react when they bid against automated agents.

The main results of our two auction scenarios present some unexpected findings. First, contrary to previous experimental studies of sealed bid vs. ascending bid auctions, where ascending auctions tend to outperform their sealed bid counterpart, in our study, the sealed bid Vickrey auctions generate significantly higher efficiency and revenue than the *i*BEA auctions. Second, unlike in second-price auctions where most participants overbid, most participant in our Vickrey package auctions either underbid or bid their true values. Finally, we find that human bidders learn from automated agents that use "intelligent" strategies.

The rest of the paper is organized as follows. Section 2 introduces the auction mechanisms. Section 3 presents the experimental design. Section 4 presents the hypotheses. Section 5 presents the analysis and main results. Section 6 concludes.

2 The Auctions

In this section, we introduce our two auction mechanisms. To do so, we first set up a simple framework that allows us to explain the auctions clearly.

[6] See, e.g., Kagel and Levin (2001) for an experiment with human bidders against automated agents in multi-unit auctions of homogeneous goods.

In this framework, let $N = \{1, \ldots, n\}$ be a finite set of bidders. Let i denote an agent, where $i = 0$ is an *auctioneer* and $i > 0$ is a bidder. $N_0 = N \cup \{0\}$ is the set of all bidders and the auctioneer. Let $K = \{1, \ldots, k\}$ represent the set of objects to be sold, and $X = \{0, 1\}^k$ represent the set of combinations of objects. Let $\mathcal{B}_i$ be a set of *bids* of bidder i. Each bid is a pair, (x, p), where $x \in X$ corresponds to the packages desired and $p \in \mathbb{R}$ is the bid price. Let $v_i : X \to \mathbb{R}_+$ be bidder i's *valuation function* that assigns a value to a package. Finally, $\delta = (\delta_1, \ldots, \delta_b)$ is an indicator vector, where $\delta_j \in \{0, 1\}$ indicates whether bid (x_j, p_j) is *winning* or *losing*, and where b is the total number of bids.

As the winner determination problem is part of every package auction design, we formally define this problem for our experiment.

Definition (Package auction winner determination problem). The package auction winner determination problem is to maximize the sum of all bids, indicating each bid as *winning* or *losing*, under the constraint that each item can be sold to at most one bidder:

$$\max_{\delta} \sum_{j=1}^{b} \delta_j p_j \quad \text{subject to} \quad \sum_{j \in \{j : \delta_j = 1\}} x_j \leq (1, 1, \ldots, 1). \tag{1}$$

One of the most important goals in multi-object auctions is for all bidders to bid truthfully. We now define a truthful bid: a bid $(\mathbf{x}, p) \in \mathcal{B}_i$ is a *truthful bid*, if $v_i(\mathbf{x}) = p$. The allocation associated with the solution to the winner determination problem maximizes the aggregate surplus, when every bidder bids for all packages and the bids are all truthful,

Although it is desirable that all bidders truthfully bid for all packages, in theory, depending on the auction mechanism, it might not always be optimal for a bidder to bid her true value, and in practice, it is not easy for bidders to calculate all values for all possible packages. To design a good auction mechanism, we would like to obtain the most efficient allocation among feasible outcomes. In the design of a multi-object auction mechanism, there are three potential problems to overcome.

1. The exposure problem: When items are not substitutes and bidders cannot bid on packages, bidders are usually exposed to the risk that they may overpay. For example, suppose that there are two items and that bidder i's valuations exhibit strong complementarity, such that $v_i((1, 0)) = v_i((0, 1)) = 1 < 3 = v_i((1, 1))$. He may bid on each item at prices more than 1, expecting that he gets both of the items. However, it is possible that he could get only one item while paying more than 1. Auctions with package bidding, such as the two auctions in this study, should overcome this problem.
2. The threshold problem: Suppose that there are four bidders and three items to trade and that $v_1((1, 0, 0)) = v_2((0, 1, 0)) = v_3((0, 0, 1)) = 1.5$ and $v_4((1, 1, 1)) = 4$. In this setting, it is efficient to allocate the items to Bidders 1, 2 and 3. However, suppose that Bidders 1, 2 and 3 bid on their desirable item at price 1, respectively, and that Bidder 4 bids on the package of all items at price 3.6. In this case, the winning bid is Bidder 4's bid, but none of Bidders 1-3 can overbid on Bidder 4's bid. The ascending bid auction, *i*BEA, overcomes the threshold problem, which we will explain in Section 2.

3. Computational problems: There are two computational problems in package auctions, the auctioneer's and the bidders'. For the auctioneer, the problem is to solve Eq. (1), which is NP-complete (Rothkopf, Pekec and Harstad 1998).[7] For the bidders, they must evaluate all possible combinations, which can be cognitively difficult. Evaluation of the latter is where experimental research can be especially valuable.

In the following sections, we first introduce the Vickrey and *i*BEA auctions, and then discuss their respective theoretical properties.

2.1 Vickrey Auction with Package Bidding

A Vickrey auction with package bidding is an extension of the more familiar second-price auction. At the beginning of each auction, each bidder selects the packages he would like to bid on, and the amount he would like to bid for each package. Each bidder can choose to bid on as many packages as he wants, and he can bid on a single object multiple times, by bidding on several packages that contain that item. However, no matter how many packages a bidder bids on, he will never win more than one package. This type of bidding is called an exclusive-or (XOR) bid. Since all items are weakly complements in our experiment, XOR is not necessary for incentive compatibility. However, to minimize the interface difference between the two mechanisms, we impose XOR bids in the Vickrey auction, as they are required in the *i*BEA auction.

Next, once all bidders have submitted their bids, the auctioneer will choose the combination of submitted bids that yields the highest sum of bids. The set of bidders winning a package are the winning bidders.

After determining the winning bidders, the auctioneer then, one at a time, chooses each winning bidder as a **pivotal bidder**. The auctioneer examines the bids again, but ignores the bids of the pivotal bidder. The auctioneer determines the allocation of goods that maximizes the sum of bids, using the same rules as before, but not considering any bids placed by the pivotal bidder. Once this new allocation has been determined, the auctioneer compares the sum of bids generated by this allocation with those generated when no bids are excluded.

At the end of the auction, the amount that the winning bidders are required to pay depends on the additional revenue that each bidder generated, which is calculated by comparing the original revenue obtained by the auctioneer versus the revenue obtained by the auctioneer when the given bidder is pivotal. The following example from the experimental instruction illustrates how this process works. The fictitious currency used in this example (and throughout the experiment) is pounds (£).

[7] A problem is in P if it is solvable in polynomial time by a deterministic Turing machine (e.g., a program on a conventional computer), and a problem is in NP if it is solvable in polynomial time by a non-deterministic Turing machine. While P-problems are also in the class of NP, it is believed that NP problems are not in P. That is, NP problems are not solvable in polynomial time by a conventional computer in the worst case. Furthermore, an NP-complete problem is the "hardest" among all NP problems in the sense that any NP problem can be reduced into the NP-complete problem using a deterministic Turing machine. In sum, in combinatorial auctions, the auctioneer needs to solve an NP-complete problem that is believed to be unsolvable in polynomial time (Weisstein 2002).

Example: Suppose that there are three bidders and four objects to allocate, and the following bids are submitted:

	Package	Price	Status
Bidder 1	**AB**	£50	winning
Bidder 2	**CD**	£40	winning
Bidder 3	**ABCD**	£60	
Bidder 3	**AB**	£30	

As shown in the table, the bids from Bidder 1 and 2 are the winning bids, because they generate the highest revenue for the auctioneer £50 + £40 = £90.

However, the auctioneer does not ask Bidder 1 to pay £50. Suppose we choose Bidder 1 as a pivotal bidder, and ignore his bids. The winning bids then become Bidder 2's bid on **CD** and Bidder 3's bid on **AB**.

	Package	Price	Status
~~Bidder 1~~	~~**AB**~~	~~£50~~	~~winning~~
Bidder 2*	**CD**	£40	winning
Bidder 3	**ABCD**	£60	
Bidder 3*	**AB**	£30	winning

In this case, the auctioneer calculates the revenue that those winning bids would generate, which is £40 + £30 = £70. Thus, the additional revenue that Bidder 1 generates is £20, since £90 – £70 = £20. This £20 is the price adjustment for Bidder 1. Therefore, Bidder 1 pays £50 and receives £20 back. His final price is £30.

The **Vickrey auction** is dominant strategy incentive compatible and implements an efficient outcome. Each bidder ends up with a Vickrey payoff in equilibrium. However, it has some shortcomings. First, it is vulnerable to collusion. Second, the revenue under a Vickrey auction can be very low. Third, the dominant strategy in a Vickrey auction might not be transparent when it is implemented with boundedly rational people. Many previous experimental studies on single unit auctions (see Kagel (1995) for a survey) demonstrate that bidders systematically overbid in single-unit Vickrey auctions. More recently, Isaac and James (2000) show that, in an experimental setting with two items (and thus three packages), Vickrey auctions with package bidding consistently generate higher efficiency than those without package bidding. However, they do not compare a Vickrey package auction with an ascending package auction. Based on previous experimental results which mostly concentrate on single-unit auctions, we expect the ascending bid auctions to achieve higher efficiency than the Vickrey auction in the multi-object setting with package bidding.

