DIVISION OF THE HUMANITIES AND SOCIAL SCIENCES
CALIFORNIA INSTITUTE OF TECHNOLOGY
PASADENA, CALIFORNIA 91125

ALLOCATING UNCERTAIN AND UNRESPONSIVE RESOURCES

Jeffrey S. Banks
University of Rochester

John O. Ledyard
California Institute of Technology

David P. Porter
University of Arizona

SOCIAL SCIENCE WORKING PAPER 680

September 1988

ALLOCATING UNCERTAIN AND UNRESPONSIVE RESOURCES

Jeffrey S. Banks
University of Rochester

John O. Ledyard
California Institute of Technology

David P. Porter
University of Arizona

ABSTRACT

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 super-computers, zero-inventory planning, and the allocation and pricing of NASA's planned Space Station. 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 percent 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 that arise naturally from auctions used in the allocation of 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 both administrative and market procedures.

ALLOCATING UNCERTAIN AND UNRESPONSIVE RESOURCES[1]

Jeffrey S. Banks, John O. Ledyard and David P. Porter

I. INTRODUCTION

Short-run demand and supply imbalances are pervasive in markets and other organizations. The economic impact of these imbalances is exacerbated when the supply cannot be inventoried or quickly changed and when there are indivisibilities in demand so that it is impossible to react to the shocks. Indivisibilities imply that reactions by economic agents are not smooth, creating a coordination problem which becomes increasingly severe as the number of commodities increases. The combination of uncertainty, indivisibilities, and unresponsive supply creates an important economic allocation problem which 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 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 both through 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 the planned earth-orbiting Space Station of NASA. 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.[2] 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 which 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 the selection of users and the scheduling (manifesting) of 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.[3]

Other examples include super-computer 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 non-storable commodities. Some of these features are stronger in some examples than in others[4] 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. Job shop scheduling in manufacturing is generally organized the latter way. With administrative procedures, an apparently more centralized method is created to solve the coordination and timing problems. Gas pipe-lines are regulated.[5] Airports are managed through complex committees.[6] 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.[7] Every economist is trained to expect the inefficiencies, inherent in the centrally administered approach to allocation, which 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 non-convexities (optimal allocations cannot be supported as equilibria even with complete contingent contracts) which means 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 is not stable because feedback through prices and (random) rationing can be misleading. It may not even be possible to be good on average. Consider the following quote 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 quote retains its validity today and captures the spirit of our approach. In what follows we have focused on a "limited but well-defined problem" which is crudely simple but which we feel 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-dimension problems, which employ a form of computer assisted coordination to significantly improve on that performance. We feel that one of these mechanisms strongly merits further study.

II. A SIMPLE ENVIRONMENT

It is fairly easy to describe the general structure of the class of problems we are interested in. It is easier, however, to understand the practical difficulties created for economic organization, in the presence of multiple decision makers with differential information, through the use of 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 successfully complete their missions. To keep the example simple, suppose 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 each Shuttle provides resources $\bar{y} \in \mathbf{R}_+^K$ with a successful launch. [Included in K are available power and manpower on orbit as well as volume—in the Shuttle—and the mass that will be lifted.] The situation is summarized in Table 1.

A natural next modeling step would be to assign probabilities to the states and think of this as a random supply problem. But allocations cannot be made solely contingent on the state because at the time A launches on date 1 it is still unknown whether B will launch on date 2. It is necessary to explicitly consider the time structure in an event-tree. We have drawn this in Figure 1 where A = "A launches," B = "B launches," and N = "Launch Delayed". The states of the world as described in Table 1 remain as events 1, 2, 3, 4 and two additional events have been identified to recognize the effect of the timing of information. Allocations, supplies, and demands are specified contingent on events. Thus, supply is modeled as a function $y : \mathbb{E} \rightarrow \mathbf{R}_+^K$ where $\mathbb{E} = \{1, \ldots, 6\}$, $y(1) = y(3) = y(5) = \bar{y}$, and $y(2) = y(4) = y(6) = 0$.

