

Facilitating Large-scale Collaboration Through Effective Mechanism Design: An Economic Examination of Resource Allocation in Collaboratories

John C. Lin, Yan Chen,
Thomas A. Finholt
School of Information,
University of Michigan
1075 Beal Avenue
Ann Arbor, MI 48109-2112
{johnclin,yanchen,finholt@umich.edu}

Kan Takeuchi
Department of Economics,
University of Michigan
238 Lorch Hall
Ann Arbor, MI 48109-1220
ktakeuch@umich.edu

ABSTRACT

In many disciplines success depends on access to scarce resources, such as unique instruments. Research on computer-supported cooperative work (CSCW) has contributed to the development of technologies, such as collaboratories, that broaden access to scarce scientific and engineering resources. While broader access is often applauded, little attention has been focused on the problem of equitable and efficient resource allocation in the face of increased demand created through collaboratory use. This paper, then, applies concepts from the economic discipline of mechanism design to compare different resource allocation schemes (RAD, Vickery, and knapsack) within a hypothetical collaboratory. Experimental results show that knapsack achieves a more equitable distribution of resources than RAD or Vickery, but that RAD and Vickery are both more efficient than knapsack. The findings highlight the need for systematic exploration of allocation mechanisms within collaboratories, where simple optimization (e.g., knapsack) is likely to produce a sub-optimal match of resources to needs. More generally, the findings illustrate the utility of economic approaches in understanding issues that emerge in large-scale collaborations, such as entire scientific and engineering communities, that have not typically been the subject of CSCW research.

Categories and Subject Descriptors

J.4 [Computer Applications, Social and Behavioral Sciences]: economics; H.5.3 [Information Systems Information Interfaces and Presentation]: Group and organization interfaces—*computer-supported cooperative work*,

*(Produces the permission block, and copyright information). For use with SIG-ALTERNATE.CLS. Supported by ACM.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

Copyright 200X ACM X-XXXXX-XX-X/XX/XX ...\$5.00.

collaborative computing; K.6.2 [Computing Milieux, Management of Computing and Information Systems]: [pricing and resource allocation]

General Terms

Economics, Experimentation, Performance

Keywords

package auction, mechanism design, cyberinfrastructure, collaboratory, resource allocation

1. INTRODUCTION

The National Science Foundation (NSF) and other federal agencies are making significant investments in expanding the ability of geographically distributed groups of scientists to conduct research via the Internet (e.g., [1]). One means to enhance the work of dispersed teams is the “collaboratory.” 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 [19]. The collaboratory idea is a descendent of earlier notions, such as Vannevar Bush’s 1945 proposal for the “memex;” Douglas Engelbart’s 1968 demonstration of graphical user-interfaces for computer-supported meetings; and the creation in 1969 of the ARPAnet. There are currently several hundred collaboratories in use within multiple scientific communities, including space physics, medicine, software engineering, and neuroscience [6].

The George E. Brown, Jr. Network for Earthquake Engineering Simulation (NEES) is a ten-year NSF program focused on accelerating and improving earthquake engineering research through the use of a collaboratory (it.nees.org). The NEES collaboratory connects earthquake engineering researchers to sixteen state-of-the-art laboratory facilities distributed around the United States. These facilities were built with the explicit intent of combining their capabilities through the Internet. For example, users of the NEES collaboratory are able to collect data across instrument modalities spanning key earthquake engineering sub-disciplines. In this approach a test of a bridge pier system might include performance data on: a) the soil the pier sits on (e.g., from

physical geotechnical simulations run using specimens in a centrifuge); b) the pier structure (e.g., from physical structural simulations run using specimens on a shake table); and c) the bridge deck (e.g., from numerical simulations). Through NEESit the results of collaboratory studies will be accessible through shared data repositories, remote participants will be able to observe experiments, and under highly controlled circumstances remote participants will be able to control experiments in real-time.

A critical element of the NEES vision, and of collaboratories generally, is that Internet-mediated methods for conducting and observing research should diminish the advantages of physical collocation. Specifically, proximity to unique instruments has frequently created differential access, both to the instruments and to the community of scientists who use the instruments ([8]; [16]). In NEES, a critical goal of making the sixteen new sites accessible over the Internet is to broaden use of the facilities, particularly among researchers at institutions that lack earthquake engineering research equipment. To ensure that broader use occurs, NSF stipulated the creation of a non-profit corporation, NEES Inc. (nees.org), as an additional element of the NEES program intended to operate the NEES collaboratory over the period 2004 through 2014.