2.2 *i*Bundle Extend & Adjust (*i*BEA) Auction

The *i*Bundle Extend & Adjust (*i*BEA) auction is proposed by Parkes and Ungar (2002). It is an ascending-price generalized Vickrey auction. It maintains non-linear and non-anonymous prices on packages, and terminates with approximately efficient allocation and Vickrey payments. To achieve these properties, the mechanism requires a myopic best response, which is an *ex post* Nash equilibrium. In an *i*BEA auction, each auction takes place in several rounds. Let ε be the price increment. The choice of ε involves a tradeoff between the speed of the auction and the closeness to efficiency of the final outcome. That is, a smaller ε can achieve more efficient outcomes, at the cost of a longer auction. An outline of the *i*BEA process is as follows:

1. The auctioneer initializes prices for all packages.
2. At the beginning of each round, for each bidder, the auctioneer announces "ask prices" $\mathbf{p}_i(t)$ for all packages to bidder i.
3. Given the prices, each bidder can bid on as many packages as she wants, and can bid on a single item multiple times, by bidding on several packages that contain that item. Each bidder's submitted bids, $\mathcal{B}_i(t)$, must satisfy the following rules:

   (a) Winning bid resubmission rule: $\exists(x, p) \in \mathcal{B}_i(t)$, such that (x, p) is a winning bid in the previous round. That is, if a bidder has made a winning bid in the previous round, she is obligated to bid on that package, at the same price, in the next round. Once a bid is losing, a bidder has no further commitment to bid on that package, unless he chooses to do so.

   (b) Last-and-final bid: The last-and-final option allows a bidder to continue to bid for a package when the bid price is narrowly above the object's value. Once a bidder chooses the last-and-final option on a package, he receives a small discount, ε, and the last-and-final bid is automatically resubmitted at the same price in every round until the auction terminates. Therefore, even if the price for that package increases, a bidder cannot increase his bid on that package.

   The last-and-final option facilitates the computation of Vickrey prices. It also reveals to the auctioneer a bidder's approximate true value for a package.

4. If there are no new bids, then the auction terminates. Otherwise, given $\{\mathcal{B}_i(t)\}_{i \in N}$, the auctioneer solves Eq. (1) and revises ask prices to each bidder in the following manner:

   (a) The ask prices for a bidder with any winning bid(s) remain the same as in the previous round.

   (b) The prices for packages with last-and-final bids remain the same as in the previous round.

   (c) The price for each package of a losing bid increases by ε.

   (d) All ask prices are adjusted to be self-consistent, i.e., the price for a package should not be less than the price for any of its subsets.

   The auctioneer proceeds to the next round $t + 1$, then returns to step 2.

In the *i*BEA auction, there are two phases to determine the auction's outcome. The auctioneer uses Phase I to determine the final allocation and Phase II to compute the final prices and Vickrey discounts. Phase I terminates when all agents who submit bids are assigned a package, and the allocation is the final allocation. The auction then proceeds to

Phase II. In each round of Phase II, the auctioneer selects a pivotal bidder and ignores this bidder's bids to compute his *externality*. When she computes all externalities, the prices are determined, from which she computes the Vickrey payoffs, in the same manner as in the sealed bid Vickrey auction.

Just as Ausubel and Milgrom (2002) rely on the assumption of a straightforward bidding strategy, in an *i*BEA auction, it is assumed that each bidder takes a *Myopic Best-Response* (MBR). We say bidder i takes a myopic best-response to the ask prices $\mathbf{p}_i(t)$ announced by the auctioneer, if

$$\mathcal{B}_i(t) = \left\{ (x,p) \middle| v_i(x,p) - p \geq \max \left\{ \max_{z \in X} \left\{ v_i(z, p(z)) - p(z) \right\}, 0 \right\} - \varepsilon \text{ and } p = p(x) \right\}. \quad (2)$$

where $p(x)$ is the ask price of package x for bidder i. The MBR strategy chooses ε-maximized packages, i.e., packages arbitrarily close to the best package.

We now explain MBR in more detail. We define a bidder's *temporary profit* as $v_i(x, p(x)) - p(x)$. Therefore, when a bidder examines his menu, he first looks for any packages on which he has a negative temporary profit. For those packages that have a negative profit, he adds ε to his temporary profit, because he knows that he can buy the package with an ε discount using the last-and-final option. He does not change his temporary profit for a package on which he currently has a positive profit. He then examines the revised temporary profits for all packages, and finds the package that gives him the greatest revised temporary profit. There are two possibilities:

1) If his greatest possible profit is greater than or equal to ε, he will bid on all packages that give him a revised temporary profit within ε of the maximum temporary profit. For example, if one package gives him a temporary profit of £17, no package gives him a profit of more than £17, and $\varepsilon = \pounds 5$, then he will bid on all packages that give him a temporary profit of at least £17 – £5, or £12.
2) If his greatest possible profit is less than ε, he will bid only on those packages that have a revised temporary profit greater than or equal to 0.

Under the assumption of a MBR, we can achieve competitive equilibrium prices and efficiency. More precisely, if bidders follow the MBR, then the vector of ask prices $\{\mathbf{p}_i(t)\}_{i \in N}$ is a competitive equilibrium price vector after the end of Phase I. As $\varepsilon \to 0$, the final allocation becomes arbitrarily close to the efficient allocation. Furthermore, Phase II allows the auctioneer to compute Vickrey payoffs. Parkes and Ungar (2002) prove that the MBR is incentive compatible: MBR is an *ex post* Nash equilibrium of *i*BEA, as $\varepsilon \to 0$.

To illustrate how the *i*BEA auction works, we use a simple three-bidder, two-item example. In this example, we assume that all bidders follow a MBR strategy. Furthermore, we set $\epsilon = 5$.

[Table 1 about here.]

Table 1 illustrates how *i*BEA works. The top panel presents Phases I and II round-by-round and the bottom panel compares the results with those of the Vickrey auction.

As shown in Table 1, for Bidder 1, the values of **A**, **B** and **AB** are (10, 0, 10). For Bidders 2 and 3, they are (0, 30, 30) and (4, 14, 22), respectively.

In round 1, the offered prices are the same. The auctioneer breaks the tie by randomly choosing as many winning bidders as possible, and Bidders 1 and 2 are selected. Since Bidder 3 does not win any package, the auctioneer raises the offered prices for Bidder 3. Specifically, the price for the losing bid on **AB** is raised by an increment of 5.

In Round 2, the highest temporary profit of Bidder 3 is $17 = 22 - 5$. Bidder 3, following the MBR, is willing to bid on any package whose temporary profit is above $12 = 17 - 5$. Bidder 3 bids on (**AB**, 5) and (**B**, 0), since the temporary profit of each package is 17 and 14, respectively. The auctioneer chooses (**AB**, 5) from Bidder 3 as the winning bid for Round 2. The process continues.

Notice, however, that Bidder 3 makes last-and-final bids on (**A**, 5) and (**B**, 15) in Round 8, because the offered prices are above the value of those packages. These bids are submitted with a discount of 5 and become *(**A**, 0)* and *(**B**,10)*, respectively. At the end of Round 9, there is no new bid. Consequently the allocation of items is finalized: Bidder 1 receives **A** and Bidder 2 receives **B**.

The auction proceeds to Phase II. In this phase, the auctioneer randomly chooses one of the winners as a pivotal bidder. Let us select Bidder 2 in this example. The auctioneer excludes Bidder 2's bids from the bids submitted in Round 9 and selects a new set of wining bids. The new winning bid is ***(AB, 20)*** from Bidder 3. So, the auctioneer raises the offered prices for Bidder 1 and proceeds to Round 10. In Round 10, (**A**,10) from Bidder 1 and (**B**, 10) from Bidder 3 are selected as the winning bids, and there is no new bid. Choosing Bidder 1 as the next pivotal bidder, the auctioneer goes through the same process. At the end of Round 11, since there is no new bid and no winner for the next pivotal bidder, the auction terminates.

Once the auction has ended, the auctioneer determines the price adjustment (rebate) for Bidders 1 and 2. Note that the revenue is 30 before any price adjustment. Based on the offered prices at the end of the auction, the auctioneer calculates the additional revenue that each winner generates. When Bidder 1 is pivotal, Bidder 2 receives **B** and Bidder 3 receives **A** and the revenue is 20. Thus, the additional revenue from Bidder 1 is $10 = 30 - 20$, which is the price adjustment for Bidder 1. Similarly, the price adjustment for Bidder 2 is 10. All payoff information is summarized in the bottom panel of Table 1.

Although *i*BEA has more desirable properties than any other ascending package auctions, it has not been tested in a laboratory setting. Past experiments on single-unit auctions show that the ascending bid auction, e.g, English clock auction, achieves higher efficiency than the sealed bid Vickrey auction, even though the solution concept is weaker (Kagel 1995). The ascending auction provides more feedbacks, which makes the optimal strategy more transparent than its sealed bid counterpart. Therefore, it is interesting to see whether this superior performance of the ascending bid auction carries over to a multi-object package bidding context.

3 Experimental Design

Our experimental design reflects both theoretical and technical considerations. Specifically, we are interested in three important questions. First, how do the Vickrey and *i*BEA auctions compare in performance? Second, how do human subjects respond to the degree of rationality in the environment? Third, do human subjects imitate bidding strategies of their robot opponents in the auction settings? We describe our experimental environment and procedures below.

3.1 The Economic Environment

In each auction, three bidders compete for four items (**A**, **B**, **C** and **D**), and thus 15 packages.[8] Each human subject competes against two automated bidders, i.e., computer programs (or "robots"). In each session, subjects are informed that they interact with robots. Using robots gives us more control over each human subject's environment. For the human participants, this environment reduces the strategic uncertainties inherent in interactions between human bidders. Furthermore, it allows us to observe human subject reactions when they interact with robots. The latter, in itself, has important implications for e-commerce (Eisenberg 2000).

