**Allocating Uncertain and Unresponsive Resources: An Experimental Approach**

Jeffrey S. Banks; John O. Ledyard; David P. Porter

*The RAND Journal of Economics*, Vol. 20, No. 1 (Spring, 1989), 1-25.

Stable URL:
http://links.jstor.org/sici?sici=0741-6261%28198921%2920%3A1%3C1%3AAUAURA%3E2.0.CO%3B2-%23

*The RAND Journal of Economics* is currently published by The RAND Corporation.

---

Your use of the JSTOR archive indicates your acceptance of JSTOR's Terms and Conditions of Use, available at http://www.jstor.org/about/terms.html. JSTOR's Terms and Conditions of Use provides, in part, that unless you have obtained prior permission, you may not download an entire issue of a journal or multiple copies of articles, and you may use content in the JSTOR archive only for your personal, non-commercial use.

Please contact the publisher regarding any further use of this work. Publisher contact information may be obtained at http://www.jstor.org/journals/rand.html.

Each copy of any part of a JSTOR transmission must contain the same copyright notice that appears on the screen or printed page of such transmission.

---

JSTOR is an independent not-for-profit organization dedicated to creating and preserving a digital archive of scholarly journals. For more information regarding JSTOR, please contact support@jstor.org.

http://www.jstor.org/
Thu Feb 2 14:38:50 2006

RAND Journal of Economics
Vol. 20, No. 1, Spring 1989

# Allocating uncertain and unresponsive resources: an experimental approach

Jeffrey S. Banks*

John O. Ledyard**

and

David P. Porter**

*We identify an important class of economic problems that arise naturally in several applications: the allocation of multiple resources when there are uncertainties in demand or supply, unresponsive supplies (no inventories and fixed capacities), and significant demand indivisibilities (rigidities). Examples of such problems include: scheduling job shops, airports, or supercomputers; zero-inventory planning; and the allocation and pricing of NASA's planned Space Station. Using experimental methods, we show that the two most common organizations used to deal with this problem, markets and administrative procedures, can perform at very low efficiencies (60–65% efficiency in a seemingly robust example). Thus, there is a need to design new mechanisms that more efficiently allocate resources in these environments. We develop and analyze two mechanisms that arise naturally from auctions used to allocate single-dimensional goods. These new mechanisms involve computer-assisted coordination made possible by the existence of networked computers. Both mechanisms significantly improve on the performance of administrative and market procedures.*

## 1. Introduction

■ Short-run demand and supply imbalances are pervasive in markets and other organizations. The economic impact of these imbalances is exacerbated when supply cannot be quickly changed or inventories cannot be held and when there are indivisibilities in demand. Moreover, the problems created by indivisibilities become increasingly severe as the number of commodities increases. The combination of uncertainty, indivisibilities, and unresponsive supply creates an important economic allocation problem that is not amenable to standard methods of analysis.

One example of this problem is found in airport scheduling and the allocation of takeoff/landing slots. Weather and mechanical failures create uncertainty. Indivisibilities

* University of Rochester.

** California Institute of Technology.

This work was partially funded by Caltech and NASA-Jet Propulsion Laboratories. We thank them for their support. They are not responsible for the content. We thank Peter Gray and Mark Olson for computer programming assistance. We also thank Charles Plott, Jim Quirk, and Stan Reiter for helpful insights and discussions.

in demand occur because the arrival or departure of a single airplane requires a fixed capacity for a slot, baggage handling, a gate, parking, etc. These capacities are not quickly adjusted. The method by which resources are allocated in this situation can have a significant effect on efficiency through both the direct allocation and the decisions of users of the system as to which aircraft to fly and when.

Another example, which motivated the research we report here, is NASA's planned earth-orbiting Space Station. The Station is to be an integrated facility providing a variety of services (e.g., data management, manpower, pressurized volume) to users over time. This will be a pioneer project with many new and untested technologies.[1] The performance of the Station and the resources it will be able to supply to users will be subject to considerable uncertainty over its lifetime. On the demand side, users will design and develop payloads that will consume station resources in varying degrees of intensity. Once designed and built, there is little scope for substitution. Thus, the overall Space Station allocation problem will involve selecting users and scheduling (manifesting) discrete payload demands within the uncertain and unresponsive operating capacities of the system. The processes by which allocations are chosen will affect payload designs and the ultimate rewards from the use of the Space Station.[2]

Other examples include supercomputer scheduling, natural gas pipeline networks, electric power grids, NASA's deep space network, job-shop scheduling, and attempts to coordinate production schedules so that one is perpetually in a state of "zero inventories." Each involves uncertainty in demand or supply (usually correlated across commodities), indivisibilities in demand, unresponsive supply, and nonstorable commodities. Some of these features are stronger in particular examples than in others,[3] but all are potentially present, especially over short periods of time.

Two generic forms of economic organization have usually been brought to bear on this economic problem: markets and administrative procedures. With markets, property rights are defined, initial endowments assigned, and contingent contracts created and freely traded either through organized markets or in a more dispersed manner as in wholesale-retail relationships. With administrative procedures, an apparently more centralized method is created to solve the coordination and timing problems. Gas pipelines are regulated.[4] Airports are managed through complex committees (Grether, Isaac, and Plott, 1981). The Space Transportation System (sometimes called STS or the Space Shuttle) and its resources are allocated through a complex system of hierarchical committees and detailed administrative rules.[5] Every economist is trained to expect the inefficiencies inherent in the centrally administered approach to allocation that arise from differential information, inappropriate incentives, and the existence of veto groups. What may not be obvious to some is that, in the environments we have described, markets can suffer similar inefficiencies. The technical reason is the existence of nonconvexities (optimal allocations cannot be supported as equilibria, even with complete contingent contracts), so that market-clearing prices do not exist. The intuitive reason, obvious to most engineers, is that attempts at quick coordination across multiple dimensions through unconnected markets create instability because feedback

---

[1] See Banks, Ledyard, and Porter (1985) and Fox and Quirk (1985) for details.

[2] For a more extensive discussion of the Space Station allocation and decisionmaking problems, see Ledyard (1986).

[3] For example, in electric networks, new power plants can be fired up to yield responsive supply, and demand uncertainties are probably uncorrelated—except for weather events such as heat waves—so that aggregate demand is predictable over time. Thus, only the indivisibilities are serious (shipping power from $A$ to $B$ requires a combination of links on the grid, some of which are no help without others).

[4] Of course this regulation is primarily intended to control pipeline owners in monopoly positions, but it is also used to coordinate demand-supply imbalances.

[5] See Appendix H of the *Space Station Operations Task Force: Panel 3 Report* (1987) for a detailed description of resource allocation in NASA programs, including the STS.

through prices and (random) rationing can be misleading. It may not even be possible to be good on average. Consider the following quotation from Koopmans and Beckman (1957):

In the light of the practical and theoretical importance of indivisibilities, it may seem surprising that we possess so little in the way of successful formal analysis of production problems involving indivisible resources. However, the mathematical difficulties that arise in attempts to construct a general theory of allocation of indivisible resources have so far seemed quite formidable. Perhaps the best chance of progress lies in isolating for detailed study a few limited but well-defined problems, proceeding gradually from crude simplicity and artificiality to more realistic complexity.

This statement retains its validity today and captures the spirit of our approach. In what follows we have focused on a "limited but well-defined problem" that is crudely simple but captures the phenomena we want to analyze. We examine the performance of well-defined but simple versions of the economic organizations identified above as markets and administrative procedures and find them wanting. We then develop and design other organizations, adapted from known principles in single-dimensional problems, which use a form of computer-assisted coordination to improve significantly on that performance. We feel that one of these mechanisms strongly merits further study.

## 2. A simple environment

■ It is fairly easy to describe the general structure of the class of problems in which we are interested. It is even easier, however, to understand the practical difficulties created for economic organization, in the presence of multiple decisionmakers with differential information, by using an example. A simplistic but representative version of the economic problem arises in the allocation of STS resources. Agents from the private and public sectors design and build payloads such as commercial satellites or scientific experiments. These payloads are then integrated into a Shuttle, which is launched into low earth orbit. Once in orbit, the payloads use resources supplied by the Shuttle, such as power and manpower, to complete their missions. To keep the example simple, suppose that there are two shuttles, $A$ and $B$, to be launched on two dates, 1 and 2, where any particular launch may be delayed. Further, suppose that each Shuttle provides resources $\bar{y} \in \boldsymbol{R}_+^K$ with a successful launch. Included in $K$ are available power and manpower on orbit as well as the mass that will be lifted.

A natural next modelling step would be to assign probabilities to the states and to think of this as a random supply problem. But allocations cannot be made solely contingent on the state, because at the time $A$ is launched on date 1 it is still unknown whether $B$ will be launched on date 2. It is necessary to consider explicitly the time structure in an event tree. We show this in Figure 1, where $A$ = $A$ launches, $B$ = $B$ launches, and $N$ = launch delayed. Allocations, supplies, and demands are specified contingent on events. Thus, we model

FIGURE 1

TREE STRUCTURE OF EVENTS

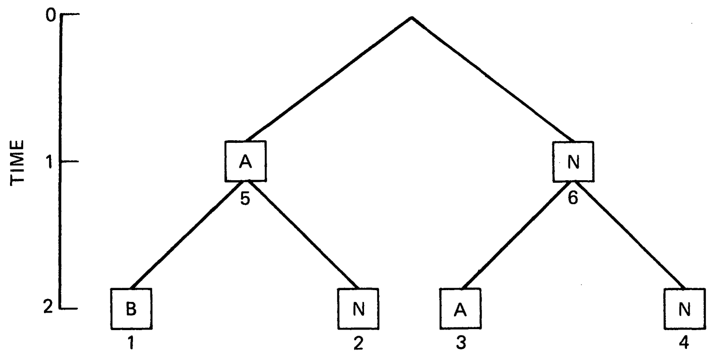

supply as a function $y : \mathbb{E} \rightarrow \boldsymbol{R}_+^K$; where $\mathbb{E} = \{1, \ldots, 6\}$, $y(1) = y(3) = y(5) = \bar{y}$, and $y(2) = y(4) = y(6) = 0$.

We assume that the probability of any launch is common knowledge and represent the probability of an event $e$ in $\mathbb{E}$ as $\eta(e)$. For example, $\eta(1)$ is the probability that $B$ launches at time 2, given that $A$ launched at time 1. Thus, $\eta(1) + \eta(2) = 1$, $\eta(3) + \eta(4) = 1$, and $\eta(5) + \eta(6) = 1$.