NEES Inc. has responsibility for governing the use of NEES facilities, as well as for allocating maintenance and operation funds. A central concern of NEES operations is to ensure shared use of the sixteen new laboratories. Specifically, a special committee of the NEES Inc. board of directors has oversight of shared use policies and mediates disputes arising from conflicts in shared use. Unfortunately, beyond the requirement that NEES accomodates shared use, NSF has not offered specific guidelines about how to achieve this goal. At an abstract level, the problem confronting NEES Inc. is similar to a class of problems in economics described as mechanism design. That is, NEES Inc. controls instrument time designated for shared use at the sixteen NEES laboratories (i.e., use by researchers in the broader earthquake engineering community, not just those collocated with the NEES equipment). This instrument time, presumably, has value for researchers, such that demand for instrument time will exceed supply. For example, in astronomy there is intense competition for telescope time ([5]; [14]).

Understanding how to achieve efficient and equitable allocation of scarce resources is a key focus of research on mechanism design. For instance, allocation solutions in other domains (e.g., the sale of frequency spectrum by the Federal Communications Commission) involve various forms of auctions. There are two reasons why it is important to explore similar approaches for resource allocation in collaboratories. First, in the absence of mechanisms it seems likely that collaboratory users will default to allocation through a scheduling czar or by insiders (e.g., instrument owners). This approach runs the risk that facilities will be under-utilized and that allocation decisions may be subject to bias (e.g., political influences). Second, absent any specific allocation mechanisms, collaboratory users may arrive at naive pricing schemes for resources, such as flat unit costs versus costs that reflect actual demand (and hence value). For example, the unit cost assumption breaks down when a user places a premium on allocation of contiguous operating intervals (i.e., ten hours together versus ten one hour segments dispersed over several days).

Collaboratory operators, such as NEES Inc., will be the immediate beneficiaries of an effort to design and demonstrate mechanisms for allocation of instrument time within collaboratories. However, more generally, research into the problem of efficient and equitable resource allocation within large-scale collaborations is important. For instance, if the NEES collaboratory is a precursor to the organization of infrastructure across a wide variety of scientific and engineering communities then the problems confronted by NEES are likely to emerge in these other settings. That is, if sharing of instruments is to become the norm, then there must be a corresponding set of mechanisms to support this sharing. It seems unlikely that simple altruism, particularly in highly competitive scientific and engineering fields, will yield satisfactory solutions. Instead, the time is right to conduct research that will produce a toolbox of candidate mechanisms, with proven characteristics, so that when communities confront breakdowns or bottlenecks in the use of instruments, or any shared resource, these difficulties can be overcome quickly or avoided completely.

The rest of the paper is divided into four sections. In the first of these sections we identify three allocation mechanisms and describe their features. In the next section we describe our experimental design including our assumptions about the hypothetical collaboratory. The next section presents our results and in the final section we discuss the results.

2. ALLOCATION MECHANISMS

In this section we outline the technical details of three allocation mechanisms. We assume that a critical feature of the instrument time allocation problem is that contiguous time slots are more valuable than the sum of separate slots. For example, in the case of the NEES collaboratory the difficulty of experiment set-up and teardown dictates a preference for consecutive time intervals to minimize installation effort. Therefore, package auctions might be an important mechanism in achieving efficient allocation of equipment time. We consider two package auction mechanisms, Vickrey and RAD, compared with an ordinal ranking mechanism, knapsack, selected as a best-case representation of how allocation is currently accomplished.

2.1 The Vickery Auction

The Vickrey auction ([17], [4] and [7]) is an important standard for nearly all mechanism design work, and for auctions in particular. It 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.

A Vickrey auction with package bidding is an extension of the more familiar second-price auction. At the beginning of each auction, bidders select the packages they would like to bid on, and the amount they would like to bid for each package.

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.

Previous laboratory experiments indicate that, in environments with a small number of packages, the Vickrey auction can achieve high allocation efficiency in multi-object auctions with package bidding [9], and that it outperforms a complex ascending bid auction in terms of efficiency and revenue [3]. It is an open question whether it can retain its excellent performance in more complex environments. This study addresses this issue by expanding the number of packages to a more realistic level.

2.2 Resource Allocation Design (RAD)

The RAD mechanism is an iterative ascending bid package auction proposed by Kwasnica et al. [12]. The RAD mechanism borrowed features from the AUSM mechanism [2] and the Milgrom FCC Design [13]. The features that RAD borrowed from the Milgrom FCC Design are: the auction is iterative instead of one shot; the use of an eligibility rule; and a price improvement rule. The feature that RAD borrowed from AUSM is the use of a package bidding language. RAD, in addition, calculates the prices on the individual months that underlie the packages. RAD is shown to outperform both parents [12].