We implement a $2 \times 3 \times 2$ design. In the first dimension, we compare the two mechanisms, the Vickrey and *i*BEA auctions. In the second dimension, we implement three combinations of robot bidding strategies. As Sincere bidding (S) leads to efficient allocation in both mechanisms, while Random bidding (R) represents zero-intelligence, we set up three different combinations of these strategies, SS, SR and RR. In SS, both robots follow Sincere bidding; in SR, one follows Sincere and the other follows Random bidding; and in RR, both follow Random bidding. As we are interested in whether a subject will imitate a certain strategy when told one of the robots follows such a strategy, we design a third dimension with two information conditions. In the low information treatment, subjects are told that they are competing against robots; however, the robot bidding strategies are not explained to the subjects. In the high information condition, we explain the robot strategies in the instructions.[9]

We now describe bidder preferences in our experimental setting. Let V_i be the value of package i. The value for each item is drawn independently from a uniform distribution on $\{0, 1, 2, \cdots, 10\}$. The value of a package is the sum of the values of the items in the package plus the bonus value for certain combinations of items due to synergy. A human bidder and the first robot derive synergy from **A** and **B**. In the experiment, we choose a CES function to represent this preference. Thus, the value of A and B together equals $V_{AB} = (V_A^\rho + V_B^\rho)^{1/\rho}$. The CES function allows us to control for the degrees of complementarity and substitutability. When $\rho > 1$, $V_{AB} < V_A + V_B$. When $\rho = 1$, $V_{AB} = V_A + V_B$. When $0 < \rho < 1$, $V_{AB} > V_A + V_B$, i.e., **A** and **B** have synergy. In the experiment, we choose $\rho = 0.9$, as we are interested in the case when synergy is present. Similarly, the second robot derives synergy from **C** and **D**. For example, if a human bidder or Robot 1 has package **ABCD**, the value equals $V_{ABCD} = V_{AB} + V_C + V_D$. If Robot 2 has the same package, the value equals $V_{ABCD} = V_A + V_B + V_{CD}$. Since the human bidder and Robot 1 have the same preference, the environment is more competitive for each than it is for Robot 2.

In the *i*BEA auction, we set $\varepsilon = \pounds 5$, since the smallest grid size is £1, and $\varepsilon = \pounds 5$ generates a reasonable speed of convergence in the lab. In Section 5, we use simulations to compare the performance of the two mechanisms using the actual grid size of $\varepsilon = \pounds 1$ for Vickrey and $\varepsilon = \pounds 5$ for *i*BEA, as well as the same grid size of $\varepsilon = \pounds 5$ for both auctions.

[8]These packages are **A**, **B**, **C**, **D**, **AB**, **AC**, **AD**, **BC**, **BD**, **CD**, **ABC**, **ABD**, **ACD**, **BCD**, and **ABCD**.

[9]The instructions are quite lengthy; they are available at http://www.si.umich.edu/~yanchen/

3.2 Experimental Procedures

Our experiment involves 8 to 10 human subjects per session. At the beginning of each session, each subject is given printed instructions. After the instructions are read aloud, subjects are encouraged to ask questions. The instruction period takes, on average, 21 minutes for *i*BEA and 14 minutes for Vickrey. After receiving the instructions, subjects take a quiz designed to test their understanding of the mechanisms. At the end of the quiz, the experimenters go through the answers with the group of subjects. The quiz takes an average of 17 minutes for *i*BEA and 11 minutes for Vickrey. At the end of the quiz, subjects randomly draw a PC terminal number. Each subject then sits in front of the corresponding terminal and starts the experiment.

In each session, each subject participates in 10 auctions. As we are interested in the effects of learning on bidding behavior, there are no practice auctions. At the beginning of each auction, the value for each item is randomly drawn from a uniform distribution on $\{0, 1, 2, \cdots, 10\}$ for each bidder.

In the Vickrey treatments, each subject is informed of the value of the 15 packages on his screen. He can enter an integer bid for any of the 15 packages.[10] Note that a zero bid on a package is treated differently from no bid. A subject can be allocated a package on which he bids zero, but can never be allocated a package that he does not bid on. Meanwhile, each robot submits a bid on each package, following the pre-assigned strategy. A robot following the Sincere strategy bids its true value for each package, while a robot following the Random strategy randomly chooses a number between 0 and 120% of its value for each package. The upper bound is chosen based on previous experimental evidence on bidding range in single-unit Vickrey auctions (Kagel 1995). The server collects all bids from each group, computes the final allocation and payoff for each bidder and sends this information back to the bidder's screen. Each human bidder gets the following information at the end of each auction: his allocation, his price and price adjustment, his value for the allocated package, his profit for this auction and his cumulative profit.

In the *i*BEA treatments, each subject is given a menu, which contains the following columns: the package, his value for each package, his price, his temporary profit, two check boxes (whether he wants to bid on a package, and whether he wants his bid to be the last-and-final bid) and the status of his previous bid (winning, losing, last-and-final and winning, and last-and-final and losing). The bidder may check either, both or neither of the two checkboxes. Recall that for winning bids and last-and-final bids, the checkboxes are automatically checked.

In the *i*BEA auction, a robot following the Sincere strategy adopts the MBR. A robot following the Random strategy examines his menu and looks for packages for which he has an original temporary profit of greater than −£5. For each package that yields at least a temporary profit of −£5, he flips a coin to decide whether or not to place a bid on that package.

Robots choose the last-and-final option for any package they place bids on that have an original negative temporary profit.

Then the auction proceeds, as described in Section 2.

[Table 2 about here.]

[10] The zTree program sets a lower bound of zero and an upper bound of 1000 for the bids.

Table 2 presents the relevant features of the experimental sessions, including mechanisms, robot strategy profiles, information conditions, the shorthand notation for each treatment, the number of subjects for each treatment and the exchange rates.[11] Overall, 12 independent computerized sessions were conducted in the RCGD lab at the University of Michigan from July to November 2003. We used zTree (Fischbacher 1999) to program our experiments. Our subjects were students from the University of Michigan.[12] No subject was used in more than one session, yielding a total of 115 subjects across all treatments. Each *i*BEA session lasted approximately two hours, while each Vickrey session lasted approximately one hour. In addition to their auction earnings, subjects could win or lose money based on their quiz answers. A subject with fully correct answers gets $5. Each mistake in the quiz costs 50 cents. The average earning (including quiz award) was $26.6 for the Vickrey auction and $47.13 for the *i*BEA auction. Data are available from the authors upon request.

4 Hypotheses

Based on the theoretical predictions and our experimental design, we identify the following hypotheses.

Hypothesis 1. *In Vickrey auctions, bidders will bid on all packages.*

Hypothesis 2. *In Vickrey auctions, bidders will bid truthfully.*

Hypotheses 1 and 2 are based on the dominant strategy of the Vickrey auction. The next two hypotheses are based on the theoretical predictions of the *i*BEA auction.

Hypothesis 3. *In iBEA auctions, bidders will bid on packages within £5 of the maximum temporary profit.*

Hypothesis 4. *In iBEA auctions, bidders will choose the last-and-final option for all packages with negative temporary profits.*

We now formulate hypotheses which compare the performance of the two mechanisms. As *i*BEA is an ascending version of Vickrey, we expect, *ex ante*, that the two mechanisms will generate the same bidder profit, auctioneer revenue and efficiency.

Hypothesis 5. *Vickrey and iBEA will generate the same amount of bidder profit.*

Hypothesis 6. *Vickrey and iBEA will generate the same amount of auctioneer revenue.*

Hypothesis 7. *Vickrey and iBEA will generate the same efficiency.*

As the environment is increasingly competitive with an increase in the number of Sincere robots, we expect that the human bidder will have the highest profit when competing against two Random Robots, and will have the lowest profit when competing against two Sincere Robots. We also expect the human bidder to learn from robot strategies in the high information treatments.

[11] At the end of iSS_ℓ, iSS_h and iSR_ℓ, the actual earnings of subjects were low; therefore, we adjust the exchange rate to $1 equal £1 for iSS_ℓ and iSS_h, and $1 equal £1.2 for iSR_ℓ. Since these adjustments happened at the end of the experiment, and no subject was used for more than one session, we do not expect them to affect the experimental results.

[12] Doctoral students in Economics are excluded from participation.

5 Results

In this section, we first examine individual bidder behavior in Vickrey and *i*BEA auctions. We then compare the aggregate performance of the two auctions.

5.1 Individual Behavior in Vickrey Auctions

Our Vickrey auction experiment consists of 60 subjects, each of whom independently plays 10 auctions with two robots in one of the six treatments. In each auction, a subject can bid on any of the 15 packages at any price between £0 and £1000.

Unlike the single item case, the strategy in a multi-item Vickrey auction with package bidding has two dimensions. The first dimension is whether to bid on a package. The second is how much to bid on a package if one decides to bid on it. A bidder's strategy on either dimension affects his profit, as illustrated in a series of simulations in Figure 1.

[Figure 1 about here.]

Figure 1 presents simulated profits for the human bidder under the three different environments in the Vickrey auction. The top panel (a) presents results for the environment with two sincere robots. The middle panel (b) presents results for the environment with one sincere and one random robot. The bottom panel (c) presents results for the environment with two random robots. The horizontal plane consists of the two strategy dimensions, the probability of bidding on a package, and the bid/value ratio for a package. In the simulation, we generate 10,000 hypothetical auctions, each of which consists of independent draws of the values for items A, B, C and D from the uniform distribution on $\{0, 1, 2, \cdots, 10\}$. In our hypothetical auctions, the preferences for the human and robot bidders, as well as the auction rules, are identical to the experimental environment. For each combination of the Probability of Bidding (drawn from the interval $[0, 1]$ with a grid size of 0.02) and Bid/Value Ratio (drawn from the interval $[0, 2]$ with a grid size of 0.04), we compute the average profit for the human bidder across these 10,000 auctions, and report it on the vertical axis. Comparing the profit from each strategy combination across environments, we find that, Sincere-Sincere is the most competitive environment, with the lowest profit level for the human bidder occurring at any given combination of the probability of bidding and the bid/value ratio. The Sincere-Random environment is the next most competitive, followed by the Random-Random environment. To analyze the tradeoffs of the two dimensions within each environment, we use the contour set for each environment.