The demand side is almost as simple. At time 0, each agent $i = 1, \ldots, I$ picks a payload design $a^i$ from a set of possible designs $A^i$. To minimize notation we shall index designs by their resource requirements. Thus, $A^i \subseteq \boldsymbol{R}_+^K$. A (contingent) allocation to $i$ is a vector $x^i = \langle x_1^i, \ldots, x_6^i \rangle$, where $x_e^i \in \boldsymbol{R}_+^K$ is the allocation of resources to $i$ to be delivered if and only if event $e$ occurs. The utility $i$ gets from the design-allocation vector $(a^i, x^i)$ is based on the assumption that $i$ is a risk-neutral, von Neumann-Morgenstern, expected-utility maximizer whose preferences are quasi-linear in a commodity we shall call money. Let

$$U^i(a^i, x^i) = G^i(a^i)\Big[\sum_{e \in \mathbb{E}} \gamma^i(a^i, x_e^i, e)\Big],$$

where $\gamma^i(a^i, x_e^i, e) = \eta(e)$ if $x_e^i \geq a^i$, $= 0$ otherwise, and $G^i(a^i)$ are the net benefits of the payload design $a^i$ if the launch is successful. If $i$ pays $b^i$ for the contingent allocation $x^i$ with design $a^i$, $i$ attains an *ex ante* expected utility at time 0 of $U^i(a^i, x^i) - b^i$. Finally, we require that $x_1^i \cdot x_5^i = 0$ since a payload may launch only once.

To summarize, the environment is given by the event tree in Figure 1, the supply function $y(\cdot)$, the design sets $A^i$, the benefit functions $G^i(a^i)$, and the restrictions $x_1^i x_5^i = 0$ for all $i$. The uncertain and unresponsive supply is found in $y(\cdot)$ and $\eta(\cdot)$. Indivisibilities enter in two ways. First, the constraint that $x_1^i x_5^i = 0$ is locational in nature (Koopmans and Beckman, 1957) since $i$ can consume only in event 1 or 5 but not both. This creates a nonconvexity in the consumption sets. Second, once $a^i$ is chosen, there is no flexibility, and $i$ therefore demands only a single fixed amount in each event. This creates a threshold nonconvexity in utility.

To evaluate the performance of alternative institutions in this environment, we must use a performance target to identify those (contingent) allocations that are considered to be desirable. Our choice for a target is the set of *ex ante* Pareto-optimal allocations.[6] That is, we consider the maximization of social value from a perspective before any realization of supply—at time 0, the top of Figure 1. Because we assume that all agents have quasi-linear, risk-neutral utility functions and the same, correct objective beliefs about the probabilities of each event, *ex ante* Pareto-optimal allocations are equivalent to solutions to the following mixed-integer, nonlinear programming problem: choose $a$, $\beta$, and $x$ to

$$\max \sum_{i=1}^{I} G^i(a^i) \sum_e \beta^i(e)\eta(e),$$

subject to

$$\sum_{i=1}^{I} \beta^i(e)a^i \leq y(e) \qquad \forall e$$

$$\beta^i(1)\beta^i(5) = 0 \qquad \forall i$$

[6] A natural alternative to this performance standard would be the set of *ex post* Pareto-optimal allocations. This standard would require that, after observing the true state of the world, the new information would not lead society to regret any previous allocation decision; there would be no reallocation that could have made everyone better off, even with the added advantage of hindsight. Unfortunately, there are many examples of the type of environment we are considering for which there are no (contingent) allocations that are both consistent with the information structure in Figure 1 and are *ex post* Pareto optimal. It is easy to construct examples such that the *ex post* optimal allocation of $A$ differs in states 1 and 2. Thus, one cannot choose an allocation contingent on event 5 that will be *ex post* optimal *in all possible branches from 5*. Therefore, we have chosen *ex ante* Pareto optimality as the primary performance standard by which we shall judge alternative allocation mechanisms.

$$\beta^i(e) = 0 \text{ or } 1 \qquad \forall i, \forall e$$
$$x_e^i = \beta^i(e)a^i \qquad \forall i, \forall e$$
$$a^i \in A^i \qquad \forall i.$$

The allocation problem is to devise a method to find an optimal allocation when only $i$ knows $\langle A^i, G^i(\cdot)\rangle$.

If a single agent possessed all the information, he could, in principle, solve this problem and compute the optimal contingent allocations. In practice, this is a multiple knapsack problem whose complexity increases dramatically in the number of commodities and number of users. If the information is dispersed and knowledge is privately held by $i$, then this information must be communicated in some way to the others so that the optimization problem can be solved. Further, this communication must occur before the realization of $e$. Economists now realize that, at the very least, this involves an incentive problem. For the environments in which we are most interested there is also a communication problem. It is not possible in practice (with bounded communications systems) for any agent to communicate the full range of his possibilities $A^i$ or preferences $G^i$ when there are multiple contingencies and multiple dimensions.

## 3. Markets

■ The first instinct of most economists would be that, since there are no externalities and a small number of states, one should be able to create enough contingent markets to allocate resources efficiently if there is enough depth to avoid noncompetitive behavior. The analysis is simple and standard. Following methods introduced in Debreu (1959, chap. 7), one creates a market for each commodity at each node of the event tree. A price $P_k(e)$ is a real number interpreted as the amount paid initially by the agent who commits himself to accept delivery of that commodity $k$ if event $e$ occurs. If $e$ does not occur, no delivery takes place, but the payment is still made. A competitive equilibrium is a vector of prices $P^* \in \boldsymbol{R}^{KE}$, where $E$ is the number of events in $\mathbb{E}$, payload designs $a^{*i}$, and event-contingent allocations $x^{*i}(e)$, for all $e \in \mathbb{E}$, such that $\sum_{i=1}^{I} x^{*i}(e) \leq y(e)$, $\forall e \in \mathbb{E}$, and $\forall i \in N$, $a^{*i}$, $x^{*i}(1), \ldots,$ $x^{*i}(E)$ solves

$$\max \left\{\sum_e U^i[a^i, x^i(e), e]\eta(e)\right\} - \sum_e P_k^*(e)x^i(e).$$

□ **The analysis.** It is easy to see that there are no externalities and that preferences satisfy local nonsatiation (because "money" is infinitely divisible and always desired). Thus, the first welfare theorem (Debreu, 1959, chap. 5) implies that the competitive equilibrium allocation $a^*$, $x^*(\cdot)$ is Pareto optimal. One is tempted to conclude that markets solve the problem.

Unfortunately, competitive equilibria may not exist. The reason is straightforward. The environment has a unique optimal allocation. By the second welfare theorem (Debreu, 1959, chap. 6), if there is enough convexity and continuity, then there are prices such that this optimum is (supportable as) a competitive equilibrium. But there are two fundamental nonconvexities in the structure of our allocation problem that prevent the application of this theorem. In fact, for most of the environments under consideration, the unique optimum cannot be supported as an equilibrium. Therefore, by the first welfare theorem there exists no equilibrium. We can trace each of the two troublesome nonconvexities to a specific type of indivisibility.

(1) There is a locational indivisibility. A payload can only be in shuttle $A$ or $B$ but not in

both. This creates a nonconvexity in the consumption set. In a two-dimensional analogy one must be on the edges of the Edgeworth box. Interior allocations are not possible consumptions. This is exactly the type of indivisibility that Koopmans and Beckmann (1957) identified and that has created problems for urban regional economics.

(2) There is a threshold indivisibility that exists because one must commit to a design before knowing the state of the world. The inability to adapt to the realizations of the random events means that, conditional on the design $a^i$, one can only choose whether to participate by buying $a^i$ or by buying 0 at each event. This creates a nonconvexity in the preferences of $i$( a nonconcavity in the utility function).

Each nonconvexity is potentially fatal in the sense that there are environments with each alone such that competitive equilibria do not exist. Together these nonconvexities make nonexistence virtually certain. But does this mean that markets will not work? Not necessarily. For example, agents might use mixed strategies that "smooth out" the discontinuities created by the nonconvexities and yield allocations that are good on average. One might not achieve 100% efficiency, but one might come close. It is possible to test this view theoretically, but strategic issues become quickly confused with the competitive hypothesis. What is the appropriate "equilibrium price" if competitive equilibria do not exist? What is the game? Are there multiple prices for the same commodity? We can no longer ignore institutional features normally abstracted from in competitive analysis.[7] But development of the theory is not necessary for this article. We tested the performance of markets in another way and turn to that now.

□ **The experiment.** To test the effectiveness of markets we created an experimental environment for which a competitive equilibrium does not exist and asked the double oral auction to allocate resources. This is a market institution resembling organized commodities and stock markets[8] that has performed at 95–100% efficiency levels in past applications (Smith, 1982; Plott, 1982) and that has become the experimental standard market mechanism.

*The environment.* The experimental environment is based on the example described in Section 2. In particular, the design involves two resources $(X, Y)$ in fixed supply, two dates $(t = 1, t = 2)$, and two possible outcomes at each date $(g, n)$. The quantity of the goods for period 1 are available $(g)$ with probability $\rho_1$ and unavailable $(n)$ with probability $(1 - \rho_1)$. Either the total quantity is available or no quantity is available for consumption. For period 2 the probability of $g$ is $\rho_2$ and of $n$ is $(1 - \rho_2)$. $\rho_2$ is independent of the period 1 outcomes. Table 1 shows the exact parameters used to represent the supply side of the experiments.

The demand side was created by using monetary functions to induce value (Smith, 1976). For subject $i = 1, \ldots, I$, values are induced by assigning to each $a^i \in A^i$ a monetary value of $M^i(a^i)$. If a subject has a von Neumann-Morgenstern, risk-neutral utility function for money, then by letting $G^i(a^i) = M^i(a^i)$ we can identify the induced values with the model in Section 2.[9] In the experiments only discrete amounts were made available for the

[7] Even so, it is easy to construct examples of marketlike institutions and game structures in which only pure-strategy equilibria exist and in which the equilibria allocations are not optimal.

[8] The rules impose a bid-ask improvement rule and provide for the public broadcast of the standing bid-ask spread.

[9] If subjects are not risk neutral, but have von Neumann-Morgenstern utilities $g^i(m)$ for money, then if they choose design $a^i$, receive resources $x^i$, and pay $b^i$, they have an induced expected utility of

$$U^*(a^i, x^i, b^i) \equiv g^i[M^i(a^i) - b^i](\sum_e \gamma^i(a^i, x^i_e, e)) + [1 - (\sum_e \gamma^i(a^i, x^i_e, e))]g^i(-b^i),$$

**TABLE 1 Supply Parameters**

| Time Period | Quantity of X | Quantity of Y | Probability of Availability |
|---|---|---|---|
| 1 | 20 | 20 | $p_1 = 2/3$ |
| 2 | 20 | 20 | $p_2 = 1/2$ |

$a$ choices.[10] In particular, each of the six subjects was given a 3 × 3 matrix of values corresponding to nine possible choices. The actual values appear in the Appendix. The specific parameters (payoffs and project sizes) chosen for the experiments required a computer search, since the number of combinations that can fit within the available capacity limits and provide action in the market is sizeable, given six 3 × 3 matrices of choices. (A total of 90 parameters must be picked.) The selection rule for the parameters in our design was quite subjective.