We assume 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 will index designs by their resource requirements. Thus, $A^i \subseteq \mathbf{R}_+^K$. A (contingent) allocation to i is a vector $x^i = \langle x_1^i, \ldots, x_6^i \rangle$, where $x_e^i \in \mathbf{R}_+^K$ 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 will call money. Let

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

$$\text{where } \gamma^i(a^i, x_e^i, e) = \begin{cases} \eta(e) & \text{if } x_e^i \geq a^i \\ 0 & \text{otherwise} \end{cases}$$

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 $x_1^i \cdot x_5^i = 0$ 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 $x_1^i \cdot x_5^i = 0$, is locational in nature (see Koopmans and Beckman (1957)) since i can consume only in event 1 or 5 but not both. This creates a non-convexity in the consumption sets. Second, once a^i is chosen there is no flexibility and, thus, i demands only a single fixed amount in each market. This creates a threshold non-convexity in utility.

To evaluate the performance of alternative institutions in this environment, we must use a performance target to identify those (contingent) allocations which are considered to be desirable. Our choice for a target is the set of *ex ante* Pareto-optimal allocations.[8] That is, we consider the maximization of social value from a perspective prior to any realization of supply—at time 0, the top of Figure 1. Because all agents are assumed to 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, non-linear programming problem; choose a, β, x to

$$\text{maximize } \Sigma_{i=1}^{I} G^i(a^i) \Sigma_e \beta^i(e)\eta(e)$$

subject to

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

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

$$\beta^i(e) = 0 \text{ or } 1 \qquad \forall i, \forall e$$

$$x_e^i = \beta^i(e) a^i \qquad \forall i, \forall e$$

$$a^i \in A^i \qquad \forall i.$$

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

If a single agent possessed all the information they 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 prior to the realization of e. Economists now realize that, at the very least, this involves an incentive problem. For the environments we are most interested in we feel 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 their possibilities A^i or preferences G^i when there are multiple contingencies and multiple dimensions. One institution which economists normally expect to perform well under these conditions is markets. At worse the performance of markets provides a benchmark against which to measure other possibilities. As we will soon see, however, our simple example provides an apparently insurmountable challenge to at least one organized market institution.

III. 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 efficiently allocate resources if there is enough depth to avoid non-competitive behavior. The analysis is simple and standard. Following methods introduced in Debreu (1959, ch. 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 $P^* \in \mathbf{R}^{KE}$, 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

$$(a)\ \Sigma_{i=1}^{I} x^{*i}(e) \leq y(e) \quad \forall e \in \mathbb{E}$$

and

$$(b)\ \forall i \in I,\ a^{*i}, x^{*i}(1), \ldots, x^{*i}(E)$$

solves

$$\text{Max } \{\Sigma_e U^i [a^i, x^i(e), e]\, \eta(e)\} - \Sigma_e P_k^*(e) x^i(e).$$

A. The Analysis

It is easy to see that there are no externalities and that preferences satisfy local non-satiation (because "money" is infinitely divisible and always desired). Thus, the First Welfare Theorem [Debreu (1959, ch. 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 straight-forward. The environment has a unique optimal allocation. By the Second Welfare Theorem [Debreu (1959, ch. 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 non-convexities in the structure of our allocation problem, which prevent the application of this theorem. In fact, for most of the environments under consideration, the unique optimum cannot be supported as an equilibrium and, by the First Welfare Theorem, therefore there exists no equilibrium. Each of the two troublesome non-convexities can be traced to a specific type of indivisibility.

(1) There is a locational indivisibility. A payload can only be in shuttle A or B but not both, which creates a non-convexity 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 Beckman (1957) identified and which has created problems for urban–regional economics.

(2) There is a threshold indivisibility which 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 or not to participate, by buying a^i or by buying 0, at each event. This creates a non-convexity in the preferences of i (a non-concavity in the utility function).

Each non-convexity is potentially fatal in the sense that there are environments with each alone such that competitive equilibria do not exist. Both non-convexities together make non-existence virtually certain. But does this mean "markets won't work?" Not necessarily. For example, agents might use mixed strategies that "smooth out" the discontinuities created by the non-convexities and yield allocations that are good "on-average". One might not achieve 100 percent 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? Institutional features normally abstracted from in competitive analysis can no longer be ignored.[9] But development of the theory is not necessary for this paper. We tested the performance of markets in another way. We turn to that now.