[Figure 2 about here.]

Figure 2 presents the contour sets for the three environments specified in Figure 1. The horizontal axis is the Bid/Value Ratio, while the vertical axis is the Probability of Bidding. Each curve represents combinations of the bid/value ratio and the probability of bidding which yield the same profit. The highest profit is achieved in each case when the probability of bidding is 1, and the bid/value ratio is also 1, i.e., bidding on every package and bidding one's true value for each package, which is the dominant strategy. Each curve is similar to an indifference curve. The "inner" (or "upper") curves represents high profit levels.

[Table 3 about here.]

Table 3 presents results for the human bidders' bidding decisions, averaged across all treatments. The Active Bids column presents the ratio of positive and zero bids over 600, which is the total number of bids if every bidder bids on every package. Note a zero bid is an active bid, while no bid is inactive. The Bid/Value Ratio is the mean ratio of the bid to the value of a package. The next three columns present the proportion of Truthful Bidding, Overbidding and Underbidding among active bids, respectively.

Even though bidding on every package is a weakly dominant strategy, the results in the second column of Table 3 show that the proportion of bids never reaches 100% on any package. Furthermore, participants bid on packages containing items AB more frequently than they do on other packages. We now investigate which factors induce a higher proportion of active bids. We use a probit model with robust clustering at the individual level. The dependent variable is Active Bids, a dummy variable, which equals one if a bidder places an active bid on a package and zero otherwise. The independent variables include Value of a package, a dummy variable D_{AB}, which equals one if a package contains both items A and B, and zero otherwise, and a dummy variable D_{HS}, which equals one if the information condition is High and there is at least one sincere robot bidder, and zero otherwise. The reason for including the latter is that participants might learn from the sincere robot strategy. Even though we are explicit in the instructions for all High Information treatments that robots bid on all packages, participants might ignore both parts of the random robot strategy, as the second part is obviously not optimal. Indeed, replacing the last dummy with a dummy D_H, which equals one with High Information and zero otherwise, results in an insignificant coefficient.

Result 1 (Whether to Bid on a Package). Bidders are significantly more likely to bid on packages with higher values. In addition, they are more likely to bid on packages with synergistic items. The proportion of active bids increases significantly in treatments with at least one sincere robots and high information.

[Table 4 about here.]

Support. Table 4 presents results from the two probit specifications described above. The coefficients are probability derivatives. The Value of a package increases the likelihood of bidding on this package by 0.8%. If a package contains the synergistic bundle AB, the likelihood that a subject bids on this package is increased by 8.5%. Compared to other information conditions, subjects increase the likelihood of bidding on packages by 16.4% with at least one sincere robot and high information. All coefficients are significant at the one- or five-percent level. ∎

Result 1 indicates that bidders are significantly more likely to bid on high value packages and those with synergistic bundles. Furthermore, bidders imitate the Sincere Robot, but not the Random Robot.

We now explore the second dimension of bidding strategy, how much to bid on a package. To investigate how much participants bid on a package in a Vickrey auction, we use a structural approach based on Hypothesis 2, which proposes that bidding one's true valuation is a weakly dominant strategy. To test this hypothesis, we use an OLS regression with clustering at the individual level. In the first specification, we use Bid as the dependent variable, and Value as the only independent variable. We do not include a constant because of the theoretical prediction. In the second specification, we add Auction, Cumulative Profit, D_{AB} and D_{HS} as independent variables. The variable Auction denotes the number of auctions in a session. Thus, this variable captures any learning effect. For each specification, we run two-sided Wald tests of the null hypothesis of bids being equal to values against the alternative hypothesis of bids not being equal to values. The results are presented in Table 5.

[Table 5 about here.]

The results in Table 5 indicate that, on average, bidders tend to underbid, rather than overbid, in Vickrey auctions. In specification (1), the coefficient for Value is 0.962, which is close to truthful preference revelation. In a similar regression conducted by Isaac and James (2000), the coefficient for Value is 0.95 and not statistically different from one.[13] To classify bidders, we repeat the first specification in Table 5 for each bidder. We then perform the Wald test for the null hypothesis that the coefficient on Value is 1 and subsequently classify bidders into the following groups.

1. Underbidder: If we can reject the hypothesis of truthful bidding at the 5% level and the coefficient is below 1.

2. Truthful Bidder: If we cannot reject the hypothesis of truthful bidding at the 5% level.

3. Overbidder: If we can reject the hypothesis at the 5% level and the coefficient is above 1.

We now summarize the analysis of bidding behavior in Vickrey in the following result.

Result 2 (Bid Price in Vickrey). Bidders in a Vickrey auction, on average, bid 96.2% of their true value. Of our participants, 57% can be classified as underbidders, 32% as truthful bidders and 12% as overbidders.

Support. Table 5 presents the OLS regression results for our Vickrey auctions. The coefficient estimates show the amount subjects bid compared to their valuations. Robust standard errors in parentheses are clustered at the individual level. A two-sided Wald test of the null hypothesis of bids being equal to values yields p-value of 0.193. The classification of bidders comes from regressions at the individual level. The average R^2 of individual regressions is 0.934, with a standard deviation of 0.131. ∎

Most previous laboratory studies of single-unit Vickrey auctions find that bidders tend to overbid in such environments (Kagel 1995). In multi-unit uniform price auctions, bidders tend to overbid on the first unit and underbid on the second unit, which is consistent with

[13] Isaac and James (2000) also include a constant, which is estimated to be -0.19 and not statistically different from zero.

the theoretical prediction of demand reduction (Kagel and Levin 2001). Our finding that most bidders either underbid or bid their true value in Vickrey auctions is in stark contrast with previous experimental results. Our study presents empirical evidence that the "robust" finding of overbidding in single-unit Vickrey auctions does not carry over to package Vickrey auctions.

In addition to studying overall bidding behavior, we examine how learning may impact bidding behavior. To do so, we first define *bidding pattern* as a combination of two different measures, the Number of Active Bids[14] in an auction and the Bid/Value Ratio. To identify the effects of prior experience while minimizing individual-specific characteristics on bidding behavior, we use a difference-in-difference approach. First, for each auction, we take the difference of the Number of Active Bids between the current and the previous auctions. We then classify all observations into two groups, a winner group where subject(s) won a package in the previous auction, and a loser group, where they did not. Finally, we compare the difference between the two groups. We analyze the Bid/Value Ratio in a similar way.

Result 3 (Effect of Prior Experience on Bidding). Losers in a previous auction are significantly more likely to change their number of active bids, compared to winners. Furthermore, losers increase their bid/value ratio, while winners decrease their bid/value ratio. The difference is significant at the one-percent level.

Support. A loser in a previous auction changes his number of active bids by 2.12, on average (with a standard error of 0.221), while a winner in the previous auction changes his number of active bids by only 1.22 (with a standard error of 0.116). The difference is statistically significant at the level of 1%. A two-sample t-test with equal variances yields a p-value of 0.0001.
Furthermore, a loser in a previous auction increases his bid/value ratio by 0.102 on average (with a standard error of 0.031), while a winner in the previous auction decreases his bid/value ratio by 0.047 (with a standard error of 0.039). The difference is statistically significant at the 1% level. A two-sample t-test with equal variances yields a p-value of 0.0084. ∎

Result 3 indicates that participants in our study learn from prior experience. The directions of change for winners and losers are intuitive, yet indicate that the dominant strategy is not transparent to a substantial number of participants, who adjust their behavior by trial and error.

To understand the individual learning dynamics in our Vickrey auctions, we also investigate individual learning through two learning models, the reinforcement learning model (Erev and Roth 1998) and the payoff assessment learning model (Sarin and Vahid 1999). Both models have been shown to track human learning behavior fairly well in a variety of games, such as relatively simple games (Erev and Roth 1998), games with complete information (Chen and Gazzale 2004) and of limited information (Chen and Khoroshilov 2003). However, we find that neither tracks learning dynamics well in the more complex Vickrey auctions. The results are available from the authors upon request.

Overall, the most surprising finding in our Vickrey auctions is that most bidders either underbid or bid their true value. Even though the dominant strategy is not transparent,

[14] This is equivalent to the likelihood of bidding when the total number of packages is fixed.

our subjects tend to adjust their behavior by trial and error. These findings in individual behavior will translate into interesting aggregate performance measures, which we analyze in Section 5.3.

5.2 Individual Behavior in *i*BEA

Our *i*BEA experiment consists of 55 subjects, each of whom independently plays 10 auctions with two robots in one of the six treatments. At any round of an *i*BEA auction, each participant strategy has two components: which package(s) to bid on, and when to check the last-and-final option.

To analyze our *i*BEA auction results, we define *adjusted temporary profit* as the temporary profit of a package after a subject considers whether to check the last-and-final option. For packages with the last-and-final option checked, the temporary profit is increased by £5. The myopic best response (MBR) strategy in an *i*BEA auction states that bidders should bid on packages within ε of the maximum adjusted temporary profit. Among packages a participant can actively bid on,[15] we group packages by their adjusted temporary profits in each round of the auction. Those packages within £5 of the maximum adjusted temporary profit are called the *MBR packages.* Sincere Robots bid only on MBR packages. The rest of the packages are called *non-MBR packages.*

[Figure 3 about here.]

Figure 3 presents the proportion of human bidders that bid on MBR and non-MBR packages among all packages a participant can actively bid on, in each of the six different treatments. Comparing the low and high information treatments within SS, SR and RR, the proportion of bids (MBR and non-MBR) increases when more information is provided. However, this increase seems to occur mostly with MBR packages in SS and SR, while in RR, the proportion of bids on both MBR and non-MBR packages increases. This increase is consistent with the hypothesis that humans learn from robot strategies. We now use probit models to formally check our impression from Figure 3.