We summarize a Monte Carlo study of efficiency levels[11] in Figure 2. Notice that the lowest possible efficiency level is in the 20–30% interval and that 85% of the distribution mass is between 40 and 75%. Hence, if we were to allocate resources randomly, we should not be surprised to see efficiency levels in the range of 40–70%. Very few allocations, however, yield efficiency levels above 80%. The allocation described in Table 2 is the unique expected-value-maximizing use of resources and design choices. Although the numbers used for the experiments are contrived, they do provide a "hard" test for any mechanism designed to coordinate demands to allocate resources efficiently.

*The institution.* To test the efficiency of markets in the context of the nonconvexities we have created, we gave subjects the option of trading in any of six markets corresponding to the appropriate contingent commodities in nonzero supply. We identify the markets in Table 3. We divided the capacity available in each market among each of the six agents as an initial endowment that they could keep or sell. We chose the specific initial allocation to minimize the efficiency level so that we did not do any of the work the mechanism is supposed to do. For our experiments the initial allocation had an efficiency level of zero. We used the standard double oral auction,[12] in which agents can make bids or offers for any number of units in any market and can accept standing bids or offers for any number of units up to the number offered or requested. To be sure that we gave this institution its

which is neither quasi-linear nor risk neutral. The theory and measurement could be modified to adjust for nonrisk-neutral subjects, but we chose not to do so since we wanted to concentrate on mechanism performance from the point of view of an unbiased, risk-neutral planner. Section 7 contains data that suggest that assuming risk-neutral subjects was not at excessive variance with reality.

[10] The $a^i$ choices confronting the subjects are designed to be similar in spirit to those faced by a Space Station payload operator who must design an instrument and use some Station resources to produce output.

[11] As a way to provide a feel for the design, we plot the distribution of the expected value of user benefits (the expected subject payments $M^i$) of a random selection of 30,000 combinations of configurations that fit in the capacity constraints. The combinations were found as follows: first an individual valuation sheet $i$ was selected at random (without replacement), and then one of its configurations $(x_i, y_i)$ was selected at random and placed in $A$. Next, another individual valuation sheet was randomly selected without replacement along with one of its configurations. This was placed in $A$ if there were room left, $B$ if there were room left, otherwise it was discarded. This process continued until the set of available valuation sheets was exhausted. We then calculated the expected value of $\Sigma_i M^i(a^i)$ and started the selection process over again. The expected value of the payments for an allocation as a percent of the maximum is a measure of the desirability of that allocation. We call this percent the (risk-neutral) *efficiency level.*

[12] We used the Caltech computerized version that allows for up to 19 simultaneous markets. For more information on this software, see Johnson, Lee, and Plott (1988).

FIGURE 2

DISTRIBUTION OF EFFICIENCY (30,000 RANDOM COMBINATIONS)

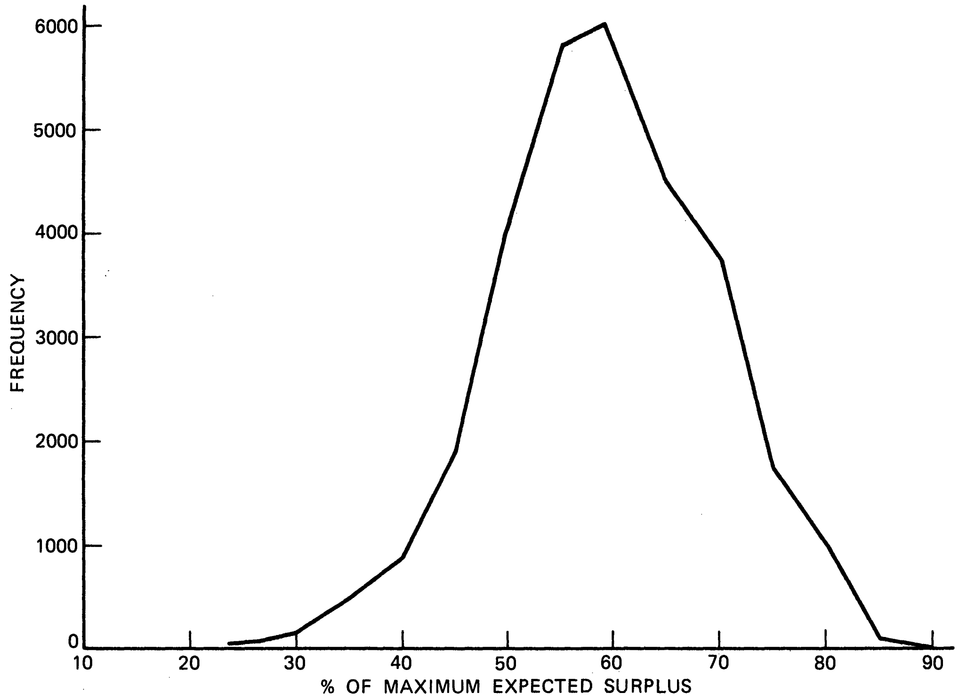


best chance to produce efficient allocations, we used subjects who were Caltech undergraduates and who were *experienced* in double oral auction experiments. They used the same personal computers and software across 19 markets. Further, we gave each subject the same payoff tables and endowments in each trial.

□ **The result.** The experimental data, some of which appear in Figure 3 and Table 4, show that this environment was overwhelming for the double oral auction. Although efficiency reaches 80% (but not until period 7), the data suggest failure. The average efficiency is 66.4% with a range of [43, 83.1]. Early periods are the worst. There are several reasons for this performance that are traceable to the nonconvexities, but the primary one seems to be the inability of the agents to decide in which market they want to be involved. There is no clear evidence of monopoly (or any other strategic) behavior, only of the inability of

**TABLE 2 Optimal Demand Configuration***

| | Contingent Allocation of $A$ $X = 20$, $Y = 20$, Capacity | | | | Contingent Allocation of $B$ $X = 20$, $Y = 20$, Capacity | | | |
|---|---|---|---|---|---|---|---|---|
| | Sheet | $X$ | $Y$ | Value | Sheet | $X$ | $Y$ | Value |
| | 1 | 12 | 9 | $3.25 | 4 | 8 | 12 | $2.75 |
| | 2 | 3 | 6 | $1.25 | 5 | 12 | 7 | $2.50 |
| | 3 | 5 | 4 | $2.00 | | | | |
| Total | | 20 | 19 | $6.50 | | 20 | 19 | $5.25 |

* Expected value = (5/6)($6.50) + (1/3)(5.25) = $7.17.

**TABLE 3 Contingent Contract Specification**

| Market # | 1 Unit of This Contract Provides |
|---|---|
| 1 | 1 unit of $X$ in event 5 ($A$ goes on date 1) |
| 2 | 1 unit of $Y$ in event 5 |
| 3 | 1 unit of $X$ in event 1 ($A$ goes on 1, $B$ goes on 2) |
| 4 | 1 unit of $Y$ in event 1 |
| 5 | 1 unit of $X$ in event 3 ($A$ goes on 2) |
| 6 | 1 unit of $Y$ in event 3 |

the markets to provide clear and predictable signals to coordinate the activities of the agents and to allocate them to the appropriate markets. There seems to be no obvious way to adjust the market institution to yield higher efficiencies. This leaves open the question, can any other existing institution do better?

## 4. Administrative processes

■ The first instinct of most engineers would be that some form of centralized project management would clearly provide the coordination needed to solve the problem. Administrative processes come in many forms, but we shall concentrate on the one we feel is not an unreasonable abstraction of NASA's STS pricing and allocation policy before the "Challenger disaster." That policy consisted of a posted price, which was zero for NASA-sponsored payloads, and an allocation policy based on exogenous priority assignment and a first-come, first-served (which is essentially random) selection of available payloads.

We consider the following mechanism. Agents pick their designs, $a^i$, and submit a resource-requirement vector $x^i$. An "administrator" then selects from $\{x^1, \ldots, x^I\}$ ran-

FIGURE 3
MEAN EFFICIENCY PER PERIOD

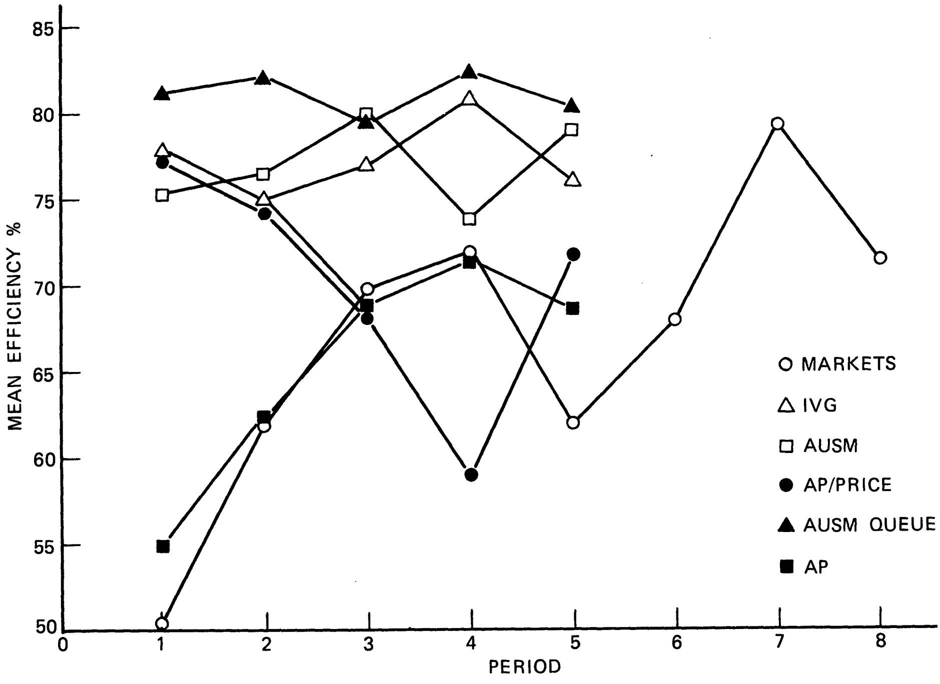

**TABLE 4 Summary of Experiments**

| Mechanism | Number of Experiments | Total Number of Periods | Subject Pool | Efficiency $\mu$ | Efficiency $\sigma$ | Efficiency Range |
|---|---|---|---|---|---|---|
| Administrative Process (AP) | 5 | 17 | Caltech, Pasadena City College | 63.5 | 10.0 | [39, 76] |
| Adaptive User Selection Mechanism (AUSM) | 4 | 20 | Caltech, Pasadena City College | 77.7 | 4.1 | [71, 86] |
| Adaptive User Selection Mechanism with Queue (AUSMQ) | 4 | 20 | Caltech | 80.8 | 4.0 | [72, 86] |
| Iterative Vickrey-Groves (IVG) | 4 | 20 | Caltech | 77.9 | 6.8 | [60, 91] |
| Markets* (M) | 3 | 24 | Caltech | 66.4 | 10.2 | [43, 83] |
| Administrative Process with Prices (APP) | 3 | 13 | Caltech | 69.9 | 8.1 | [58, 81] |

* Complete contingent markets and subjects experienced in multiple markets.

domly without replacement. At each draw the agent receives a contract for resources $x^i$ contingent on $A$'s launching, if such a contract will not require more resources than are available if $A$ launches. If it would, they receive a contract for $x^i$ contingent on $B$, if the contingent supply is still available. If not, they receive no contract. Thus, contingent resources are assigned by a first-come, first-served principle where the arrival time is under the control of neither the agents nor the administrator. We call this mechanism the administrative process.