B. The Experiment

To test the effectiveness of markets, we created an experimental environment for which 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,[10] which has performed at 95-100 percent efficiency levels in past applications (see Smith (1982) or Plott (1982)) and which has become the experimental standard market mechanism.

1. The Environment

The experimental environment is based on the example described in section II. 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 time period 1 are available (g) with probability ρ_1 and unavailable (n) with probability $(1-\rho_1)$. Either the total quantity is available or no quantity is available for consumption. For time period 2 the probability of g is ρ_2 and of n is $(1-\rho_2)$. ρ_2 is independent of the time period 1 outcomes. Table 2 shows the exact parameters used to represent the supply side of the experiments.

The demand side was created using monetary functions to induce value (see Smith (1976)). For subject $i=1,\ldots,n$ 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 $G^i(a^i)=M^i(a^i)$, the induced values are identified with the model in Section II.[11] In the experiments only discrete amounts were made available for the a choices.[12] In particular, each subject was given a 3×3 matrix of values corresponding to nine possible choices. If the mechanisms we are considering work well in this environment, they can easily be modified for operation in a more continuous demand structure. The actual valuation tables used in the experiments can be found in Appendix A. Subjects could only use the nine discrete choices available to them on the valuation sheets. We used six subjects per experiment.

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 sizable 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.

As another way to provide a feel for the design, we calculated 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 which 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. The expected value of $\Sigma_i M^i(a^i)$ was then calculated and the selection process started over again.

If all subjects are risk-neutral then an *ex ante* Pareto-optimal allocation of resources is one that maximizes the expected value of $\Sigma_i M^i(a^i)$. The allocation described in Table 3 is the unique expected value maximizing use of resources and design choices. 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*, and summarize the distribution of efficiency levels in Figure 2. Notice that the lowest possible efficiency level is in the 20–30 percent interval and that 85 percent of the distribution mass is between 40 percent and 75 percent. Hence, if we were to randomly allocate resources we should not be surprised to see efficiency levels in the range 40 percent–70 percent. Very few allocations, however, yield efficiency levels above 80 percent. Even though 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.

2. The Institution

To test the efficiency of markets in the context of the non-convexities we have created, we gave subjects the option of trading in any of six markets corresponding to the appropriate contingent commodities in non-zero supply. The markets are identified in Table 4. The capacity available in each market was divided among each of the six agents as an initial endowment which they could keep or sell. The specific initial allocation was chosen to minimize the efficiency level so 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. The standard Double Oral Auction institution was used[13] 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 we gave this institution its best chance to produce efficient allocations, we used subjects who were Caltech undergraduates and who were *experienced* in DOA experiments using the same PCs and software across 19 markets. Further all subjects were given the same payoff tables and endowments in each trial.

C. The Result

The experimental data, some of which is displayed in Figure 3 and Table 7, shows that this environment was too much for the Double Oral Auction. Although efficiency reaches 80 percent (but not until period 7), the data suggest failure. The average efficiency is 66.4 percent with a range of [43,83.1]. Early periods are the worst. There are several reasons for this performance, traceable to the non-convexities, but the primary one seems to be the inability of the agents to decide which market they want to be involved in. There is no clear evidence of monopoly (or any other strategic) behavior, only of the inability of the markets to provide clear and predictable signals to coordinate the activities of the agents and 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?

IV. 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 will concentrate on the one we feel is not an unreasonable abstraction of NASA's STS pricing and allocation policy prior to 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 $\{x^1, \ldots, x^I\}$ randomly without replacement. At each draw the agent receives a contract for resources x^i contingent on A 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 not under the control of either the agents or the administrator. We call this mechanism the Administrative Process.

A. The Analysis

The mechanism is not designed to achieve efficient allocations. To see why, from a theoretical perspective, let us consider the allocations which 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*. [14] 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 3, 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.[15] 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. However, even if they do, the efficiency performance of the mechanism is still only 83 percent of the maximum. This is because 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.[16] Presumably with inefficient designs, overall efficiency deteriorates even more.