[Table 6 about here.]

Table 6 presents six probit specifications which examine factors affecting the likelihood of MBR bids. In all specifications, the dependent variable is PlaceBid, a dummy variable which equals one if a bid is placed on a package and zero otherwise. The independent variables are Temporary Profit, D_{ab}, D_{info} in specifications (1), (3) and (5). In specifications (2), (4) and (6), we add an independent variable, D_{info}*MBR, which equals one if the information condition is high and the package belongs to the set of MBR packages, and zero otherwise. This variable controls for the information condition on MBR package bidding. Note that D_{ab} and D_{info}*MBR are positively correlated. We summarize the results in the following discussion.

[15] Recall that last-and-final packages and winning packages from previous rounds are automatically checked in the current round; therefore, a participant cannot act on them. Consequently, we do not consider them as part of the choice set for the subject.

Result 4 (Bidding Decision in *i*BEA). In our *i*BEA auctions, we find that bidders are significantly more likely to bid on packages with a higher temporary profit. We also find that additional information on robot strategies
(1) induces significantly more bids on both MBR and non-MBR packages under RR, but not under SS or SR;
(2) induces significantly more bids on MBR packages under SS and SR, but not under RR.

Support. In Table 6, the coefficients on Temporary Profit are positive and significant in all six specifications. The coefficient on D_{info} is positive and significant in specification (5), but not in (1) - (4) or (6). The coefficient of D_{info}*MBR is positive and significant in (2) and (4), but not in (6). ∎

Result 4 confirms our intuition from Figure 3 that knowledge of robot strategies significantly impacts human bidding behavior. While the Sincere Robot induces more MBR bids, the Random Robot tends to increase bids on both types of packages. This suggests that, in a complex auction, participants might not be able to figure out the optimal bidding strategy, and thus are susceptible to learning. Therefore, in the actual implementation of an *i*BEA auction, teaching bidding strategies may make bidding more effective.

We now explore specifically whether subjects learn the robot strategy regarding the last-and-final option. In all *i*BEA treatments, all robots follow the equilibrium strategy in checking the last-and-final option, i.e., they check this option for a package whose temporary profit is negative. In our high information treatments, subjects are explicitly told when robots check the last-and-final option. We are interested in two questions. First, do subjects learn to use this option? Second, when they use this option, do they use it at the right time?

[Table 7 about here.]

Table 7 presents summary statistics on the mean temporary profit of all last-and-final bids in a treatment, the number of such bids where the temporary profits are negative, or non-negative, the total number of last-and-final bids and the proportion of negative last-and-final bids. The average temporary profit is £1.96 for all treatments, £0.77 for high information treatments and £4.40 for low information treatments. The difference between the high and low information treatments is highly significant overall, as well as within each environment (p-value < 0.01 for two sample t-tests with equal variance in all four cases). This indicates that subjects tend to check this option earlier than would be optimal. However, they learn from the robot strategies in the high information treatments. Another interesting result is that, among the three high information treatments, the temporary profit of the last-and-final bids in the RR environment is significantly higher than that in the SS and SR treatments (p-value < 0.01 for two sample t-tests with equal variance). This could be due to two reasons. First, because the environment RR is less competitive, the auction lasts fewer rounds than the SR or SS environments. Therefore, the SR and SS environments give bidders more time to learn. Second, the first dimension of a random robot's strategy (which packages to bid on) is obviously not optimal, which might have spillover effects on the second dimension, i.e., when to check the last-and-final option. To evaluate these causes, we use the following specifications.

[Table 8 about here.]

Table 8 reports the results of two OLS specifications, each of which controls for clustering at the individual level. In both specifications, the dependent variable is the Temporary Profit of the last-and-final bids. In specification (1), the independent variables are Round (the round when the last-and-final option is checked), D_{info} (a dummy variable which equals one under high information treatments, and zero otherwise), and a constant. As we expect the temporary profit to be lower the longer an auction lasts, we use Round to control this effect, while we study the effects of information conditions. The results in Table 8 show that the coefficients of both Rounds and D_{info} are negative and significant. In specification (2), we add another independent variable, D_{HS}, a dummy which equals one if the information condition is high and there is at least one Sincere robot, and zero otherwise. When this variable is added, the coefficient for D_{info} is no longer significant at the 5% level. This indicates that information on sincere robot strategies, rather than on random robot strategies, leads bidders to learn optimal bidding strategies.

Overall, in our *i*BEA auctions, consistent with the MBR strategy, bidders are more likely to bid on packages with higher temporary profits. Furthermore, bidders seem to imitate bidding strategies from the Sincere robots, and not the Random robots.

5.3 Aggregate Performance of the Two Mechanisms

We now examine the aggregate performance of the two mechanisms in three aspects: bidder profit, auctioneer revenue and efficiency.

In examining bidder profit, we look at both human bidder profit and aggregate human and robot profit in each treatment. Auctioneer revenue in this analysis follows the standard definition.

In single-unit auctions, efficiency is often measured by a ratio of the number of auctions where the object goes to the bidder with the highest valuation to the total number of auctions. In the multi-unit context, this definition is not applicable. Instead, we use the definition of efficiency developed by Kagel and Levin (2001). In this definition, **efficiency** of an auction is the ratio of the total surplus of the allocation to the highest possible surplus among all possible allocations, where total surplus is the sum of bidder profit and auctioneer revenue. In our environment, we have 256 possible allocations in total, as there are four items and three bidders ($256 = 4^4$). Note that it is not sufficient to consider $3^4 = 81$ allocations, since some allocations which leave item(s) unsold may yield the highest revenue to the auctioneer. For example, suppose that Bidder 1 bids on A at $1 and Bidder 2 bids on B at $1, and that Bidder 3 bids for C at $2 and CD at $1. In this case, the auctioneer does not sell D to anyone. Therefore, the problem is equivalent to allocating four items among four agents (three bidders and an auctioneer), yielding 256 possible allocations.

[Figure 4 about here]

Figure 4 presents the distribution of actual observed efficiencies in all auctions, pooling across all treatments. The left panel is the efficiency distribution under the *i*BEA auctions, and the right panel is the efficiency distribution under the Vickrey auctions. From Figure 4, it is clear that the Vickrey auction, on average, generates higher efficiency.

[Table 9 about here.]

Table 9 presents summary statistics on the average human profit, total bidder profit, auctioneer revenue and efficiency under each auction mechanism. Recall that, in the experiment, the price increment in the *i*BEA auction is £5 due to the length of the auction, while in the Vickrey auction we use a grid size of £1. We take seriously the possibility that the performance of *i*BEA might be disadvantaged due to this difference in grid size. Therefore, under each category, we present two columns for the Vickrey auction, one representing the observed results and a second representing the hypothetical results. In this case, observed efficiency (or bidder profit or revenue) is the actual experimental observation using a grid size of £1, while the hypothetical measure uses a grid size of £5. In computing the latter, we round the actual bids up or down to the closest integer on a grid of size 5. When we adjust the Vickrey auction grid size, *i*BEA and Vickrey auctions are more comparable in these measures. Of course we do not rule out the possibility that participant behavior might be psychologically affected by grid size. Thus, we compare the *i*BEA auction with both the hypothetical and actual Vickrey auction results. This comparison shows that the *i*BEA auction yields better human and total profit, while the Vickrey auction achieves higher Revenue and Efficiency. Since session average does not control for a number of exogenous factors, we use the following specifications to model factors affecting the performance of each mechanism.

[Table 10 about here.]

Table 10 presents four OLS specifications. The dependent variable in each specification is (1) Human Profit, (2) Total Profit, (3) Revenue and (4) Efficiency. The independent variables are Mechanism, which equals 1 for the *i*BEA auction and 2 for the Vickrey auction, D_{info}, the number of random robots, quiz score and a constant. Note that we use a hypothetical grid size of 5 in computing each of these measures under the Vickrey auction, so that it is comparable to the *i*BEA auction. All results hold if we use the actual Vickrey grid size of 1. We summarize our findings below.

Result 5 (Bidder Profit). Human bidder profit, as well as total profit, are significantly higher in the *i*BEA auction than in the Vickrey auction. Furthermore, the number of random robots significantly increases human profit, and weakly increases total profit. A higher quiz score significantly increases human profit.

Support. In Table 10, for specification (1), the coefficient for Mechanism is negative and significant. The coefficients for the number of random robots and quiz score are both positive and significant. In specification (2), the coefficient for Mechanism is negative and significant. The coefficient for the number of random robots is weakly significant at the 10% level. ∎

By Result 5, we reject Hypothesis 5. As an ascending version of the Vickrey auction, the iBEA auction generates significantly high profits for the human bidder. We find it interesting that a high quiz score at the end of the instruction significantly increases bidder profit, indicating that those who have a better understanding of the rules of the auction end up doing better. Note that part of this result may reflect a stronger ability to imitate sincere robot strategies as well as better understanding of the rules.

Result 6 (Revenue). Vickrey auctions generate significantly higher revenue than *i*BEA auction. In addition, the number of random robots significantly decreases revenue in both auctions.

Support. In Table 10, for specification (3), the coefficient for Mechanism is positive and significant. The coefficient for the number of random robots is negative and significant. ∎

By Result 6, we reject Hypothesis 6. Unlike the theoretical prediction of equal revenue performance, sealed-bid Vickrey auctions generate significantly higher revenue than do ascending *i*BEA auctions. This result is particularly interesting in light of the concern in the literature that Vickrey auctions might generate low revenue. Our finding of superior Vickrey auction results could be due to the fact that each human bidder competes against two robots in our experimental setting, which makes it impossible to collude. This points to a fruitful area of further research, which is to implement Vickrey package auctions among human bidders to check if this result still holds.

In our study, we measure efficiency through bidder profit and seller revenue. Using this measure, we find the following results.