□ **The analysis.** The mechanism is not designed to achieve efficient allocations. To see why this is true from a theoretical perspective, let us consider the allocations that result from Nash-equilibrium behavior in the game in which each agent picks his design $x^i$, given the designs of others, to maximize his expected utility, where the uncertainty includes that *created by the process.*[13] It is difficult to say much in general about the characteristics of these equilibria, but we can say something fairly interesting for our specific example—the experimental design. For that specific environment, if agents $1, \ldots, 5$ choose the design appropriate for maximum efficiency (see Table 2) and if agent 6 chooses $x^6 = (7, 11)$, then these choices constitute a Nash equilibrium for the game created by the first-come, first-served mechanism.[14] Thus, it is possible *in some environments* for this mechanism to induce agents to select the correct designs if they can find the Nash equilibrium. Even if they do, however, the efficiency performance of the mechanism is still only 83% of the maximum. This is so since the random selection process sometimes assigns agent 6 to shuttle $A$ and sometimes agent 1 is rejected. Although efficient designs are selected, efficient allocations of resources are not made.[15] Presumably, with inefficient designs, overall efficiency deteriorates even more.

□ **The experiment.** The real issue, of course, is to determine what might happen in practice. Can agents find good allocations? How does the mechanism perform? To answer these

[13] We do not consider the uncertainty created by the incomplete information about others' payoff sheets, since we want to show what happens even if there is much common information.

[14] The calculation is tedious but can be done on a personal computer.

[15] We do not know whether the vector of efficient payloads is the only equilibrium. We also do not know whether there are other equilibrium joint payload designs $(x^1, \ldots, x^6)$ that yield higher efficiencies than this equilibrium.

questions in a controlled experiment we took the environment described in Section 3 and allocated resources by using the administrative process. Subjects were told that there were two markets, each with capacity $X = 20$, $Y = 20$. A contract in market 1 corresponded to a contract to deliver contingent on states 1, 2, or 3 in Figure 1. This type of contract is called a *priority contract* (Chao and Wilson, 1987; Harris and Raviv, 1981). Such a contract specifies the order in which resources are dispatched in case a curtailment is necessary when there is excess demand. There are many practical examples of this type of contract. A contract in market 2 corresponded to a contract to deliver contingent on state 2 in Figure 1 (a lower priority). Subjects could submit only one order consisting of an $x$ and $y$ configuration and a preference ranking over markets 1 and 2. The orders were collected and randomly selected, one-at-a-time, from a box and placed in the first market with capacity available in accordance with the stated preference rankings. When all the capacity or orders were exhausted, a die was rolled. The orders in market 1 were filled if any of the numbers 1 through 5 appeared, and the orders in market 2 were filled if the number 1 or 2 appeared. These probabilities are taken from Table 1 and were known by all participants before they placed their orders. If a participant's order were filled, she received the value associated with the configuration ordered. Since prices were not used to allocate resources, subjects did not need to pay anything. The process was repeated for a number of periods: subjects were allowed to change their orders between periods.

□ **The result.** How does this administrative process perform? The answer seems to be, for the experimental environment we created, that this mechanism does neither better nor worse than markets. There appears to be no statistically significant difference in the average efficiencies achieved by either process. Further, their performance over time also appears to be identical.

To see whether posting prices for priority contracts would improve the administrative process, we modified the experimental institution in a way that parallels NASA's STS pricing policy. Prices for $X$ and $Y$ were posted, and subjects were told that if they agreed to pay these prices, they would be first to be allocated resources. After payers were allocated contracts, nonpayers were then allocated contracts. Although policymakers do not have the information required to post "correct" prices, we decided to give posted pricing its best chance by posting those prices that would lead subjects 1, 2, and 3 to choose their (efficient) payload designs in Table 2, if they treated prices as given. We let $P_x = 9$ and $P_y = 10$.

What happened is not surprising. Although some subjects did pay in the initial periods, later all decided a free ride with a lower probability of inclusion was preferable to a costly ride with a higher probability of inclusion. In the end, posting prices had no effect. Periods 4 and 5 of the administrative process with prices are similar to periods 1–3 of the administrative process without prices. Performance is the same as if no price were posted. Total efficiencies are not statistically different.

One must conclude that neither the administrative process nor markets perform very well in this environment. This leaves open the question, can a designed institution do better than the naturally evolved institutions have done?

## 5. A Vickrey-Groves mechanism

■ It has been well known from mechanism theory that there exists an optimal mechanism design for the environments in which we are interested when social efficiency is the performance standard. That mechanism, introduced by Vickrey (1961) and analyzed more generally by Groves (1973), is based on a modification of the second price (or first rejected bid) auction for single-dimensional problems. The key feature of this mechanism is that the price charged to any user is a function only of what the *other* participants bid. The mechanism asks each buyer to report his entire willingness-to-pay (*ex ante* utility) function. The allocation is then selected to maximize the sum of reported utilities. Each buyer $i$ pays an amount equal to the difference between the total utility the others would have gotten had $i$ not participated and what they actually get, according to their reported functions. Various lump-sum payments can also be levied. This creates the correct incentives (it is a dominant strategy) for participants to reveal correctly their willingness-to-pay for all contracts. Further, *ex ante* efficient allocations will result when participants use these (truthful) dominant strategies. Given these facts, it would be natural to apply this Vickrey-Groves mechanism to our problem. There are, however, at least two characteristics of that mechanism that can create problems in applications.

First, as Vickrey (1961) recognized, balancing payments is generally not possible. For public enterprises, such as the Space Station, this is not serious if, as we have modelled it, the supply decisions have already been made. As long as the government is interested in (short-run) efficiency and not monopoly revenue, one can ignore the problem of balancing the budget. After all, current policy that is similar to the administrative process suffers from the same failure. For private-sector applications, balancing may be more important. Nevertheless, we were willing to ignore this difficulty.[16]

The second characteristic of the Vickrey-Groves mechanism—that each bidder must report an *entire payoff function*—cannot be ignored. In most applications with environments characterized by multiple units and multiple dimensions, such as the Space Station, the informational requirements of such mechanisms are likely to prove intractable. Not only is communication difficult, but also much of the information required from an agent may not be readily available to that agent. Although any demander might be able to describe his benefits from any particular configuration of resources, given enough time, each usually knows best only those configurations in a neighborhood of what he expects to receive (Hayek, 1945; Marschak, 1972). To confront these informational realities, we modified the standard Vickrey-Groves mechanism to require buyers to specify only one configuration and a willingness-to-pay per report, i.e. a single demand point. Our intuition, based on others' research with iterative mechanisms (Smith, 1980; Ferejohn, Forsythe, and Noll, 1979), was that if prices at each iteration were consistent with the Vickrey-Groves logic, then with sufficient iterations information will be generated that would guide users to the efficient allocations in an incentive-compatible manner.

□ **The analysis.** The new process, which we call the iterative Vickrey-Groves process, proceeds as a communication *tâtonnement* until a particular stopping rule is applied. At each iteration $t = 1, 2, \ldots$, each $i \in N$ submits a "bid" $(d^i, b^i, f)$. $d^i$ indicates the vector of resources requested, $b^i$ is an amount $i$ is willing to pay for $d^i$, and $f \in \mathbb{F}$ describes the conditions under which delivery occurs. We call $f$ the contract type and consider $\mathbb{F}$ part of the mechanism design. For example, one could let $\mathbb{F} = \{A, B\}$, where $f = A$ means resources $d^i$ are delivered if and only if $A$ launches. Or $\mathbb{F} = \{1, \ldots, 6\}$, where $f = 3$ means delivery if and only if event 3 in Figure 1 occurs.

Once the bids are received, individualized charges are computed as follows. For each $f$ define $N_f$ as those users submitting a request for a contract of type $f$. Let

$$\Gamma(f) = \{\gamma \subseteq N_f : \sum_{j \in \gamma} d^j \leq \bar{y}\},$$

$$K_f = \underset{\gamma \in \Gamma(f)}{\text{argmax}} \sum_{j \in \gamma} b^{jf},$$

$$\gamma_i(f) = \underset{\substack{\gamma \in \Gamma(f) \\ i \notin \gamma}}{\text{argmax}} \sum_{j \in \gamma} b^{jf}$$

[16] If all agents viewed the mechanism as an incomplete information game and if they independently (without interaction) agreed that $G^i = G^i(a_i, \theta^i)$ for all $i$, that $\theta = (\theta^1, \ldots, \theta^I)$ is distributed according to $g(\theta)$, and that $g(\theta)$ represents *objective* common knowledge—that is the game is the same from all players' points of view (Harsanyi, 1967–1968, 1980) for these points—then one could adapt the approach of d'Aspremont and Gerard-Varet (1979) to balance the budget. We do not do this because we do not think there is any basis upon which to identify an *objective* common knowledge prior $g(\cdot)$ for the applications in which we are interested.

subject to $\sum_j d^j + d^i \leq \bar{y}$ and $\beta^i(f) = \max_{\gamma \in \Gamma(f)} \sum_{i \in \gamma} b^{jf} - \sum_{r \in \gamma_i(f)} b^{rf}$. If the process were to stop at this iteration, then $i$ would pay $\sum_{f \in F} \beta^i(f)$ if $i \in K_f$ and 0 otherwise. To understand what $\beta^i$ is, one first notices that, for a contract of type $f$, $\Gamma(f)$ identifies all feasible coalitions; i.e., all groups of users whose collective bids are feasible, while $K_f$ selects the coalition with the maximum sum of bids. If $i \in K_f$, then $\gamma_i(f)$ is simply $K_f - \{i\}$, while if $i \notin K_f$, then $\gamma_i(f)$ identifies the coalition in $\Gamma(f)$ that would remain feasible if $i$ were added and that maximizes the sum of bids of its members. Thus, joining with $\gamma_i(f)$ is $i$'s "best chance" of acquiring a contract of type $f$, given the behavior of the other participants. Given a vector $d$ of resource demands, the "price" $\beta_i(f)$ that trader $i$ faces for contract $f$ is equal to either the social cost of $i$'s being in $K_f$ in terms of revenue foregone by $i$'s inclusion, or the minimum amount $b^i$ needed to become a member of $K_f$ when other traders' bids are held constant. In the former the first term on the right-hand side of the price equation is the amount generated if $i$ did not participate; subtracting the bids by other members of $K_f$ gives an equivalent version of the Vickrey-Groves "second-price" auction for a single unit of a good. Given $d$ and assuming risk neutrality, bidding one's expected value for a contract is the dominant strategy. In the latter, the first term is simply the sum of bids of the members of $K_f$. Thus, subtracting the bids of $i$'s best-chance coalition gives the amount $i$ would have to bid to be allocated resources in $f$.