The details of the RAD mechanism are explained in [12]. Here, we will briefly explain the different features on a global level. An iterative auction is one where the bidding activities are separated into different rounds. Within a round, bidders can place bids. At the end of a round, bidders are given feedback information that can be used before bids are placed in the next round. In the case of RAD, bidders are told at the end of a round if the bids they placed were provisionally winning or losing. At the end of the last round, there is no provisional winning bid. A winning bid at the end of the last round is the binding winning bid.

The effect of the eligibility rule can be briefly summarized as a “use it or lose it rule”. Roughly, the more packages a bidder bids on, the higher is the eligibility number for that round. In the subsequent round, the number of packages a bidder can bid on depends on the eligibility number from the previous round. The higher the eligibility number, the more packages a bidder can bid on.

The price improvement rule specifies a minimum price for each package based on the bids that were submitted in the previous round and a price improvement factor. The price improvement rule, in conjunction with the eligibility rule, helped to drive an auction to a close.

The difference between RAD and the FCC Design is RAD allows the placement of package bids. In the hypothetical collaborative context, this means the ability to say a bidder wants a three month package starting in the first month, or a two month package starting in the second month, but not both. Note that the bidder who placed the above two bids

can only win at most one of the two placed packages, not both. Furthermore, if a bidder wins a package, the bidder is guaranteed to win the complete package, not part of a package. With the FCC design, bidders can only place bids on single months. So even though bidders can place bids that are equivalent to a two month package starting in the second month, it is possible for the bidder to win only part of the package, namely, the second or the third month, instead of the package containing both months. Thus bidders can be exposed to the risk that they may overpay. This is known as the “exposure problem” in the economic literature.

Lastly, the novel aspect of RAD is its ability to calculate prices for the underlying months that the packages are constructed out of.

Single unit iterative ascending auctions, such as the English auction, tend to outperform their sealed bid counterpart in terms of efficiency. Therefore, in the package bidding context, we expect that RAD will generate higher efficiency than Vickrey.

2.3 Knapsack with Ordinal Ranking

In many scientific communities, allocation of instrument time can be determined by a scheduling czar, a committee [14] or some formal or informal optimization procedure which uses the ordinal ranking information from potential instrument users. The knapsack mechanism is an idealized representation of the latter. In this mechanism, everyone submits ordinal rankings of packages, from the top choice (# 1) to the last choice (# n for n packages). The top choice is awarded n points, second choice $n - 1$ points, ..., and the last choice 1 point. and the computer will allocate goods in such a way that maximizes the total number of points.

This class of knapsack mechanisms with ordinal ranking has two problems. First, truth-telling might not be a dominant strategy. One can easily construct examples to show that truthful ranking based on one’s valuations for the packages might not be optimal. Second, bidders cannot express the intensity of their preferences. Therefore, we expect that the knapsack mechanism will generate lower efficiency than Vickrey and RAD. Because of the prevalence of knapsack-like mechanisms in scheduling, we use it as a useful benchmark of typical allocation approaches in order to evaluate the two auction mechanisms.

3. METHOD: EXPERIMENTAL DESIGN

Our experimental design reflects both theoretical and technical considerations. Specifically, we are interested in three important questions. First, how do the Vickrey, RAD and Knapsack mechanisms compare in performance? Second, how do subjects respond to the incentives in each mechanism? In this section we describe our hypothetical collaborative and experimental procedures.

3.1 The hypothetical collaborative

The hypothetical collaborative was designed to capture the essential aspects of a scientific or engineering collaborative, using the NEES collaborative as a model. In some areas, simplifications were made to keep the problem tractable in the experimental setting, where subjects were only available for a limited amount of time in each session. Participants were told that they represented research groups seeking instrument time via a collaborative within a 24-month

window. We defined participants in terms of their resource level (i.e., small vs. big projects) and in terms of their time preference (i.e., early, indifferent, or late in the 24-month window). Each participant was asked to consider sixty-three (for a big project) or sixty-six (for a small project) packages of instrument time. The packages were defined by the starting month and the duration of the use. For example, a package that started on the first month and lasted for three months contained month 1, month 2, and month 3. For simplicity, packages were denoted in the form of an ordered pair (starting month, length of package) or (1,3) for the example above.

A small project bidder had values for packages that lasted for 2 to 4 months. A big project bidder had values for packages that lasted for 3 to 5 months. As a result, a small project bidder was able to bid on packages that started on month 1 to month 23 (i.e., the first number in the ordered pair is a number between 1 to 23), with lengths of 2, 3, or 4 (i.e., the second number in the ordered pair is a number between 2 to 4). Big project bidders were also able to bid on packages starting from month 1 to month 22 (i.e., the first number in the ordered pair is a number between 1 to 22), but with lengths of 3, 4, or 5 (i.e., the second number in the ordered pair is a number between 3 to 5). Note a big project bidder was bidding on sixty-three packages, while a small project bidder was bidding on sixty-six packages.