Result 7 (Efficiency). Vickrey auction generates significantly higher efficiency than does the *i*BEA auction. Furthermore, the number of random robots significantly decreases efficiency, while a higher quiz score significantly improves efficiency.

Support. In Table 10, for specification (4), the coefficient for Mechanism is positive and significant at the one percent level. The coefficient for the number of random robots is negative and significant, while that for the quiz score is positive and significant. ∎

By Result 7, we reject Hypothesis 7. Empirical evidence from past laboratory studies of Vickrey and its strategically equivalent ascending bid auctions shows that, in the single unit case and the multi-unit homogeneous object case, the ascending auction usually achieves higher efficiency (Kagel 1995). This difference in performance is usually attributed to feedback in the ascending bid auction which makes bidding strategies more transparent (Kagel, Kinross and Levin 2003). Our results show that, with package bidding, the Vickrey auction generates significantly higher efficiency than does the ascending *i*BEA auction.

6 Conclusion

Package auctions have become increasingly popular in procurement and complex resource allocation contexts and thus have stimulated a large body of theoretical research in combinatorial auctions in economics and computer science. Because of the concerns over the deficiencies of Vickrey auctions, several new ascending package auctions have been proposed to implement efficient allocations under various assumptions. One of the most prominent new ascending package auctions is the *i*BEA auction, which achieves approximately efficient allocation and implements Vickrey payments under minimal assumptions on preferences.

As a first step in using this auction as an actual economic process that solves naturally occurring problems, we observe the performance of the *i*BEA auction in the context of simple situations that can be created in a laboratory. We then assess its performance relative to a natural and important benchmark, the sealed-bid Vickrey auction.

As the first experimental study of the iBEA auction in comparison with the Vickrey auction, we use a simple environment where each human bidder competes against two robots with different levels of bidding intelligence. This implementation creates an environment free from the strategic uncertainties inherent in interactions between human bidders.

Our experiment yields several surprising findings. First, unlike the single-unit Vickrey auction where bidders tend to overbid in the laboratory, most of our bidders either underbid or bid their true value. A simple dynamic adjustment model captures the learning dynamics reasonably well. Second, in terms of aggregate performance, while bidder profit is significantly higher in the iBEA auction, the Vickrey auction generates significantly higher revenue and efficiency than does the iBEA auction. This result is particularly interesting in light of the general concern in the literature (Milgrom 2004) that package Vickrey auctions might generate low revenue. Admittedly, our result on revenue might be a consequence of our experimental setting where a human bidder competes with two robots, making it impossible to collude. Nonetheless, it has important implications for the increasingly popular use of automated agents. It also points to a natural next step of the research, which is to let human bidders compete against other humans to check if they collude and thus lower revenue. Lastly, we find that when human bidders compete against robots in a complex environment, they learn from their robot opponents when the robot strategies are intelligent (e.g., myopic best responses).

With this first laboratory study of the iBEA auction, we identify several issues that warrant further empirical study. The first issue is the tradeoff between the speed and efficiency of the auction as manipulated through setting different price increments. The larger the auctioneer sets the price increment, the faster the auction converges. However, the final allocation might be further away from the efficient allocation. In an auction with a large number of bidders, we expect the second phase of the iBEA auction to take longer. Thus, it is important to quantify the speed-efficiency tradeoff. The second issue is whether bidders can detect the second phase of the auction and thus collude to increase bidder profit. As in Vickrey auctions, the use of robots in our experiment makes second phase collusion impossible. Therefore, one should interpret our results as empirical findings in the absence of collusion.

Empirical investigations of package auctions have important implications not only in procurement and privatization, but also potentially in the allocation of scarce equipment time in large scientific collaboratories. Past studies, such as Kwasnica et al. (2005), show that, with synergies, package auctions tend to outperform auctions without package bidding. The natural next step is to evaluate the pool of package auction mechanisms under a variety of environments in the laboratory, select those which perform robustly well, and test them in a field setting. This study contributes to the laboratory evaluation of two important package auction mechanisms, the Vickrey and iBEA auctions. In light of past findings, the results are surprising. Thus, they point to fruitful areas of future research.

		bidder 1			bidder 2			bidder 3		
Package		A	B	AB	A	B	AB	A	B	AB
Value		10	0	10	0	30	30	4	14	22
Phase	Round	(Bid, Price)			(Bid, Price)			(Bid, Price)		
I	1	**(A, 0)**		(AB, 0)		**(B, 0)**	(AB, 0)			(AB, 0)
I	2	(A, 0)		(AB, 0)		(B, 0)	(AB, 0)		(B, 0)	**(AB, 5)**
I	3	**(A, 5)**		(AB, 5)		**(B, 5)**	(AB, 5)		(B, 0)	(AB, 5)
I	4	**(A, 5)**		(AB, 5)		**(B, 5)**	(AB, 5)		(B, 5)	(AB,10)
I	5	**(A, 5)**		(AB, 5)		(B, 5)	(AB, 5)	(A, 0)	**(B,10)**	(AB,15)
I	6	**(A, 5)**		(AB, 5)		(B,10)	(AB,10)	(A, 0)	**(B,10)**	(AB,15)
I	7	**(A, 5)**		(AB, 5)		**(B,15)**	(AB,15)	(A, 0)	(B,10)	(AB,15)
I	8	**(A, 5)**		(AB, 5)		**(B,15)**	(AB,15)	*(A, 0)*	*(B,10)*	(AB,20)
I	9	**(A, 5)**		(AB, 5)		**(B,15)**	(AB,15)	*(A, 0)*	*(B,10)*	*(AB,20)*
II		(A, 5)		(AB, 5)		~~**(B,15)**~~	~~(AB,15)~~	*(A, 0)*	*(B,10)*	***(AB,20)***
II	10	**(A,10)**		(AB,10)		~~**(B,15)**~~	~~(AB,15)~~	*(A, 0)*	***(B,10)***	*(AB,20)*
II		~~**(A,10)**~~		~~(AB,10)~~		(B,15)	(AB,15)	*(A, 0)*	*(B,10)*	***(AB,20)***
II	11	~~**(A,10)**~~		~~(AB,10)~~		**(B,20)**	(AB,20)	***(A, 0)***	*(B,10)*	*(AB,20)*

Notes:
1. Boldface indicates the winning bid after the current round.
2. Italics indicate bids with the last-and-final option.
3. ~~(Bid,Price)~~ indicates that the (Bid, Price) pair is excluded.

iBEA	Package	Price	Piv. Revenue	Rebate	Final Price	Profit
bidder 1	A	10	20	10	0	10
bidder 2	B	20	20	10	10	20
bidder 3	-	-	-	-		0
auctioneer	-	-	-	-		10
Vickrey	Package	Price	Piv. Revenue	Rebate	Final Price	Profit
bidder 1	A	10	34	6	4	6
bidder 2	B	30	24	16	14	16
bidder 3	-	-	-	-		0
auctioneer	-	-	-	-		18

Note: Piv. Revenue refers to total revenue when bidder i is excluded.

Table 1: A Simple Example of an iBEA Auction Process

Mechanism	Robot Strategies	Information	Notation	# of Subjects	Exchange Rates
		High	iSS_h	10	2
*i*BEA	Sincere, Sincere	Low	iSS_ℓ	10	2
		High	iSR_h	9	1.5
*i*BEA	Sincere, Random	Low	iSR_ℓ	8	1.5
		High	iRR_h	8	1
*i*BEA	Random, Random	Low	iRR_ℓ	10	1
		High	vSS_h	10	1.25
Vickrey	Sincere, Sincere	Low	vSS_ℓ	10	1.25
		High	vSR_h	10	1.25
Vickrey	Sincere, Random	Low	vSR_ℓ	10	1.25
		High	vRR_h	10	1.25
Vickrey	Random, Random	Low	vRR_ℓ	10	1.25

Table 2: Features of Experimental Sessions

Package	Active Bids	Bid/Value	Underbidding	Truthful Bidding	Overbidding
A	0.700	0.989	0.538	0.271	0.190
B	0.663	0.995	0.503	0.307	0.191
C	0.668	1.184	0.544	0.284	0.172
D	0.653	0.875	0.531	0.296	0.173
AB	0.860	1.029	0.572	0.217	0.211
AC	0.708	0.964	0.607	0.224	0.169
AD	0.702	0.908	0.620	0.219	0.162
BC	0.713	0.959	0.605	0.206	0.189
BD	0.697	0.948	0.610	0.208	0.182
CD	0.730	0.917	0.584	0.249	0.167
ABC	0.835	0.978	0.605	0.202	0.194
ABD	0.802	0.953	0.613	0.200	0.187
ACD	0.743	0.954	0.596	0.231	0.173
BCD	0.722	0.962	0.580	0.238	0.182
ABCD	0.827	1.005	0.597	0.198	0.206

Table 3: Proportion of Active Bids, Bid/Value Ratio and Proportion of Under-, Truthful, and Overbidding in the Vickrey Auction

Dependent Variable: Active Bids		
	(1)	(2)
Value	0.008	0.008
	(0.002)***	(0.003)***
D_{AB}	0.085	0.086
	(0.021)***	(0.022)***
D_{HS}	0.164	
	(0.072)**	
D_{info}		-0.017
		(0.077)
Observations	9000	9000

Notes:

1. Coefficients are probability derivatives.
2. Robust standard errors in parentheses are adjusted for clustering at the individual level.
3. Significant at: ** 5% level; *** 1% level.

Table 4: Probit: Factors Affecting the Likelihood of Active Bids in the Vickrey Auction

	Dependent Variable: Bid	
	(1)	(2)
Value	0.962	0.937
	(0.029)***	(0.028)***
Auction		0.045
		(0.115)
D_{HS}		-0.874
		(0.634)
Cum. Profit		-0.053
		(0.042)
D_{AB}		0.526
		(0.327)
Constant		0.875
		(0.789)
Observations	6614	6614
R-squared	0.81	0.52

Notes:

1. Robust standard errors in parentheses are adjusted for clustering at the individual level.
2. Significant at: *** 1% level.