Before proceeding to the next iteration, each $i$ observes $\{\beta^i(f)\}$ as well as $d$. Thus, at each trial, the participants gain information concerning the demand not only for resources and contracts that they (provisionally) acquire, but also for those that they do not. In this way, as the iterations proceed and bidders search for their "best" alternative, the mechanism may lead to an efficient outcome, if participants adopt the "short-run" dominant strategy of bidding their true value.

The process stops at $t$ if $K_f$ at $t$ is the same as $K_f$ at $t-1$, $\forall f \in F$. That is, the process stops when the set of participants acquiring resources for each contract type remains unchanged. It is relatively easy to see that if users bid their true expected values for $(d^i, f)$, then this mechanism will select the efficient combination of uses, $K_f$, given the requested $d^i$. It is not at all obvious, however, how a user either should or would select $d^i$ and $f$. Inefficient choices might occur. Nevertheless, we expected that this process would improve the administrative process described in Section 4, since the iterative Vickrey-Groves process at least selects efficient combinations of configurations, even if those configurations may not be fully optimal.

□ **The experiment.** Determining allocations and calculating individual prices for this mechanism is an enormous task that cannot be done by hand in an effective manner. Thus, for testing this mechanism, all communication and calculations were made by using a network of personal computers. The computers were connected on a local area network with a central controller computer that received messages, calculated prices, determined allocations, and sent messages.

At the beginning of a trial in a market period, an individual would submit a configuration $a^i$ and select a priority, either market 1 or market 2 (but not both). As in the administrative process, market 1 corresponds to the priority contract (deliver contingent on states 1, 2, and 3), while market 2 corresponds to the contract (deliver contingent on state 2). The subject would then enter a bid for the configuration. After each individual sent his message to the center, it calculated the provisional allocation and prices for each participant. Each individual user was then informed of the provisional traders in each contract and of their configurations, on the basis of the trial messages. In addition, each subject received a private price message that described his potential payment if he was part of the provisional allocation, or (if he was not) the amount he would have to bid to have his configuration included in the provisional allocation. The algorithm used to calculate allocations and prices and to return this information to subjects took less than one second to transmit after the last message was entered. Furthermore, each subject had displayed on his screen the history of the last three trials, including provisional allocations and prices.

The stopping rule for allocating the contracts was partially sequential. In particular, the process stopped if the same subjects and configurations occurred in the markets three times in a row (rule $A$). Otherwise, market 1 closed after $t_1$ trials were exhausted; market 2 closed after $t_2$ trials if rule $A$ were not executed, where $t_2 > t_1$.

The only restrictions on the individuals' messages were that $b^i > 0$ (and integer-valued) for each trial, that the $a^i$ must be one of $i$'s nine choices, and that a subject could submit a bid for market 1 or market 2, but not both markets in the same trial. There was no ratchet (improvement) rule for individual bids in the process. Thus, a bid was not necessarily binding because one can bid "very high" in trial $t$ and then bid almost zero in trial $t + 1$. We chose this set of rules to allow individuals to search "easily" for combinations.[17]

□ **The result.** The iterative Vickrey-Groves mechanism yielded higher efficiency levels than either markets or the administrative process. The iterative mechanism attained a mean efficiency of 78% with a range of [60, 91], a significant increase from 64%. Of particular interest was the fact that the iterative mechanism dominated both markets and the administrative process on a period-by-period basis: it attained high efficiencies early and maintained them.

Unfortunately, an efficiency of 100% was not achieved. In fact, the iterative mechanism did not even produce stable levels of efficiency. It dipped as low as 60% (easily obtainable at random). One explanation for this instability is the fact that the process was pushed to the final trial in many cases, and this led the final bid for $d^i$ to be somewhat random. Coordination across user designs is lacking. Although one trial produced an efficiency level of 90%, the reliability of the process is suspect.[18] We provide more details of performance in Section 7.

## 6. An ascending-bid auction

■ Feeling unhappy with the complexity of the rules for pricing in the iterative Vickrey-Groves mechanism, but feeling that iteration would be helpful if more coordination were evident and if some form of commitment to designs were required, we next designed a mechanism by modifying the English (or ascending-bid) auction commonly used to sell art objects, livestock, and tobacco.[19] We call our version the adaptive user selection mechanism (AUSM). This mechanism does not require all participants to be in the same room (as in Sotheby's art auction); they can communicate "bids" through an electronic bulletin board. Nor does it require a rapid sequence of bids to be made (as in the art auction);

---

[17] The rules used were developed in response to early testing. Initially, we did not close the markets sequentially; however the subjects pushed the process to the last trial in most instances (even with 40 trials). Since an individual could only bid for market 1 or 2 (not both), there was substantial excess supply at the close. Also, we had instituted a rule in the early testing that required individuals to better *their* previous bids in the market in which they were ordering. This rule caused individuals to be cautious in their bidding or "locked" them into larger projects that yielded low efficiency levels and so we eliminated it.

The instructions for this and other experiments reported below are available from the authors on request.

[18] One modification of the iterative Vickrey-Groves mechanism, proposed by seminar participants at New York University, that might stabilize its performance and improve its efficiency would be to allow the algorithm that computes prices and allocations to accumulate the data. That is, at each iteration all past bids would be used; more and more points on the demand surfaces would come into play. In the limit (with infinite iterations) one might achieve the full Vickrey-Groves scheme.

[19] See Cox, Roberson, and Smith (1982) and Milgrom and Weber (1982) for descriptions and analyses of ascending-bid auctions.

participants can be allowed any length of time thought to be desirable to consider their demands. AUSM is not a spot market and requires no auctioneer. It is a decentralized mechanism that guides coordination in design and that selects high-yield projects.

□ **The analysis.** The English auction, upon which our mechanism is based, is a non-*tâtonnement* process that is widely used to auction single items of uncertain value to multiple bidders. At each instant during this type of auction there is a *potential allocation,* which is common knowledge. Any agent can enter a *bid* at any time. The bid is common knowledge. There is a common *update rule* that specifies how a new bid can create a new potential allocation. The process stops when no new bid is made soon enough after the last bid. The potential allocation is then the actual allocation.

For auctions of single items the potential allocation is usually expressed as "the item goes to the current highest bidder who will pay his bid," and a bid is "a stated willingness-to-pay." The update rule is that the person bidding becomes "the current highest bidder at that bid" if his bid is higher than that of the current highest bidder. If not, no change occurs.

For multiple contracts of multiple dimensions the principle is exactly the same. There is a supply of each of $\mathbb{F}$ contracts to be allocated. The capacity of each is $\bar{y} \in \boldsymbol{R}_+^k$. We can easily modify this to accommodate an environment in which $y$ depends on $f \in \mathbb{F}$. For our experimental environment, $\mathbb{F} = \{1, 2\}$ (the priority contracts), and $\bar{y} \in \boldsymbol{R}_+^2$. A potential allocation is a feasible collection of contracts. A bid is simply a proposed contract $(d^i, b^i, f)$. A bid replaces a contract (or group of contracts) in the potential allocation if and only if the $b^i$ is higher than the sum of the bids offered by those being replaced. More formally, let $K_f$ be the agents who hold contracts in the current potential allocation of $f$ and let $R \subseteq K_f$. If $Z_f + \sum_{j \in R} d^j \geq d^i$ and $b^i \geq \sum_{j \in R} b^j$, where $Z_f = \bar{y} - \sum_{w \in K_f} d^w$, then $(d^i, b^i, f)$ replaces the collection $\{(c^j, b^j, f)\}_{j \in R}$. If there is no such $R$, then the new allocation equals the old (i.e., $i$'s bid is rejected). If there are more than one such $R$, we assume that $i$ replaces the $R$ with the smallest value of $\sum_{j \in R} b^j$.

The potential allocation can be publicly displayed on a (computerized) bulletin board as, for example, in Table 5. For this example, if bidder 2 wanted contract 1 in the amount of $(x, y) = (10, 3)$, he could do so by bidding $(10, 3, 0)$. If bidder 2, on the other hand, wanted $(x, y) = (12, 6)$, he would have to bid at least 201 (to displace bidder 3). If bidder 2 wanted $(x, y) = (12, 11)$, he would have to bid at least 501 (to displace bidder 7) and if bidder 2 wanted $(6, 16)$, he must bid 701 (to displace both 3 and 7).

We chose this basic mechanism for several reasons: (1) the practical success of the single-unit English auction, as signalled by its widespread use; (2) the feeling, based on experimental experience, that in an environment in which the bases for common knowledge are little understood or controlled, iterations with commitment allow subjects to "feel their way" in a manner in which sealed-bid, one-shot auctions do not;[20] and (3) a theoretical analysis of its properties. Let us briefly expand on the last.

We emphasize two facts about our adaptive user selection mechanism-bulletin board. First, given a proposed allocation, any bidder $i$ can, with a high enough bid, change the proposed allocation to one in which $i$'s contract $f$ is for any amount less than or equal to $y$. Second, the proposed allocation puts a lower bound on how much $i$ must bid to achieve any desired allocation on contract $f$. Let $\psi^*$ represent a set of contracts or, alternatively, a potential allocation. Let $\xi^i(\psi^*)$ represent the set of allocations to which $i$ can unilaterally cause $\psi^*$ to be changed with some bid. We call $\psi^*$ a *simple equilibrium* if $\psi^*$ is feasible and if $\forall i = 1, \ldots, n, \xi^i(\psi^*) \cap \{\psi \mid V^i(\psi) > V^i(\psi^*)\} = \varnothing$, where $V^i(\psi)$ represents the (expected)

[20] Of course, if enough contingent bids could be submitted in a sealed-bid auction, it could mimic an iterative process, but in an informationally more complex manner.