The value of a package was determined by the type and the time preference factor of a bidder, as well as a v parameter that specified the value of an additional month for the participant. As a first step, we determined the value of a base package, (1, 2) for a small project bidder and (1, 3) for a big project bidder. For a small project bidder the value of the base package was drawn from the uniform distribution [20, 100]. For a big project bidder the value of the base package was drawn from the uniform distribution [20, 150].

Once we determined base package values, we determined the value of packages that start in the first month. A v parameter was drawn from the uniform distribution of [10, 20] for all the participants. For example, for a small project bidder, where the value of the base package (1, 2) was 50 and the v drawn for that particular player was 15, the values for the packages (1, 3) and (1, 4) were 65 and 80, respectively. We determined the values for the other packages from the packages that start in the first month. As an illustration, the value of the (2, 2) package was the value of the (1, 2) package times the time preference factor. So if a bidder had the time preference factor of 1.2, the (2, 2) package was more valuable than the (1, 2) package (in general, this means the participant prefers packages that start later rather than earlier). If the participant had the time preference factor of 1, the (2, 2) package was worth the same as the (1, 2) package (in general, this means the participant was indifferent between the starting time of the package). If the participant had the time preference factor of 0.9, the (2, 2) package was less valuable than the (1, 2) package (in general, this means the participant prefers packages that start earlier rather than later).

In summary, the values of the packages were determined by the value of the base package, the time preference factor of the bidder, and the value of an additional month, beta. The value of the base package in turn was determined by the bidder's type.

We settled on nine bidders in each experimental session

Bidder ID	Project Type	Time Preference	Package Size
1	Big	Prefer later	2 to 4 months
2	Big	Indifferent	2 to 4 months
3	Big	Prefer earlier	2 to 4 months
4	Small	Prefer later	3 to 5 months
5	Small	Indifferent	3 to 5 months
6	Small	Indifferent	3 to 5 months
7	Small	Prefer earlier	3 to 5 months
8	Small	Prefer earlier	3 to 5 months
9	Small	Prefer earlier	3 to 5 months

Table 1: Design Parameters: Bidder Preferences

to mimic the number of projects funded to use the NEES collaboratory in the first year of operation (2004-05). Of these nine there was a distribution of big and small projects, which we operationalized as three big projects and six small projects. Of the three big project bidders all three time preference factors were represented among the bidders. Of the six small project bidders, one had the time preference of 1.2, two had the time preference of 1, and three had the time preference of 0.9. Earlier months are more competitive than later ones, to capture the stylized fact in most scientific communities that earlier experiment, and hence, potential discoveries, are more valuable to scientists. Table 1 summarizes the main feature of bidder preferences.

3.2 Experimental Procedures

Each experimental session required exactly nine participants. The participants were recruited from the University of Michigan, including both the graduate and undergraduate population. The participants had declared majors in the fields of science, math, or engineering. Once the participants arrived they reviewed and signed an informed consent form. At the beginning of each session, each participant was given printed instructions. Before the instructions were read aloud by one of the experimenters, the participants drew from a deck of index cards for a player ID that determined their project type (big or small) and their time preference factor (1.2, 1, or 0.9). The instructions were then read aloud. Participants were encouraged to ask questions during and after the experiment. The instruction period averaged around 30 minutes. Then participants took a quiz designed to test their understanding of the mechanism. At the end of the quiz, the experimenters collected, graded, and returned the quiz to each of the participants. The experimenters then reviewed the answers with the participants. The participants were given 10 minutes to do the quiz for the RAD treatment, and 7 minutes to do the quiz for the Vickrey treatment and the knapsack treatment.

There were no practice auctions in any of the experimental conditions. In the Vickrey and knapsack condition, participants had seven minutes per auction to input their bids, while in the RAD condition, the participants had four minutes for the first round and two minutes for all subsequent rounds in each auction. There were a total of eight auctions in each of the Vickrey and knapsack sessions, and three to five auctions in each of the RAD sessions.¹

At the conclusion of auctions in each condition partici-

¹Recall that RAD is an iterative auction, therefore, the number of rounds in each auction is endogenous and varies from auction to auction.