Table 5: OLS: Bidding Decision in the Vickrey Auction

	Dependent Variable: PlaceBid					
	(1)	(2)	(3)	(4)	(5)	(6)
	SS	SS	SR	SR	RR	RR
Temp. Profit	0.013	0.010	0.016	0.014	0.013	0.013
	(0.002)***	(0.002)***	(0.002)***	(0.002)***	(0.002)***	(0.002)***
D_{AB}	0.045	0.020	0.031	0.020	0.050	0.046
	(0.015)***	(0.013)	(0.015)**	(0.014)	(0.015)***	(0.016)***
D_{info}	0.078	-0.025	0.095	0.019	0.074	0.054
	(0.057)	(0.048)	(0.053)*	(0.058)	(0.038)**	(0.050)
D_{info}*MBR		0.271		0.144		0.041
		(0.096)***		(0.027)***		(0.034)
Observations	23643	23643	16506	16506	13680	13680

Notes:

1. Coefficients are probability derivatives.
2. Robust standard errors in parentheses are adjusted for clustering at the individual level.
3. Significant at: * 10% level; ** 5% level; *** 1% level.

Table 6: Probit: Likelihood of MBR Bidding in the iBEA Auction

Treatment	Mean Temp Profit	Negative	Non-Negative	Total	Correct Rate
SS_ℓ	3.97	38	79	117	32.48%
SR_ℓ	3.86	5	80	85	5.88%
RR_ℓ	10.16	0	19	19	0.00%
SS_h	1.36	128	108	236	54.24%
SR_h	-0.25	97	107	204	47.55%
RR_h	4.56	2	16	18	11.11%

Table 7: Last-and-Final Summary Statistics for the iBEA Auction

Dependent Variable: Temp. Profit of Last-and-Final Bids		
	(1)	(2)
Round	-0.390	-0.399
	(0.101)***	(0.102)***
D_{info}	-3.786	-3.061
	(1.367)***	(1.586)*
D_{HS}		-1.501
		(1.072)
Constant	6.943	7.001
	(1.525)***	(1.525)***
Observations	679	679
R-squared	0.18	0.19

Notes:

1. Robust standard errors in parentheses are adjusted for clustering at the individual level.
2. * Significant at: * 10% level; ** 5% level; *** 1% level

Table 8: OLS: Temporary Profit of Last-and-Final Bids

	Human Profit			Total Profit		
Treatment	*i*BEA	Vickrey	Vickrey (grid 5)	*i*BEA	Vickrey	Vickrey (grid 5)
SS_ℓ	3.04	2.72	2.23	16.29	10.85	8.75
SR_ℓ	3.56	1.80	2.46	18.15	8.56	8.70
RR_ℓ	7.02	3.30	2.96	18.27	9.58	9.33
SS_h	2.71	2.48	2.14	16.84	10.06	9.18
SR_h	3.75	2.96	2.89	17.77	9.75	8.99
RR_h	7.05	2.73	2.67	17.85	10.10	9.53
	Revenue			Efficiency		
robots	*i*BEA	Vickrey	Vickrey (grid 5)	*i*BEA	Vickrey	Vickrey (grid 5)
SS_ℓ	13.87	20.12	20.90	92.2%	98.4%	94.02%
SR_ℓ	8.86	20.04	19.70	84.8%	90.3%	89.60%
RR_ℓ	7.13	18.88	18.75	78.3%	89.7%	88.28%
SS_h	13.05	20.44	20.45	93.0%	97.4%	94.49%
SR_h	9.60	19.92	20.20	87.3%	93.0%	91.40%
RR_h	5.82	17.81	17.80	73.2%	88.0%	86.34%

Table 9: Aggregate Performance of the Two Mechanisms

	(1)	(2)	(3)	(4)
Dependent Variable:	Human Profit	Total Profit	Revenue	Efficiency
Mechanism	-2.422	-8.389	9.386	0.043
	(0.460)***	(0.503)***	(0.487)***	(0.010)***
D_{info}	0.002	0.291	-0.351	0.003
	(0.351)	(0.431)	(0.406)	(0.008)
# of Random Robots	1.150	0.505	-2.354	-0.060
	(0.212)***	(0.258)*	(0.277)***	(0.005)***
Quiz Score	0.816	-0.223	0.850	0.023
	(0.378)**	(0.459)	(0.436)*	(0.010)**
Constant	2.246	26.225	-0.746	0.764
	(1.352)*	(1.822)***	(1.634)	(0.041)***
Observations	1150	1150	1150	1150
R-squared	0.09	0.28	0.42	0.19

Notes:
1. Robust standard errors in parentheses are adjusted for clustering at the individual level.
2. Significant at: * 10%; ** 5%; *** 1% level.

Table 10: OLS: Factors Affecting Aggregate Performance of the Two Mechanisms

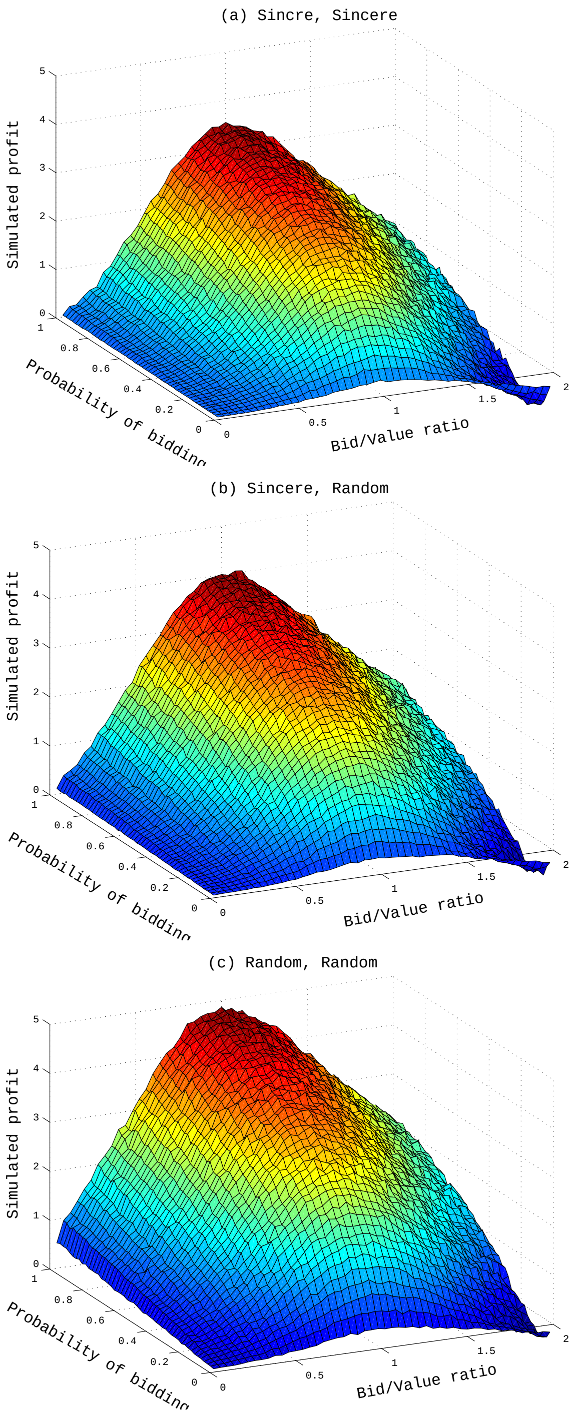

Figure 1: Simulated Profit for Human Bidder under Three Environments in the Vickrey Auction

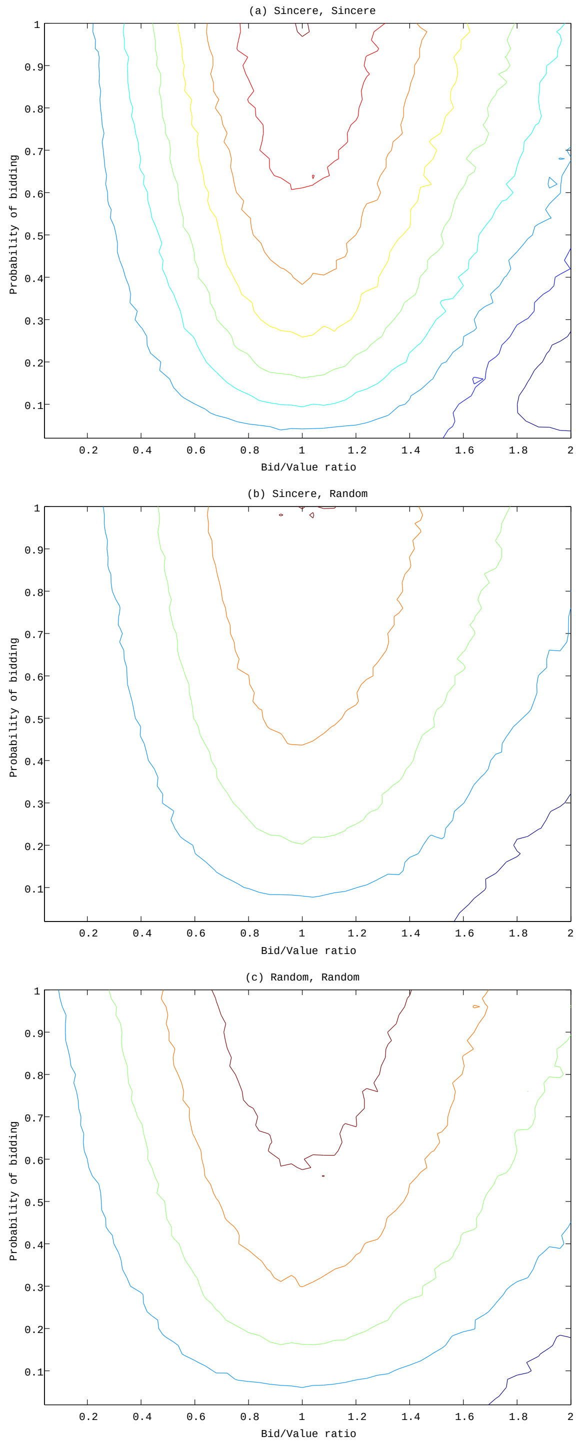

Figure 2: Contour of Simulated Profit for Human Bidders under the SS Environment in the Vickrey Auction