**TABLE 5** **An Adaptive User Selection Mechanism Configuration**

| Contract 1 | | | | Contract 2 | | | |
|---|---|---|---|---|---|---|---|
| Bidder | $x$ | $y$ | Bid | Bidder | $x$ | $y$ | Bid |
| 7 | 2 | 10 | 500 | 1 | 15 | 2 | 50 |
| 3 | 6 | 5 | 200 | 4 | 1 | 10 | 75 |
| | | | | 5 | 4 | 8 | 100 |
| Supply | 20 | 20 | | Supply | 20 | 20 | |
| Slack | 12 | 5 | | Slack | 0 | 0 | |

utility $i$ receives if the potential allocation $\psi$ is actually provided. That is, no $i$ can unilaterally improve his position, since any bid high enough to cause $i$'s quantity $d^i$ to be included in the allocation of contract $f$ will be higher than the value of the benefits attained from those $d^i$ units of contract $f$. Simple equilibria are contract allocations that are individually "stationary" allocations of AUSM. This is a fairly large set, not all of which are desirable. Further, we feel that reasonably well-informed traders will be able to avoid some of them. To see how, consider a slightly different mechanism.

Suppose that each $i$ chooses a contract $m_i = (d^i, b^i, f)$. Given $m = (m_1, \ldots, m_n)$, a potential allocation of contracts $\psi^*(m)$ is chosen as follows: for each $f$ pick $K_f$ to maximize $\sum_{i \in K_f} b^i$ subject to $\sum_{i \in K_f} d^i \leq \bar{y}$. Then allocate to $i$ the contract $f$ in the amount of $d^i$, $b^i$ if $i \in K_f$. One can think of this as a game, $G$, with strategies $m^i$ and outcome function $\psi^*(m)$, with allocations picked to maximize the aggregate *stated* willingness-to-pay. It could be used as a "sealed-bid" mechanism. We call $\psi^*$ a *noncooperative equilibrium allocation* of $G$, if $\psi^* = \psi^*(m^*)$ and for each $i$, $V^i(\psi^*) \geq V^i(\psi^*(m^*/m^i))\ \forall m^i$, where $(m^*/m^i)$ is the vector $m^*$ with $m_i^*$ replaced by $m^i$. It is obvious that if $\psi^*$ is a noncooperative equilibrium allocation of $G$, then $\psi^*$ is a simple equilibrium of AUSM. The converse is not necessarily true.

On the basis of previous experimental experience with games such as $G$, it would not be unreasonable to expect in experimental testing with replications that the final allocations would be noncooperative equilibrium allocations of $G$. Not *all* simple equilibria will occur in replicated situations when subjects can learn to avoid "bad" dynamics. Of course, what the mechanism designer is really interested in is not the equilibrium, but the efficiency of the equilibrium allocations. Unfortunately, even if only noncooperative equilibria of our mechanism occur, the associated allocations may not be desirable. Because of the lumpy nature of the users' projects, there are noncooperative equilibrium allocations of $G$ that are not efficient contract allocations. There may be changes in those allocations involving several traders *simultaneously* that can make all better off. In particular, if during the auction there is a large user who is part of the current potential allocation and who has a fairly high bid, it may be too costly for any one small user to displace him, even if it is possible that several small users can together receive more benefits than the single larger user. In this situation unilateral actions by one user are not sufficient to drive the mechanism to a more efficient allocation of contracts.[21]

In our initial testing of the adaptive user selection mechanism-bulletin board, we had hoped that the subjects would overcome complications caused by the variable size demands. But early data (reported in detail below) suggested efficiency levels of only 75–85%. We

[21] If omitted users could replace just the marginal units of those users in the potential allocation, then it would not be costly to displace part of a large user. To do this, however, users would have to be allowed to express a bid for each unit they wish to buy, which is an entirely different mechanism.

therefore felt it important to try to overcome this limitation of the mechanism. To do so we had to improve the ability of the mechanism to recognize when to replace one large user with two or more small users. Our solution was not only to allow small users to coordinate their bids, but to encourage them to do so. Thus, we modified AUSM in the following ways. We allowed a public "standby" queue, in which any agent could post a "proposed bid" $(d^i, b^i, f)$ that he would be willing to have included in a coalitional bid. Because of the possibility of joint bids from a group $\gamma$ of individual agents, we expected outcomes to change when a standby queue is made available.

Let $\gamma \subseteq N$ be an arbitrary coalition of agents and let $(m^*/m^\gamma)$ be the vector $m^*$ with $m_i^*$ replaced by $m_i$ for all $i \in \gamma$. We call $\gamma^*$ a *strong noncooperative equilibrium* of $G$ if $\psi^* = \psi^*(m^*)$ and for each coalition $\gamma$ and each $m^\gamma \neq m^{*\gamma}$ there is at least one $i \in \gamma$ such that $V^i(\psi^*(m^*)) > V^i[\psi^*(m^*/m^\gamma)]$. If $\psi^*$ is a strong noncooperative equilibrium of $G$, then $\psi^*$ is a noncooperative equilibrium of $G$. The converse is not necessarily true.

Our hope was that offering the subjects the opportunity to coordinate publicly their bids through the queue would lead them to strong noncooperative equilibria of the game $G$ (if such equilibria existed). If that occurred, then this variation in AUSM rules would solve our problem since those equilibrium allocations of $G$ are efficient.

□ **The experiment.** As before, we created two markets (priority contracts) with market 1 corresponding to $A$ in Figure 1 and market 2 corresponding to $B$. When the markets opened, subjects would submit an order consisting of a market or the standby queue, an $x$ and $y$ choice, and a bid. Their order would be accepted if it could fit within the available capacity of the market requested, if it could displace existing orders with lower bids, or if the standby queue were requested. If a subject wanted to use the standby queue, he had to indicate for which market the bid was tendered. Furthermore, if a subject's bid in the standby queue were combined with another order, then any standing order the subject had in a market was cancelled. Finally, to aid in the search process for the best configurations, subjects were allowed to move existing configurations to other markets or change their configuration and bid if the new bid could fit in the available capacity. If a subject did change his configuration in a market, however, he had to improve the bid of the total orders he was displacing *including* (if necessary) his original order. For example, suppose the orders in market 1 were as follows:

Market 1

| Subject | $X$ | $Y$ | Bid |
|---|---|---|---|
| 2 | 12 | 9 | 150 |
| 4 | 5 | 4 | 100 |
| 5 | 3 | 6 | 75 |

If subject 2 wanted to change his configuration to $x = 12$, $y = 13$, he would have to bid more than 225. The minimum bid improvement increment was set at 5. If an order was displaced, the subject was allowed to reorder through the process above and submit any feasible order he wanted. The auction stopped when there were no new orders or order changes within 30 seconds of the last order. When the market closed, a die was rolled. On the basis of the data in Table 1, if the number 1 through 5 appeared, the orders in market 1 were filled. If the numbers 1 or 2 appeared, the orders in market 2 were filled. If an order was filled, the subject was given his redemption value minus his bid. If a subject's order was not filled, his bid was subtracted from his accumulated earnings. If a subject did not have an order in a market, he received zero earnings for the market period. Since losses were possible,[22] at the beginning of the experiment each subject was given $7 of working capital to add to his earnings. We also tested our mechanism without a queue.

□ **The result.** AUSM operated smoothly. When run with a queue, it outperformed markets, the administrative processes, and the iterative Vickrey-Groves mechanism. The mean efficiency was 81% with a range from 72–86%. Our mechanism with a queue dominated all others in each period. Two other interesting facts were the small variance in the levels of efficiency achieved and the high level of revenue collected.

While AUSM yielded more efficient allocations than the others, it also did not reach 100%, and, in fact, never even reached 90%. It is important to realize, however, that the experimental environment made it particularly difficult to better 85%. There are few configurations (see Figure 2) that generate efficiencies higher than 85%, and those that do are very sensitive to the decisions of a single participant. Conversely, it is easy for efficiencies to fall below 70%. Thus, both the iterative Vickrey-Groves mechanism and AUSM systematically found low probability, but high efficiency, allocations.

## 7. Experimental results

■ In this section we provide the details that support our earlier observations. We measure three aspects of mechanism performance: efficiency, revenue, and individual behavior. The overall performance of the mechanisms is determined by using the risk-neutral expected values of (*ex ante*) efficiency. We also consider the revenue generated by each mechanism. Finally, we evaluate the data to see whether we can learn anything about individual behavior.

□ **Efficiency.** The mean efficiency (percent of the maximum expected value ($\mu$)) and the associated standard deviation ($\sigma$) for each mechanism appear in Table 4. The nature of the underlying distribution of combinations found in Figure 2 suggests that we should use nonparametric methods for our statistical analysis (Lehmann, 1975). We use the Wilcoxon rank sum to test the equality of distributions of efficiency generated by each mechanism. In particular, the $z$ scores are derived from testing the hypothesis of equality of distributions versus strict inequality of distributions. The Wilcoxon rank-sum test for each mechanism is in Table 6. We have pooled the data from each experiment. While this affords more degrees of freedom, we may be biasing results if substantial learning occurs in a mechanism. Our statistical measures in Table 6 support the following ranking of the mechanisms by efficiency as follows: markets = administrative process with prices = administrative process < iterative Vickrey-Groves mechanism = adaptive user selection mechanism < adaptive user selection mechanism with a queue.

The mean efficiency per period for each mechanism appears in Figure 2. Notice that for the administrative mechanism and markets the level of efficiency tends to increase over

[22] We made two other variations in design in response to pilot testing. Originally, we had a stopping rule that the auction would run for $T$ minutes and the allocation at $T$ would be final. This had undesired effects. In the pilot experiments very little bidding occurred until $T - \epsilon$, when a flurry of bids were made. Allocations were essentially random. We easily solved this by changing the rule to the more traditional one in which the auction ends if no new bids occur after $S$ seconds, where $S$ is a design choice. The other variation concerned the commitment entailed in placing an order in a market. In one pilot experiment we used the same ordering process as above, except that subjects could remove existing orders and raise or lower bids while in a market. There was no queue, but combining to move to different markets was allowed. This had an effect similar to the first stopping rule. Without commitment nothing serious happened until $T - \epsilon$. We fixed this by revising the bidding rules so that each bid was a potentially binding contract. Further, we added an explicit improvement rule for bids. Finally, in our initial design of the standby queue we allowed participants with orders in the standby queue to veto proposals combining with their order. We abandoned this rule in favor of committed bids in the queue when we found no vetos.