Mechanism	# sessions	# participants	Exchange Rate
Vickrey	5	45	12
RAD	4	36	4
Knapsack	4	36	20

Table 2: Features of Experimental Sessions

pants tallied their cumulative earnings, filled out a short demographic survey, and wrote down the strategies that they used. We use the induced value method [15]. Participants were paid based on their experimental profits.

Table 2 presents the relevant features of the experimental sessions, including mechanisms, number of experimental sessions, the number of participants in each condition and the exchange rates. Overall, 13 independent computerized sessions were conducted in the RCGD lab at the University of Michigan from July 2005 to February 2006. No participant was used in more than one session, yielding a total of 117 participants across all treatments. Each session lasted approximately two and a half hours. In addition to their auction earnings, participants could win or lose money based on their quiz answers. A participant with fully correct answers gets up to \$5. The average earning (including quiz award and a \$10 showup fee) was \$33.64, and the standard deviation was \$16.39. Data are available from the authors upon request.

4. RESULTS

In this section, we first examine individual bidder behavior in each allocation mechanism. We then compare the aggregate performance of the two auctions in terms of efficiency and equity.

	Vickrey		Knapsack
	Active Bids%	Bid/Value	Active Bids%
Auction 1	0.290	0.593	0.503
Auction 2	0.419	0.626	0.588
Auction 3	0.415	0.651	0.575
Auction 4	0.433	0.724	0.596
Auction 5	0.446	0.729	0.594
Auction 6	0.479	0.737	0.577
Auction 7	0.526	0.740	0.574
Auction 8	0.548	1.126	0.588

Table 3: Bidding Under Vickrey and Knapsack

We first examine two aspects of the bidding behavior in Vickrey, the proportion of active bids and the bid/value ratio for those active bids. Recall it is a weakly dominant strategy for each bidder to bid on all packages, and to bid their true value on each package.

The first two columns of Table 3 present the proportion of active bids and the bid/value ratio under the Vickrey auction. Bidders in a Vickrey auction, on average, bid 73.00% of their true value. Overall, both measures are smaller than the theory predicts. However, both the proportion of active bids and the Bid/Value ratio are increasing, indicating that bidders were learning the weak dominant strategy as they repeated the auctions.

The third column of Table 3 presents the proportion of active bids under the knapsack mechanism. Unlike the Vickrey

mechanism, the proportion of active bids remained in the (0.5, 0.6) range during the eight auctions. Not bidding on every package can sometimes be an optimal strategy under knapsack.

We also classified Vickrey bidders into three categories: Under bidder, Truthful bidder and Over bidder. Specifically, we ran the following simple OLS regression on active bids for each bidder with robust clustering at the auction level.

$$\text{Bid}(x_i) = \beta \text{Value}(x_i) + \epsilon_i. \quad (1)$$

Then, we tested the null hypothesis: $\hat{\beta} = 1$. Based on the result, we classified each bidder into one of the following.

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.

Analysis shows that bidders in a Vickrey auction, on average, bid 73.0% of their true value. Of the 45 participants, 68.9% were classified as underbidders, 22.2% as truthful bidders and 8.9% as overbidders. This is surprising, as most previous laboratory studies of single-unit Vickrey auctions find that bidders tend to overbid in such environments [10]. In multi-unit uniform price auctions, bidders tend to overbid on the first unit and underbid on the second unit, which is consistent with the theoretical prediction of demand reduction [11]. Our finding that most bidders either underbid or bid their true value in Vickrey auctions is consistent with [3].

Figure 1 presents the scatter plot of raw bids under Vickrey for small project bidders. Big project bidder bids followed similar patterns. We notice earlier months are much more competitive, and therefore, most of the bids which were over value (above the 45 degree line) were from researchers who preferred earlier months.

Individual behavior under RAD was straightforward. Bids were mostly between the minimal price and value for a package.

These individual behavioral patterns had direct implications on the aggregate performance of the mechanisms.

We used two indices to compare the aggregate performance of the three mechanisms: efficiency and the Gini coefficient. These indices were calculated for each auction outcome, which was characterized by the allocation of months and the value of those months as a package.

To derive the efficiency for an auction outcome, we needed to find the *optimal allocation* of packages. At the optimal allocation, the sum of the value of allocated package that each participant receives is maximized. Then the efficiency is the ratio of realized total bidder value and maximum potential total bidder value. It is formally defined as follows,

$$\text{Efficiency} = \frac{\sum_i v_i(x_i)}{\sum_i v_i(x_i^*)} \quad (2)$$

where x_i is the package that is actually allocated to participant i in the outcome and x_i^* is the package that would be allocated in the optimal allocation. That is, efficiency is 100% if the auction results in the optimal allocation.