B. 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 questions in a controlled experiment, we took the environment described in III B and now allocated resources 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 Table 1. This type of contract is called a *priority contract* (see Chao and Wilson (1987) or 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 Table 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 2 and were known by all participants prior to placing 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.

C. The Result

How does this Administrative Process perform? The answer seems to be, for the experimental environment we created, that this mechanism does no better and no worse than markets do. 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, non-payers were then allocated contracts. Although policy makers do not have the information required to post "correct" prices, we decided to give posted-pricing its best shot by posting those prices which would lead subjects 1, 2, and 3 to choose their (efficient) payload designs in Table 3 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 somewhat lower probability of inclusion was preferable to a costly ride with a higher probability of inclusion but a lower net return. In the end, posting prices had no effect. Periods 4 and 5 of AP/price are similar to periods 1–3 of AP. 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 environments with uncertainty, unresponsive supply, and indivisibilities in demand. This leaves open the question, can a designed institution do better than the naturally evolved institutions have done?

V. 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 2nd price (or 1st 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 their 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 correctly reveal 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 which 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 modeled 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 budget-balance problem. After all, current policy which 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.[17]

We did not feel comfortable ignoring the second characteristic which, although not usually analyzed theoretically, is more critical for many applications. Vickrey–Groves mechanisms require each bidder to report an *entire payoff function*, thus rendering the informational tractability of such a mechanism problematic in most applications with environments characterized by multiple units and multiple dimensions. 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 the benefits from any particular configuration of resources—given enough time, they usually know best only those configurations in a neighborhood of what they expect to receive.[18] To confront these informational realities, we decided to modify the standard Vickrey–Groves mechanism to require buyers to specify only one configuration and a willingness-to pay per report, a single demand point. Our intuition, based on others' research with iterative mechanisms,[19] was that if prices at each iteration were consistent with the Vickrey–Groves logic, then with sufficient iterations information will be generated which would guide users to the efficient allocations in an "incentive-compatible" manner.

A. The Analysis

The new process which we call the Iterative Vickrey–Groves process (IVG) proceeds as a communication tatonnement 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 $\mathbb{F}$ is considered 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

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

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

$$\gamma_i(f) = \underset{\substack{\gamma \in \Gamma(f) \\ i \notin \gamma}}{argmax} \sum_{j \in \gamma} b^{jf} \qquad st. \ \sum_j d^j + d^i \le \bar{y},$$

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 $\Sigma_{f \in \mathbb{F}} \beta^i(f)$ if $i \in K_f$ and 0 otherwise. To understand what β^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 $\Gamma(f)$ which 1) would remain feasible if i were added, and 2) 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 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 holding other traders' bids constant. In the former, the first term on the *RHS* of the price equation is the amount generated if i did not participate; subtracting off the bids by other members of K_f gives an equivalent version of the Vickrey–Groves "second price" auction of a single unit of a good. Given d and assuming risk neutrality, bidding one's expected value for a contract is a dominant strategy. In the latter, the first term is simply the sum of bids of the members of K_f, thus subtracting off the bids of i 's "best chance" coalition gives the amount i would have had to have bid to have been 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 not only the demand for resources and contracts that they (provisionally) acquire, but also for those which 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 $t-1, \forall f \in I\!F$. That is, the process stops when some sort of stability has been reached wherein 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 value 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 designs might occur. Nevertheless it was our expectation that this process would improve on the Administrative Process, described in Section IV, since the Iterative Vickrey–Groves process at least selects efficient combinations of configurations even if those configurations may not be fully optimal.

B. The Experiment

The determination of allocations and the calculation of individual prices for this mechanism is an enormous task which cannot be done by hand in an effective manner. Thus, for testing this mechanism, all communication and calculations were made using a network of personal computers (PCs). The PCs were connected on a local area network with a controller PC being the center where messages were received, prices were calculated, allocations were determined, and from which messages were sent.