**TABLE 6** **Rank-Sum Test (All Periods)***

| | IVG | AUSM | AUSMQ | AP | M |
|---|---|---|---|---|---|
| AP | $z = 4.39$<br>$\alpha = .000$ | $z = 4.36$<br>$\alpha = .000$ | $z = 4.72$<br>$\alpha = .000$ | $z = 1.81$<br>$\alpha = .07$ | $z = .63$<br>$\alpha = .53$ |
| IVG | | $z = -.004$<br>$\alpha = .480$ | $z = 1.58$<br>$\alpha = .057$ | $z = 2.63$<br>$\alpha = .009$ | $z = 3.45$<br>$\alpha = .001$ |
| AUSM | | | $z = 2.12$<br>$\alpha = .017$ | $z = 2.65$<br>$\alpha = .009$ | $z = 3.97$<br>$\alpha = .000$ |
| AUSMQ | | | | $z = 3.95$ | $z = 4.68$<br>$\alpha = .000$ |
| APP | | | | | $z = 1.12$ |

* $\alpha$ indicates the level of significance for the test that the efficiency of the mechanism in the column equals that in the corresponding row.

time. In particular, the mean and standard deviation of efficiency for periods 3 and above is 69.4 and 5.4 for administrative, and markets have a mean and standard deviation of 73 and 6.7, respectively. Thus, we see that efficiency increases with repetition in the administrative process and markets. Further, the reduction in the standard deviation in the later periods suggests that these higher efficiencies will be maintained. For the adaptive user selection mechanism with and without a queue and for the iterative Vickrey-Groves mechanism, there is no significant effect of repetition. The mean efficiency for periods 1 and 2 was 77.1 for AUSM, 81.6 for AUSM with a queue, and 76.8 for the iterative Vickrey-Groves mechanism, while the mean efficiency for periods 3 and above was 78.1 for AUSM, 80.2 for AUSM with a queue, and 78.6 for the iterative mechanism. The administrative process with nonzero price shows a decrease over time with a mean efficiency of 63.3 and standard deviation of 5.8 for periods 3 and above. These observations on learning strengthen our statistical results on efficiency from the pooled samples.

Since the inefficiency of markets may surprise some, let us look more closely at the data on market prices. Recall that our experimental design creates a condition in which there is no competitive equilibrium. One might still hope that prices could still stabilize in the markets. The mean and 95% confidence intervals for each market/period appear in Table 7. We can make three observations: (1) contract prices decrease over time and tend to equate across product dimensions ($X$ and $Y$) for each contract; (2) variance in contract prices falls over time, but is still high; and (3) the lowest probability (1/6) state markets (5 and 6) command higher mean prices than those of markets 3 and 4 (1/3 probability). Although this last observation is curious, notice that markets 1 and 2 combined with

**TABLE 7** **Mean and 95% Confidence Intervals on Prices in Francs***

| | Periods 1–4 | | Periods 5–8 | |
|---|---|---|---|---|
| Market | $\mu$ | (95%) | $\mu$ | (95%) |
| 1 | 8.4 | (7.6, 9.2) | 7.8 | (7.1, 8.5) |
| 2 | 7.5 | (6.1, 8.9) | 6.1 | (5.2, 7.0) |
| 3 | 5.3 | (3.8, 6.8) | 2.6 | (1.9, 3.3) |
| 4 | 4.1 | (3.0, 5.2) | 2.7 | (2.0, 3.4) |
| 5 | 3.2 | (2.3, 4.1) | 3.2 | (2.5, 3.9) |
| 6 | 4.2 | (3.3, 5.1) | 3.3 | (2.5, 4.1) |

* 1 Franc = 2 cents.

5 and 6 make up a priority contract. Subjects may simply be trying to form such contracts. The main conclusion we draw is that prices are not behaving as they do when equilibria exist. They are not able to smooth out these nonconvexities.

□ **Revenue.** AUSM is more efficient with a queue than without a queue. Counteracting this advantage is the possibility that agents will form "coalitions" via the standby queue and reduce the revenue. Data on the mean revenue (total and by market) for each treatment, and the associated standard deviation appear in Table 8. Very little revenue is generated in the administrative process since most of the requests (77%) were for nonpaying status. The addition of the standby queue to AUSM results in a higher mean revenue and a shift in the support to the right. If we look at the revenue generated market by market, we see that with a queue the volatility of revenue is fairly low for market 1 and high for market 2. Without a queue revenue from each market is relatively volatile. Of course, market 1 contracts received higher bids. Specifically, market 2 contracts have a mean bid ⅓ that of market 1 contracts (approximately the difference in probability of each market's being filled). Table 9 supplies the rank-sum and $t$-tests for the overall revenue generated by each of the adaptive user selection mechanism and iterative mechanism treatments.

We see that the existence of the standby queue results in significantly higher revenues, and this comes from higher revenue generated in both markets. We report the mean for periods 1 and 2 and periods 3+ for each mechanism in Table 10. There is a statistically significant upward trend in the revenue in both AUSM and the Vickrey-Groves mechanism that can be traced to the revenue generated in market 1. We find no such trend by examining the time series for AUSM with a queue.

□ **Individual choice behavior.**

*Administrative process.* As detailed in Section 4, we would expect the administrative process to generate choices that are consistent with "scaled-down" projects. Specifically, from the redemption value sheets in the Appendix, notice that if each individual chose the largest project, only one order would fit per market. If each subject chose his smallest project, however, everyone could fit in one of the markets. Actual orders balanced between these extremes. Of the 102 orders filled in the administrative process, only six orders submitted were an individual's smallest project, while nine orders had the largest project submitted. The average sizes ($X$, $Y$) submitted by subjects over time are in Table 11, which provides evidence of an updating phenomenon on the part of the subjects. As a final note, *the ranking of markets by probability* (see Section 5) *was never violated by a subject in the experiments.*

*Adaptive user selection mechanism.* While we did not design the experiments to test whether the three equilibrium concepts used in Section 6 were consistent with reality, we can extract some evidence from the data. In a strict sense all three are rejected. None of the 20 allocations achieved by our mechanism were even simple equilibria, much less noncooperative or strong noncooperative. Only four of the 20 allocations chosen by our mechanism with a

**TABLE 8** **Revenue**

| | Market 1 | | | Market 2 | | | Total | | |
|---|---|---|---|---|---|---|---|---|---|
| Treatment | $\mu$ | $\sigma$ | Range | $\mu$ | $\sigma$ | Range | $\mu$ | $\sigma$ | Range |
| AUSM | 302.2 | 52.9 | [154, 365] | 102.3 | 22.6 | [70, 145] | 404.5 | 48.7 | [284, 475] |
| AUSMQ | 353.7 | 36.2 | [300, 425] | 122.0 | 27.5 | [75, 185] | 475.7 | 52.0 | [375, 560] |
| IVG | 284.1 | 85.1 | [160, 449] | 108.1 | 66.3 | [00, 240] | 388.4 | 118.2 | [210, 656] |

**TABLE 9** **Rank-Sum and *t*-Test for Overall Revenue Generated***

| | AUSMQ | | IVG | |
|---|---|---|---|---|
| AUSM | $z = 3.69$<br>$\alpha = .00$ | $t = 3.59$<br>$\alpha = .07$ | $z = -.65$<br>$\alpha = .25$ | $t = -.55$<br>$\alpha = .29$ |
| AUSMQ | | | $z = -2.85$<br>$\alpha = .00$ | $t = -2.95$<br>$\alpha = .00$ |

* $\alpha$ indicates the level of significance for the test that the revenue generated by the mechanism in the column equals that in the corresponding row.

queue were simple equilibria. Some of this might be explained by risk-averse behavior, but we feel that risk is not so important in this environment.

We can gain a better insight by examining which agents were following best-response strategies when the auction closed. Individuals in market 1 are almost always (58 out of 59) best responding. 27% (15/56) of the nonbest responses come from subjects in market 2. These facts suggest that subjects may be "standing pat" in market 2. Almost 70% of the nonbest responses are from subjects who are not yet allocated space in any market. These data lead us to conclude, albeit tentatively, that noncooperative equilibrium allocations of the full-information game are likely to occur in AUSM, especially if the stakes were to be increased. Over 80% of the responses are within 10¢ of the theoretical best response. Aggressive behavior, competitive pressures, and the iterative nature of the process with commitment all seem to help lead to stationary allocations that look like Cournot-Nash, full-information, noncooperative equilibria.

*Iterative Vickrey-Groves mechanism.* We were especially interested in two aspects of individual behavior in this mechanism: (1) whether individuals bid their expected value (was it demand revealing?), and (2) whether there were any discernable "strategic" bids made during a market period. Many of the bids made during a market period were above expected values (40% of the bids), and this was even true on the last trial in a period (31% of the bids). This behavior did lead some individuals making those bids to pay more than their value to get a contract. Those taking such a loss (or after obtaining prices above expected values) generally did not repeat this bidding behavior.

Two particular aspects of strategic behavior used by subjects involved attempts to gain information about other individuals' bidding behavior. First, some participants would bid zero to see what their price would be and so had no effect on the prices of individuals who were provisionally selected. Second, the participants typically used all of the possible trials, which left resources to be allocated in a one-shot game in the final trial (85% of the time for market 1 and 75% of the time for market 2). To see whether the number of trials makes a difference in the efficiency of the mechanism, we conducted an experiment with 40 trials. For the first period the efficiency was 75.5% and revenue was 250; the second period concluded with 78% efficiency and 495 for revenue. Both periods used all 40 trials.

**TABLE 10** **Mean Revenue**

| | Periods 1 and 2 | Periods 3+ |
|---|---|---|
| AUSM | 374 | 425 |
| AUSMQ | 470 | 480 |
| IVG | 351 | 413 |
| AP | 178 | 141 |

**TABLE 11** **Average Project Size Submitted per Period**

| Period | Mean ($X$, $Y$) |
|---|---|
| 1 | (9.4, 8.9) |
| 2 | (8.5, 9.9) |
| 3 | (7.7, 8.3) |
| 4 | (6.9, 6.8) |
| 5 | (8.1, 8.7) |

*Markets.* Although there were only six participants in the experiment, trading in the markets was very active.[23] The trading patterns over the course of an experiment were not entirely stable, however. That is, individuals *did not* choose the same projects and use the same contracts in each period. We call this phenomena *market switching.* Evidence of this market switching is that very few subjects (28%) stayed with the same contract across periods; 60% of the subjects that did selected no contracts (they sold their endowments). But 73% of the subjects switched markets at least once and 56% switched at least twice in the last four periods. This evidence is consistent with the fact that markets could not coordinate allocations in our experiments.

## 8. Conclusion

■ In this article we have designed and analyzed six different mechanisms for solving a resource allocation problem in an environment similar to many scheduling applications. The two key characteristics of the environment are that demands are lumpy and ill-fitting and that supply is uncertain and unresponsive at the time of contracting. The mechanisms analyzed included markets and an administrative process, as representatives of institutions that have naturally evolved, and two newly designed mechanisms, the iterative Vickrey-Groves mechanism and the adaptive user selection mechanism.

The experimental results show that simply setting up enough markets is not the best way to proceed if the environment is characterized by significant indivisibilities and, *a fortiori,* nonconvexities. In the experiment markets performed no better than an *ad hoc* administrative process. Both mechanisms, modelled on naturally evolved institutions, yielded low efficiency levels.