Figure 1: Raw Bids in the Vickrey Auction

For example, suppose we have two packages, A and B, and two bidders, John and Tom. Package A is worth \$30 to John and \$10 to Tom, while package B is worth \$20 to John and \$30 to Tom. The efficient allocation is to give A to John and B to Tom, which yields a total value of \$60 and 100% efficiency. If, for some reason, John gets B and Tom gets A, the total value is only \$30, and the efficiency is $30/60=50\%$. Efficiency requires that the package goes to the bidder who values it the most.

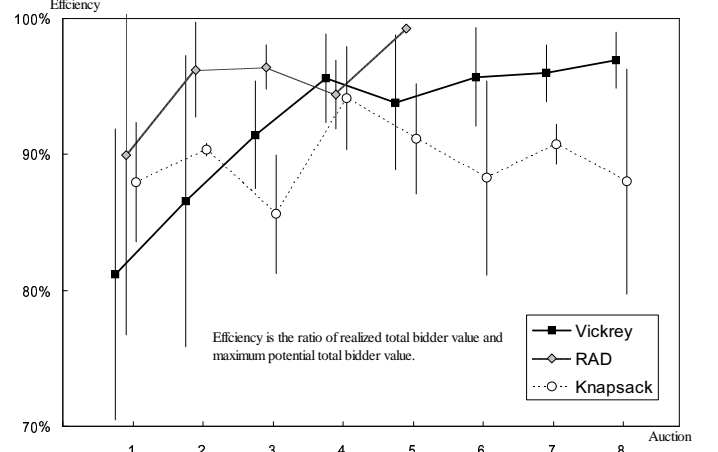


Figure 2: Efficiency comparison across mechanisms

Figure 2 presents the average time series efficiency and standard deviation (error bars) of the three mechanisms. Consistent with individual learning under the Vickrey auction, the average efficiency under Vickrey increased over time and was higher than knapsack. RAD also generated higher efficiency than knapsack. Note that RAD only had 3-5 auctions in each session, while both Vickrey and knapsack had eight auctions.² While RAD generates higher efficiency in the earlier auctions, Vickrey catches up by auction #4.

Statistical tests confirmed our impressions from Figure 2. We computed the session average, average of the last two auctions in each session, as well as auction by auction efficiency comparisons across mechanisms. We used the permutation tests³ to compare the performance of the mechanism, where the number of independent observations was the number of the sessions. For session average efficiency comparisons, we had the following results, where the p-value for each one-sided test is written under the inequality sign.

$$\begin{aligned} \text{Knapsack} &< \text{Vickrey} < \text{RAD}, \\ &\quad 0.21 \quad 0.11 \\ \text{Knapsack} &< \text{RAD}. \\ &\quad 0.05 \end{aligned}$$

Taking into account the effects of learning, we looked at the average efficiency of the last two auctions in each session,

²To avoid laying the error bars on top of each other, there is a slight shift for each mechanism, which has no numerical significance.

³The permutation test, also known as the Fisher randomization test, is a nonparametric version of a difference of two means t-test.

and found the following result:

$$\begin{aligned} \text{Knapsack} &\overset{0.01}{<} \text{Vickrey} \overset{0.33}{<} \text{RAD}, \\ \text{Knapsack} &\overset{0.03}{<} \text{RAD}. \end{aligned}$$

We concluded that efficiency was significantly higher under RAD than under knapsack. In the last two auctions, Vickrey and RAD each generated significantly higher efficiency than knapsack. The efficiency comparison between RAD and Vickrey was insignificant, except in earlier auctions, e.g., $\text{RAD} > \text{Vickrey}$ at a significance level of 2.39% in the third auction.

While the efficiency index measures whether time slots are allocated to the bidders who value them the most, we used the Gini coefficient to measure the distributional equality among bidders. The Gini coefficient is widely used to measure income inequality in a country and has been used by CSCW researchers to measure equality of participation in group discussions (e.g., [18]). We derived the Gini coefficient for each auction outcome based on the value of the allocated package to each participant. Note that a higher Gini Coefficient corresponds to greater inequality. In the extreme cases, it is 0 if every participant gets the same value from the assigned package (perfect equality), while it is approximately 1 if one participant receives some months while other participants receive nothing (perfect inequality).