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 markets). 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, based on the trial messages. In addition, each subject received a private price message which described their potential payment if they were part of the provisional allocation, or (if they were not) the amount they would have had to bid in order to have had their 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 their 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, the a^i must be one of i's nine choices, and 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 "easily" search for combinations.[20] The instructions for this experiment can be found in Appendix B.

C. The Result

The Iterative Vickrey–Groves mechanism yielded higher efficiency levels than either Markets or the Administrative Process. The IVG attained a mean efficiency of 78 percent with a range of [60, 91], a significant increase from 64 percent. Particularly of interest, was the fact that IVG dominated both Markets and the Administrative Process on a period by period basis: IVG attained high efficiencies early and maintained them.

Unfortunately, an efficiency of 100 percent was not achieved. In fact IVG did not even produce stable levels of efficiency, dipping as low as 60 percent (easily obtainable at random). One explanation for this instability is the fact that the process was pushed to the final trial in many cases which led the final bid for d^i to be somewhat random. Coordination across user designs is lacking. Even though one trial produced an efficiency level of 90 percent, the reliability of the process is suspicious.[21] More details of performance are provided below in Section VII.

VI. 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 designed a mechanism by modifying the English (or ascending-bid) auction commonly used to sell art objects, livestock, and tobacco. (See Cox, Roberson and Smith (1982) or Milgrom and Weber (1982) for descriptions and analyses of ascending bid auctions). We call our version the *Adaptive User Selection Mechanism* (*AUSM*)-*Bulletin Board*. AUSM 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); 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 which guides coordination in design and which selects high yield projects.

A. The Analysis

The English auction, upon which AUSM is based, is a non-tatonnement process that is commonly and 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* which 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 their 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 $\bar{y} \in \mathbf{R}^k_+$. [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 $\bar{y} \in \mathbf{R}^2_+$. 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 $\Sigma_{j \in R} b^j$.

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 2, on the other hand, wanted $(x, y) = (12, 6)$, 2 would have to bid at least 201 (to bump 3). If 2 wanted $(x, y) = (12, 11)$, 2 would have to bid at least 501 (to bump 7) and if 2 wanted (6,16), 2 must bid 701 (to bump both 3 and 7).

We chose this basic mechanism for several reasons: (1) the practical success of the single unit English auction as signaled 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,[22] and (3) a theoretical analysis of its properties. Let us briefly expand on the last.

We emphasize two facts about the AUSM-Bulletin Board. First, given a proposed allocation any 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 in order to achieve any desired allocation on contract f. Let ψ^* represent a set of contracts; a potential allocation. Let $\xi^i(\psi^*)$ represent the set of contract allocations to which i can unilaterally cause ψ^* to be changed with some bid. We call ψ^* a *simple equilibrium* if ψ^* is feasible and if $\forall i = 1, \ldots, n$, $\xi^i(\psi^*) \cap \{\psi \mid V^i(\psi) > V^i(\psi^*)\} = \emptyset$, where $V^i(\psi)$ represents the (expected) utility i receives if the potential allocation ψ 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 which are individually "stationary" allocations of AUSM. This is a fairly big 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 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 max $\sum_{i \in K_f} b^i$ 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 ψ^* a *non-cooperative 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 an obvious fact that if ψ^* is a non-cooperative equilibrium allocation of G, then ψ^* is a simple equilibrium of AUSM. The converse is not necessarily true.

Based on 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 non-cooperative 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 non-cooperative equilibria of AUSM occur, the associated allocations may not be desirable. Because of the lumpy nature of the users' projects, there are non-cooperative equilibrium allocations of G which are not efficient contract allocations. There may be changes in those allocations involving several traders *simultaneously* which 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 large user. In this situation unilateral actions by one user are not sufficient to drive the mechanism to a more efficient allocation of contracts.[23]

In our initial testing of the AUSM-Bulletin Board, we had hoped that these complications caused by the variable size demands would be overcome by the subjects. But early data (reported in detail below) suggested efficiency levels of only 75-85 percent. We 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 big user with two or more little users. Our solution was not only to allow small users to coordinate their bids but to encourage them to do so. We created a new mechanism by modifying AUSM in the following ways. A public "standby" queue was allowed in which any agent could post a "proposed bid" (d^i, b^i, f) which they would be willing to have included in a coalitional bid. Because of the possibility of joint bids from a group γ of individual agents, we expected different outcomes with the queue than we had hypothesized would arise in AUSM without the queue.