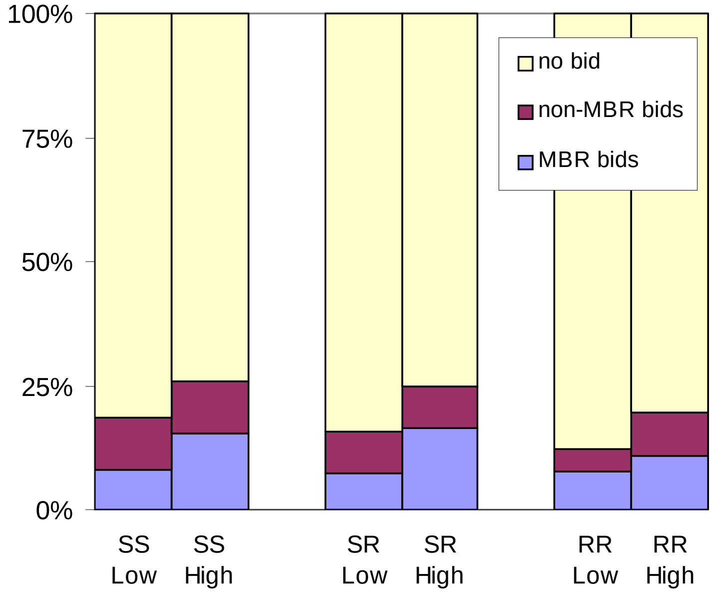

Figure 3: Proportion of MBR and non-MBR Bids in the *i*BEA Auction

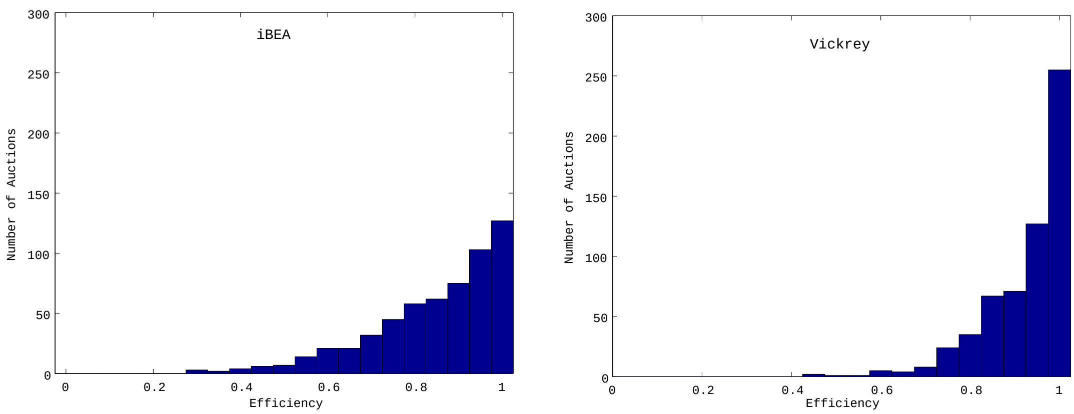

Figure 4: Distribution of Observed Efficiency in the *i*BEA and Vickrey Auctions

References

Ausubel, Lawrence M. and Paul R. Milgrom, "Ascending Auctions with Package Bidding," *Frontiers of Theoretical Economics*, 2002, *1* (1), 1–43. available at http://www.bepress.com/bejte/frontiers/vol1/iss1/art1.

Banks, Jeffrey S., John O. Ledyard, and David P. Porter, "Allocating Uncertain and Unresponsive Resources: An Experimental Approach," *Rand Journal of Economics*, 1989, *20*, 1–25.

Bykowsky, Mark M., Robert J. Cull, and John O. Ledyard, "Mutually Destructive Bidding: The FCC Auction Problem," *Journal of Regulatory Economics*, 2000, *17* (3), 205–228.

Cantillon, Estelle and Martin Pesendorfer, "Auctioning Bus Routes: The London Experience," in Peter Cramton, Yoav Shoham, and Richard Steinberg, eds., *Combinatorial Auctions*, Cambridge: MIT Press, 2006.

Chen, Yan and Robert S. Gazzale, "When Does Learning in Games Generates Convergence to Nash Equilibrium? The Role of Supermodularity in an Experimental Setting," *American Economic Review*, December 2004, *94* (5), 1505–1535.

__ **and Yuri Khoroshilov**, "Learning under Limited Information," *Games and Economic Behavior*, July 2003, *44*, 1–25.

Clarke, Edward H., "Multipart Pricing of Public Goods," *Public Choice*, 1971, *11*, 17–33.

Eisenberg, Anne, "In Online Auctions of the Future, It'll Be Bot vs. Bot vs. Bot," *New York Times*, August 2000.

Erev, Ido and Alvin E. Roth, "Predicting How People Play Games: Reinforcement Learning in Experimental Games with Unique, Mixed Strategy Equilibria," *American Economic Review*, September 1998, *88* (4), 848–881.

Finholt, Thomas A., "Collaboratories," in B. Cronin, ed., *Annual Review of Information Science and Technology*, Vol. 36 American Society for Information Science Washington, D.C. 2002, pp. 74–107.

Fischbacher, Urs, "z-Tree: A Toolbox for Readymade Economic Experiments," 1999. University of Zurich Working Paper No. 21.

Groves, Theodore, "Incentives in Teams," *Econometrica*, 1973, *41* (4), 617–31.

Isaac, R. Mark and Duncan James, "Robustness of the incentive compatible combinatorial auction," *Experimental Economics*, June 2000, *3* (1), 31–53.

Kagel, John H., "Auctions: A Survey of Experimental Research," in J. Kagel and A. Roth, eds., *Handbook of Experimental Economics*, Princeton, New Jersey: Princeton University Press, 1995, pp. 131–148.

__ **and Dan Levin**, "Behavior in Multi-Unit Demand Auctions: Experiments with Uniform Price and Dynamic Vickrey Auctions," *Econometrica*, March 2001, *69*, 413–454.

__ **, Scott Kinross, and Dan Levin**, "Implementing Efficient Multi-Object Auction Institutions: An Experimental Study of the Performance of Boundedly Rational Agents," *Ohio State University Manuscript*, 2003.

Katok, Elena and Alvin E. Roth, "Auctions of Homogeneous Goods with Increasing Returns: Experimental Comparison of Alternative "Dutch Auctions," *Management*

Science, August 2004, *50* (8), 1044–1063.

Kwasnica, Anthony M., John O. Ledyard, David Porter, and Christine DeMartini, "A new and improved design for multi-object iterative auctions," *Management Science*, March 2005, *51* (3), 419–434.

Ledyard, John O., David P. Porter, and Antonio Rangel, "Experiments Testing Multi-Object Allocation Mechanisms," *Journal of Economics and Management Strategy*, 1997, *6*, 639–675.

__ , Mark Olson, David P. Porter, Joseph A. Swanson, and David Torma, "The First Use of a Combined-Value Auction for Transportation Services," *Interfaces*, 2002, *32* (5), 4–12.

Lucking-Reiley, David, "Auctions on the Internet: What's Being Auctioned, and How?," *Journal of Industrial Economics*, 2000, *48*, 227–252.

__ , "Vickrey Auctions in Practice: From Nineteenth-Century Philately to Twenty-First-Century E-Commerce," *Journal of Economic Perspectives*, 2000, *14* (3), 183–192.

Milgrom, Paul R., *Putting Auction Theory to Work*, New York: Cambridge University Press, 2004.

Ockenfels, Axel and Alvin E. Roth, "Last-Minute Bidding and the Rules for Ending Second-Price Auctions: Evidence from eBay and Amazon Auctions on the Internet," *American Economic Review*, September 2002, *92* (4), 1093–1103.

Parkes, David C. and Lyle H. Ungar, "An Ascending-Price Generalized Vickrey Auction," June 2002. Presented at Stanford Institute for Theoretical Economics (SITE) Summer Workshop: The Economics of the Internet, `http://www.eecs.harvard.edu/~parkes/pubs/iBEA.pdf`.

Rassenti, Stephen J., Vernon L. Smith, and Robert L. Bulfin, "A combinatorial auction mechanism for airport time slot allocation," *Bell Journal of Economics*, 1982, *XIII*, 402–417.

Rothkopf, Michael H., Aleksandar Pekec, and Ronald M. Harstad, "Computationally Manageable Combinational Auctions," *Management Science*, August 1998, *44* (8), 1131–1147.

Sarin, Rajiv and Farshid Vahid, "Payoff Assessments without Probabilities: A Simple Dynamic Model of Choice," *Games and Economic Behavior*, August 1999, *28* (2), 294–309.

Vickrey, William S., "Counterspeculation, Auctions and Competitive Sealed Tenders," *Journal of Finance*, 1961, *16*, 8–37.

Weisstein, Eric W., *CRC Concise Encyclopedia of Mathematics*, London, UK: CRC Press, 2002.

Wulf, Wm A., "The collaboratory opportunity," *Science*, 1993, *261* (5123), 854–855.