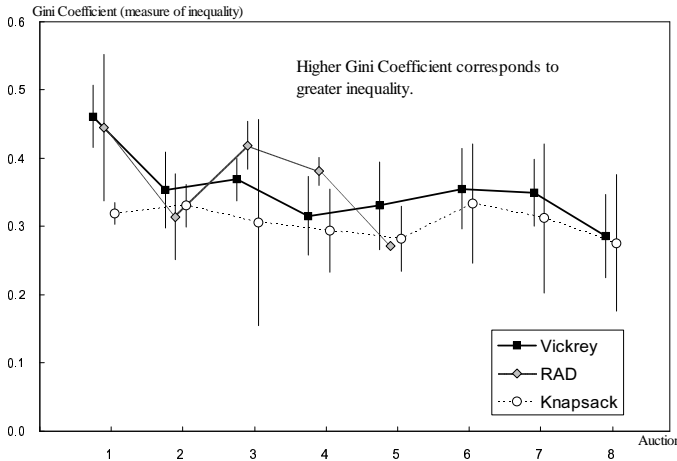


Figure 3: Equity comparison across mechanisms

Figure 3 presents the average time series Gini coefficients and the standard deviation (error bars) for each of the three mechanisms. One striking feature is that knapsack was roughly the lower envelope of the observations, indicating more equitable distribution of time slots than the other two mechanisms. Permutation tests of session average and last two auction average across mechanisms confirmed this impression.

$$\begin{aligned} \text{Knapsack} &\overset{0.04}{<} \text{Vickrey} \overset{0.06}{<} \text{RAD}, \\ \text{Knapsack} &\overset{0.03}{<} \text{RAD}. \end{aligned}$$

Taking into account the effects of learning, we looked at the average Gini coefficient of the last two auctions in each session, and found the following result:

$$\begin{aligned} \text{Knapsack} &\overset{0.34}{<} \text{Vickrey} \overset{0.04}{<} \text{RAD}, \\ \text{Knapsack} &\overset{0.06}{<} \text{RAD}. \end{aligned}$$

Therefore, we concluded that the Knapsack mechanism was significantly more equitable than the Vickrey or RAD mechanism. Vickrey, in turn, was weakly more equitable (i.e., at the 10% significance level) than RAD.

5. DISCUSSION

This paper represents an exploration of aspects of collaboration at a larger scale than is typically studied in CSCW research. Specifically, the emergence of community-level technologies such as collaboratories (i.e., serving tens to thousands of users), introduces new kinds of phenomena that are not easily captured using the techniques commonly employed in CSCW studies (e.g., ethnomethodology or social psychology). As an alternative we proposed a hypothetical collaboratory where we experimentally compared competing economic approaches for allocation of scarce resources, such as instrument time.

Our hypothetical collaboratory was designed to mimic key aspects of actual collaboratories, such as the NEES collaboratory. That is, we captured variation in the resources of collaboratory users (i.e., big vs. small projects) and in the time preferences of users (i.e., early, middle, or late in a 24-month period). Experimental participants were then assigned randomly to project types and asked to bid for instrument time using one of three mechanisms – two that were auction-based (Vickrey and RAD) and one based on ordinal ranking (knapsack).

Results showed that knapsack was more equitable than either Vickrey or RAD. However, both Vickrey and RAD were more efficient than knapsack. Intuitively, by using ordinal ranking, the knapsack mechanism does not allow bidders to express the intensity of their preferences. In the hypothetical collaboratory, this favored small project bidders and made it easier for them to obtain slots. Therefore, the knapsack allocation was more equitable, to the extent that everyone got something. Knapsack was not good at giving slots to those who valued them the most, however, where both of the auction mechanisms were better at giving the right people the right slots, at the expense of equity. Therefore, a choice among these mechanisms roughly boils down to a tradeoff between efficiency and equity, and which has more weight in the designer's objectives.

There are a number of limitations to this study, such as the usual caveats about the external validity of laboratory experiments used to simulate real world phenomena. However, there is extensive evidence that experimental participants do accurately represent economic behavior within the domain of auctions – and this is the basis for a growing sub-discipline within economics, called experimental economics (recently recognized through the award of a Nobel prize to Vernon Smith, a pioneer of experimental economic methods). Similarly, while the hypothetical collaboratory can't model all the nuance of an actual collaboratory, it is possible to stylize key aspects of collaboratory use relevant to specific problems (e.g., allocation of scarce instrument time).

and to operationize these aspects as parameters within an experimental design. Perhaps most important, it was not our purpose to use experimental methods to capture the full fidelity of collaborative use. Rather, the experimental approach gave us a sufficiently realistic arena to examine particular aspects of collaborative use with an eye toward future work that might focus on allocation in actual laboratories.