Let $\gamma \subseteq \{1, \ldots, n\}$ be an arbitrary coalition of agents and let (m^* / m^γ) be the vector m^* with m_i^* replaced by m_i for all $i \in \gamma$. We call ψ^* a *strong non-cooperative equilibrium* of G if $\psi^* = \psi^*(m^*)$ and for each coalition γ 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 ψ^* is a strong non-cooperative equilibrium of G, then ψ^* is a non-cooperative equilibrium of G. The converse is not necessarily true.

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

B. The Experiment

As before, we created 2 markets (priority contracts) with 1 corresponding to A in Figure 1 and 2 corresponding to B. When the market 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 canceled. Finally, to aid in the search process for the best configurations, subjects were allowed to move existing configurations to other markets and/or change their configuration and bid if it could fit in the available capacity. However, if a subject did change his configuration in a market 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 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 or she 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. Based on the data in Table 2, 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 the subject received zero earnings for the market period. At the beginning of the experiment each subject was given 7 dollars of working capital to add to his earnings since losses in any market period were possible.[24]

The AUSM mechanism without a queue was also tested. The instructions for both AUSM mechanisms can be found in Appendix B.

C. 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 percent with a range from 72 percent to 86 percent. AUSM with queue dominated all others each day. Two other interesting facts were (1) the small variance in the levels of efficiency achieved and (2) the high revenue collected.

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

We now turn to a more detailed analysis of the performance of each mechanism in the experimental test-bed. We then conclude with some final observations.

VII. 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. This does not necessarily measure the *ex ante* efficiency of the mechanisms since it does not account for the risk preferences of the subjects; nevertheless, it does measure the *ex ante* system performance from the point of view of a risk-neutral planner. We also consider the extent to which revenue is generated by each mechanism. Finally, we evaluate the data to see whether we can learn anything about individual behavior.

Table 6 contains relevant information describing each experimental session.

A. Efficiency

The mean efficiency (percent of the maximum expected value (μ)) and the associated standard deviation (σ) and coefficient of variation (v) for each mechanism can be found in Table 7. The nature of the underlying distribution of combinations found in Figure 2 suggests the use of nonparametric methods for our statistical analysis. We will use the Wilcoxon Rank Sum to test the equality of distributions of efficiency generated by each mechanism. In particular the z scores to be reported are derived from testing the hypothesis of equality of distributions versus strict inequality of distributions. The Wilcoxon Rank Sum Test for each mechanism are in Table 8. 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 8 support the following ranking of the mechanisms by efficiency as follows:

$$M = AP = A < IG = AUSM < AUSMQ$$

The mean efficiency per period for each mechanism can be found in Figure 3. We notice that for the Administrative mechanism and markets the level of efficiency tends to increase over 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 AUSM with and without a queue and for IVG, there is no significant effect of repetition. The mean efficiency for periods 1 and 2 was 77.1 for AUSM, 81.6 for AUSMQ, and 76.8 for IVG, while the mean efficiency for periods 3 and above was 78.1 for AUSM, 80.2 for AUSMQ, and 78.6 for IVG. 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 percent confidence intervals for each market/period are in Table 15. Three observations can be made: (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 that of markets 3 and 4 (1/3 probability). Although this last observation is curious, notice that markets 1 and 2 combined with 5 and 6 makeup 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 non-convexities.

B. Revenue

AUSM with a queue is more efficient than AUSM without a queue. Counteracting this advantage of the queue 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 and coefficient of variation are shown in Tables 9 and 10. The histogram of the overall revenue generated from the processes is in Figure 4. Very little revenue is generated in AP since most of the requests (77 percent) were for non-paying status. From Table 9 we see that 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 1/3 that of market 1 contracts (approximately the difference in probability of each market being filled). Table 11 supplies the rank sum and *t* tests for the overall revenue generated by each of the AUSM and IVG treatments, while Table 12 provides these same tests for each of the markets.