Two newly designed institutions, the adaptive user selection mechanism and the iterative Vickrey-Groves mechanism, using priority contracts, significantly outperformed markets, especially in the early periods, when there is no information and no basis for common-knowledge priors. This is an especially important consideration for nonrepeatable operations such as the Space Station. Neither mechanism could consistently generate efficiencies greater than 85% in the demanding experimental environment.

Designing mechanisms for applications, as in designing America's Cup sailboats or airplanes like the Shuttle, is an art form to be guided by theoretical developments and experience. Our use of experimental methods has allowed us to accumulate some experience before the application.[24] Several observations based on that experience deserve mention.

---

[23] Recall that our participants were experienced with multiple markets and computerized bid and ask procedures.

[24] Indeed, the results of this study have already led to the development of a vastly more elaborate experimental environment to demonstrate the feasibility and performance of the various allocation mechanisms for use on NASA's Space Station. The more elaborate testbed featured more continuous demands and the existence of a unique competitive equilibrium. AUSM produced efficiencies at the 80% level, while markets were 89% efficient (80% in early periods) and the administrative process was only 60% efficient.

First, the transparency of a mechanism—the ease with which an agent is able to anticipate the results of any particular strategy—is important in achieving more efficient allocations earlier. Market prices in these environments in early periods were unpredictable and caused inefficient resource purchases. The administrative process is simple-minded, but there was so much randomness that subjects had difficulty coordinating designs. The iterative process is "obscure"; subjects were unsure what "prices" would be. Only AUSM was straightforward and stable.

Second, while *tâtonnement*-like iterations without commitment may guide coordination (it is a form of "cheap-talk"), our experience shows that some commitment is needed to stabilize responses and to speed convergence. Iteration allows feedback, reaction, and learning about the possibilities. In the administrative process in early periods subjects have to submit designs when they know nothing about their environment. The results were very low efficiencies. AUSM allows early mistaken guesses about what would fit, etc., to be changed in response to "tentative" allocations. Of course, if there is no commitment, the information generated by iteration would be useless. Thus, mechanisms need a delicate balance between commitment without an ability to adjust (as in the administrative process) and adjustment without the ability to commit (as in the iterative process). AUSM works because it attains such a balance.

Third, markets cannot smooth over all nonconvexities, especially in environments with little prior information. In our environment, markets exhibited the worst performance on day 1—worse than first-come, first-served. Eventually, as a common basis for "rational expectations" emerges, the performance of markets improves, but never reaches that of the newly designed mechanisms. In the theoretical analysis of mechanisms, the assumption that there are objective common knowledge priors obscures this type of consideration.[25]

The fundamental open question is, of course, whether there exist other allocation mechanisms that outperform those analyzed. There are "optimal auctions," designed to attack incentive difficulties in simpler problems, that abstract from the more ugly parts of our environment but which might nevertheless be adaptable. Examples can be found in Harris and Raviv (1981). Chao and Wilson (1987) identify institutions for implementing priority pricing when there are no indivisibilities. There are algorithms from combinatorial optimization, designed to attack informational difficulties in even more complex problems, that ignore incentive issues but that might be adapted to do so. Examples can be found in French (1982), Kirkpatrick, Gelatt, and Vecchi (1983), Reiter (1966), Reiter and Sherman (1962), and Rassenti, Smith, and Bulfin (1982). The adaptive user selection mechanism and the iterative Vickrey-Groves mechanism combine features of each of these approaches. Unanswered is whether there are other mixtures that are better.

## Appendix

■ The redemption value sheets follow.

Valuation Sheet 1

| X \ Y | 3 | 9 | 13 |
|---|---|---|---|
| 4 | 100 | 150 | 175 |
| 7 | 175 | 225 | 250 |
| 12 | 250 | 325 | 335 |

Valuation Sheet 2

| X \ Y | 6 | 10 | 14 |
|---|---|---|---|
| 3 | 125 | 150 | 175 |
| 9 | 175 | 190 | 200 |
| 15 | 200 | 225 | 250 |

[25] See Daughety and Forsythe (1987) for an initial attempt to use experimental methods to study this issue.

Valuation Sheet 3

| X \ Y | 2 | 4 | 9 |
|---|---|---|---|
| 3 | 75 | 100 | 125 |
| 5 | 100 | 200 | 225 |
| 12 | 175 | 250 | 275 |

Valuation Sheet 4

| X \ Y | 8 | 10 | 12 |
|---|---|---|---|
| 6 | 100 | 150 | 200 |
| 8 | 150 | 200 | 275 |
| 12 | 175 | 250 | 300 |

Valuation Sheet 5

| X \ Y | 7 | 10 | 13 |
|---|---|---|---|
| 6 | 175 | 225 | 250 |
| 9 | 225 | 275 | 300 |
| 12 | 250 | 300 | 325 |

Valuation Sheet 6

| X \ Y | 7 | 9 | 11 |
|---|---|---|---|
| 7 | 75 | 150 | 175 |
| 9 | 125 | 175 | 200 |
| 11 | 150 | 200 | 225 |

## References


BANKS, J.S., LEDYARD, J.O., AND PORTER, D. "Pricing, Evolution, and Design Planning of Space Station under Uncertainty." JPL Economic Research Series No. 22. Pasadena, Ca.: Jet Propulsion Laboratories, 1985.

CHAO, H. AND WILSON, R. "Priority Service: Pricing, Investment, and Market Organization." *American Economic Review,* Vol. 77 (1987), pp. 899–916.

COX, J.C., ROBERSON, B., AND SMITH, V.L. "Theory and Behavior of Single-Object Auctions" in V.L. Smith, ed., *Research in Experimental Economics,* Vol. 2, Greenwich, Ct.: JAI Press, 1982, pp. 1–44.

D'ASPREMONT, C. AND GERARD-VARET, L.A. "On Bayesian Incentive-Compatible Mechanisms" in J.-J. Laffont, ed., *Aggregation and Revelation of Preferences,* Amsterdam: North-Holland, 1979, pp. 269–288.

DAUGHETY, A. AND FORSYTHE, R. "Complete Information Outcomes without Common Knowledge." Working Paper No. 87-24, Department of Economics, University of Iowa, 1987.

DEBREU, G. *Theory of Value: An Axiomatic Analysis of Economic Equilibrium.* New York: John Wiley and Sons, Inc., 1959.

FEREJOHN, J., FORSYTHE, R., AND NOLL, R. "An Experimental Analysis of Decisionmaking Procedures for Discrete Public Goods: A Case Study of a Problem in Institutional Design" in V.L. Smith, ed., *Research in Experimental Economics,* Vol. 1, Greenwich, Ct.: JAI Press, 1979.

FOX, G. AND QUIRK, J. "Uncertainty and Input-Output Analysis." JPL Economic Research Series No. 23. Pasadena, Ca.: Jet Propulsion Laboratories, 1985.

FRENCH, S. *Sequencing and Scheduling: An Introduction to the Mathematics of the Job-Shop.* West Sussex, England: Ellis Horwood Limited, 1982.

GRETHER, D., ISAAC, M., AND PLOTT, C. "The Allocation of Landing Rights by Unanimity among Competitors." *American Economic Review,* Vol. 71 (1981), pp. 166–171.

GROVES, T. "Incentives in Teams." *Econometrica,* Vol. 41 (1973), pp. 617–631.

HARRIS, M. AND RAVIV, A. "A Theory of Monopoly Pricing Schemes with Demand Uncertainty." *American Economic Review,* Vol. 71 (1981), pp. 347–365.

HARSANYI, J.C. "Games with Incomplete Information Played by 'Bayesian' Players." *Management Science,* Vol. 14 (1967–1968), pp. 159–182, 320–334, 486–502.

———. "Bayesian Probability in Game Theory" in D.H. Mellor, ed., *Science Belief and Behavior: Essays in Honor of R.B. Braithwaite,* Cambridge: Cambridge University Press, 1980, pp. 192–196.

HAYEK, F. "The Use of Knowledge in Society." *American Economic Review,* Vol. 35 (1945), pp. 519–530.

JOHNSON, A., LEE, H.Y., AND PLOTT, C.R. "The Multiple Unit Double Auction Users' Manual." Working Paper No. 676, California Institute of Technology, 1988.

KIRKPATRICK, S., GELATT, C.D., JR. AND VECCHI, M.P. "Optimization by Simulated Annealing." *Science,* Vol. 220 (1983), pp. 671–680.

KOOPMANS, T.C. AND BECKMANN, M.J. "Assignment Problems and the Location of Economic Activities." *Econometrica,* Vol. 25 (1957), pp. 53–76.

LEDYARD, J.O. "The Economics of Space Station" in M.K. Macauley, ed., *Economics and Technology in U.S. Space Policy,* Washington, D.C.: Resources for the Future, 1987, pp. 127–170.

LEHMANN, E. *Nonparametrics: Statistical Methods Based on Ranks.* San Franciso: Holden Day, 1975.

MARSCHAK, T. "Computation in Organizations: The Comparison of Price Mechanisms and Other Adjustment Processes" in C.B. McGuire and R. Radner, eds., *Decision and Organization,* Amsterdam: North Holland, 1972.

MILGROM, P. AND WEBER, R. "A Theory of Auctions and Competitive Bidding." *Econometrica,* Vol. 50 (1982), pp. 1089–1122.

NATIONAL AERONAUTICS AND SPACE ADMINISTRATION. *Space Station Operations Task Force: Panel 3 Report.* Washington, D.C.: U.S. Government Printing Office, 1987.

PLOTT, C.R. "Industrial Organization Theory and Experimental Economics." *Journal of Economic Literature,* Vol. 20 (1982), pp. 1485–1528.

RASSENTI, S.J., SMITH, V.L., AND BULFIN, R.L. "A Combinatorial Auction Mechanism for Airport Time Slot Allocation." *Bell Journal of Economics,* Vol. 13 (1982), pp. 402–417.

REITER, S. "A System for Managing Job-Shop Production." *Journal of Business,* Vol. 39 (1966), pp. 371–393.

——— AND SHERMAN, G.R. "Allocating Indivisible Resources Affording External Economies or Diseconomies." *International Economic Review,* Vol. 3 (1962), pp. 108–135.

SMITH, V.L. "Experimental Economics: Induced Value Theory." *American Economic Review, Proceedings,* Vol. 66 (1976), pp. 274–279.

———. "Experiments with a Decentralized Mechanism for Public Goods Decision." *American Economic Review,* Vol. 70 (1980), pp. 584–600.

———. "Microeconomic Systems as an Experimental Science." *American Economic Review,* Vol. 72 (1982), pp. 923–955.

VICKREY, W. "Counter Speculation, Auctions, and Competitive Sealed Tenders." *Journal of Finance,* Vol. 16 (1961), pp. 8–37.