We see four critical next steps in terms of advancing beyond this paper. First, we believe the results will have greater validity when participants are members of authentic scientific and engineering communities. That is, instead of using undergraduates, we would like to use practicing scientists and engineers – in the expectation that these participants will identify more strongly with the instrument allocation task and therefore have greater investment in bidding outcomes (i.e., valuations will be more truthful and accurate). Second, we would like to conduct surveys of scientific and engineering communities that depend on scarce resources (e.g., astronomy) to better identify current mechanisms used to allocate instrument time. In this paper, for instance, we adopted knapsack as a proxy for typical allocation mechanisms based on evidence that some variant of knapsack is used to allocate time on the Chandra X-ray observatory. Assuming we find other mechanisms in use, we would like to include these as comparisons in future experiments. Finally, our larger ambition is to use the body of experimental results to inform adoption of specific allocation mechanisms within scientific and engineering communities, like NEES, and then study the consequences of the use of the allocation mechanism. Ideally we would like to find multiple communities to conduct quasi-experiments contrasting varying mechanisms. Finally, from a design perspective, an important goal of future work will be the creation of effective and usable interfaces to the various auction mechanisms. For instance, a limitation on RAD is that bidders can't quickly scan bids to maintain an intuitive sense of where overall bidding may stand.

6. ACKNOWLEDGMENTS

We thank Mark Ackerman, Tony Kwasnica, John Ledyard and Jeff MacKie-Mason for helpful discussions, Sasha Achen and Benjamin Taylor for excellent research assistance. Chen gratefully acknowledges the financial support from the National Science Foundation through grant no. SES-0339587. Any remaining errors are our own.

7. REFERENCES

- [1] D. E. Atkins, K. K. Droegemeier, S. I. Feldman, H. Garcia-Molina, M. L. Klein, D. G. Messerschmitt, P. Messina, J. P. Ostriker, and M. H. Wright. Revolutionizing science and engineering through cyberinfrastructure. In *Report of the National Science Foundation blue-ribbon advisory panel on cyberinfrastructure*, 2003.
- [2] J. S. Banks, J. O. Ledyard, and D. P. Porter. Allocating uncertain and unresponsive resources: An experimental approach. *Rand Journal of Economics*, 20:1–25, 1989.
- [3] Y. Chen and K. Takeuchi. Multi-object auctions with package bidding: An experimental comparison of ibea and vickrey. University of Michigan Manuscript, 2005.
- [4] E. H. Clarke. Multipart pricing of public goods. *Public Choice*, 11:17–33, 1971.
- [5] J. Etherton, I. A. Steele, and C. J. Mottram. A free market in telescope time? In P. J. Quinn and A. Bridger, editors, *Optimizing Scientific Return for Astronomy through Information Technology*. SPIE, 2004.
- [6] T. A. Finholt. Collaboratories. In B. Cronin, editor, *Annual Review of Information Science and Technology*, volume 36, pages 74–107, Washington, D.C., 2002. American Society for Information Science.
- [7] T. Groves. Incentives in teams. *Econometrica*, 41(4):617–31, 1973.
- [8] W. O. Hagstrom. *Scientific community*. Basic Books, New York, 1965.
- [9] R. M. Isaac and D. James. Robustness of the incentive compatible combinatorial auction. *Experimental Economics*, 3(1):31–53, June 2000.
- [10] J. H. Kagel. Auctions: A survey of experimental research. In J. Kagel and A. Roth, editors, *Handbook of Experimental Economics*, pages 131–148. Princeton University Press, Princeton, New Jersey, 1995.
- [11] J. H. Kagel and D. Levin. Behavior in multi-unit demand auctions: Experiments with uniform price and dynamic vickrey auctions. *Econometrica*, 69:413–454, March 2001.
- [12] A. M. Kwasnica, J. O. Ledyard, D. Porter, and C. DeMartini. A new and improved design for multi-object iterative auctions. *Management Science*, 51(3):419–434, March 2005.
- [13] P. R. Milgrom. *Putting Auction Theory to Work*. Cambridge University Press, New York, 2004.
- [14] M. Olson and D. Porter. An experimental examination into design of decentralized methods to solve the assignment problem with and without money. *Economic Theory*, 4:11–40, 1994.
- [15] V. L. Smith. Microeconomic systems as an experimental science. *American Economic Review*, 72:923–955, 1982.
- [16] S. Traweek. *Beamtimes and lifetimes: The world of high energy physics*. Harvard University Press, Cambridge, MA, 1992.
- [17] W. S. Vickrey. Counterspeculation, auctions and competitive sealed tenders. *Journal of Finance*, 16:8–37, 1961.
- [18] S. Weisband, S. Schneider, and T. Connolly. Computer-mediated communication and social information: status salience and status difference. *Academy of Management Journal*, 1995.
- [19] W. A. Wulf. The collaborative opportunity. *Science*, 261(5123):854–855, 1993.