We see that the existence of the standby queue results in significantly higher revenues and this comes from higher revenue generated in both markets 1 and 2. The mean for periods 1 and 2 and periods 3+ for each mechanism and the relevant statistical tests are provided in Tables 13 and 14. We notice a slight upward trend in the revenue in both AUSM and IVG which is traceable to the revenue generated in market 1. No such trend is found by examining the time series for AUSM with a queue. We also notice that the revenue generated by IVG is quite volatile.

C. Individual Choice Behavior

Administrative Process. As detailed in Section IV, we would expect the Administrative Process to generate choices which are consistent with "scaled down" projects. Specifically, from the redemption value sheets in Appendix A, notice that if each individual chose the largest project only one order would fit per market. However, if each subject chose his smallest project everyone could fit in one of the markets. Actual orders balanced between these extremes. Out of the 102 orders filled in the Administrative Process, only six orders submitted were an individual's smallest project while only nine orders had the largest project submitted. The average sizes (X, Y) submitted by subjects over time are in Table 16. This table provides evidence of an updating phenomena on the part of the subjects. As a final note, *the ranking of markets by probability* (see Section V) *was never violated by a subject in the experiments.*

AUSM. While we did not design the experiments to test whether the three equilibrium concepts used in Section VI A were consistent with reality, some evidence can be extracted from the data. In a strict sense, all three are rejected. None of the 20 allocations achieved by AUSM were even simple equilibria much less non-cooperative or strong non-cooperative. Only 4 of the 20 allocations chosen by AUSM with queue were simple equilibria. Some of this might be explained by risk–averse behavior, but we feel that risk is not that important in this environment.

A better insight can be gained 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 percent (15/56) of the non-best responses come from subjects in market 2 for AUSM and 33 percent (11/33) in AUSM with queue. These facts suggest subjects may be "standing pat" in market 2. Most, almost 70 percent, of the non-best responses are from subjects who are not yet allocated space in any market. These data lead us to conclude, albeit tentatively, that non-cooperative equilibrium allocations of the full-information game are likely to occur in AUSM, especially if the stakes were to be increased. Over 80 percent 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, non-cooperative equilibria.

Iterative Vickrey–Groves. We were especially interested in two aspects of individual behavior in the IVG 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 percent of the bids) and this was even true on the last trial in a period (31 percent of the bids). This behavior did lead some individuals employing those bids to pay more than their value to get a contract. Those making such a loss (or after obtaining prices above expected values) generally did not repeat this bidding behavior.

Two particular aspects of strategic behavior employed 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, leaving resources to be allocated in a one-shot game in the final trial, (85 percent of the time for market 1 and 75 percent of the time for market 2). In order to see if 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 percent and revenue was 250; the second period concluded with 78 percent efficiency and 495 for revenue. Both periods used all 40 trials.

Markets. Even though there were only six participants in the experiment trading in the markets was very active.[25] However, the trading patterns over the course of an experiment were not entirely stable. That is, individuals *did not* choose the same projects and use the same contracts in each period. We call this phenomena *market switching.* As evidence of this market switching we noticed very few subjects (28 percent) stayed with the same contract across periods; 60 percent of the subjects that did selected no contracts (they sold their endowments). However, 17 percent of the subjects switched markets at least once and 56 percent switched at least twice in the last four periods. This evidence is consistent with the fact that there was no stable set of prices to coordinate allocations in our market experiments.

VIII. CONCLUSION

In this paper 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 which have naturally evolved, and two newly designed mechanisms, the Iterative Vickrey–Groves mechanism and AUSM.

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*, non-convexities. In the experimental test-bed, markets performed no better than an *ad hoc* administrative process. Both mechanisms, modeled on naturally evolved institutions, yielded low efficiency levels.

Two new designed institutions AUSM and IVG, using priority contracts were created and analyzed as alternatives to markets and administrative processes. Both significantly outperformed markets in the testbed, especially in the early periods when there is no information and no basis for common knowledge priors. This is an especially important consideration for non-repeatable operations such as the Space Station. Neither could consistently generate efficiencies greater than 85 percent in the demanding experimental environment.

Designing mechanisms for applications, as in designing America's Cup sailboats or airplanes like the Shuttle, is an artform to be guided by theoretical developments and experience. Our use of experimental methods has allowed us to accumulate some experience prior to the application. Several observations based on that experience deserve mention. 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 causing inefficient resource purchases. The Administrative Process is simple-minded but there was so much randomness subjects had difficulty coordinating designs. The IVG process is "obscure"; subjects were unsure what "prices" would be. Only AUSM was straight-forward and stable. Second, while tatonnement-like iterations without commitment may guide coordination (it is a form of "cheap-talk"), our experience shows that some commitment is desirable 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, knowing 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 IVG). AUSM works because it attains such a balance. Third, markets cannot smooth over all non-convexities, 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.[26]

This study is part of a larger study to help NASA develop appropriate mechanisms to allocate its Space Station resources among a variety of diverse users. Indeed, the results contained in this study have already lead to the development of a vastly more elaborate experimental environment (testbed) to demonstrate the feasibility and performance of various allocation mechanisms[27] (see Plott and Porter (1988)). There are other modifications of the environment that would bring the test-bed into closer congruence with the Space Station problem. For example it would increase realism by inducing the preferences of some or all of the users to be a function of the *time* at which they are allocated resources. This would model better the importance of phenomena like "launch windows" in the timing in Figure 1. Adding a time dimension to user preferences would also highlight the need for an analysis of contingent and futures contracting in allocating resources. Being first would no longer necessarily be better. This modification would allow a meaningful comparison of mechanisms across contract types, as well as contract types across mechanisms. Another change in the environment would alter the subjects' problems to correspond more closely to that of a NASA engineer and less that of a profit-maximizing firm. In particular, the payload selection process (including budgeting, peer review, and project management) might be better modeled with budget constraints on designs and payoffs for scientific benefits rather than with net benefits as we have done in this paper.[28] A fundamental change would be to allow supply to be endogenous. This significantly alters the structure of the problem and may raise problems which neither AUSM nor IVG can handle effectively.[29]

The fundamental open question is, of course, whether or not there exist other allocation mechanisms which 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–Raviv (1981). Chao and Wilson (1987, p. 914) identify institutions for implementing priority pricing when there are no indivisibilities. See also Reitman (1985) and Pitblado (1987). There are algorithms from combinatorial optimization, designed to attack informational difficulties in even more complex problems, which ignore incentive issues but which might be adaptable in any case. Examples can be found in French (1982) Kirkpatrick, Gelatt and Vecchi (1983) Reiter (1966) Reiter and Sherman (1962) and Ressenti, Smith and Bulfin (1982). AUSM and IVG combine features of each of these approaches. Unanswered is whether there are other mixtures which are better.

APPENDIX A

REDEMPTION VALUE SHEETS AND SCREEN DISPLAY

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

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

APPENDIX B

INSTRUCTIONS

INSTRUCTIONS

(This portion is the same for all mechanisms)

You are about to participate in an experiment designed to provide insight into certain features of decision processes. If you follow the instructions carefully and make good decisions, you might earn a considerable amount of money. You will be paid in cash.

In this experiment we are going to conduct a market in which you will make decisions which will be used to determine the market outcomes. You will be given a Redemption Value Sheet, which describes the value to you of the decisions you might make. You are not to reveal this information to anyone. It is your own private information.

The type of currency used in this market is francs. All transactions will be in terms of francs. Each franc is worth _____ dollars to you. Do not reveal this number to anyone. At the end of the experiment your francs will be converted into dollars at this rate, and you will be paid in dollars.

On your Redemption Value Sheet you have one project which has 9 possible X and Y configurations associated with it along with a redemption value stated in francs. Suppose for example that your Redemption Value Sheet were as follows:

X \ Y	3	6	12
5	100	200	300
10	200	400	500
15	300	450	550

Then for your project with the configuration X=5 and Y=12, you would have a redemption value of 300 francs; for your project with configuration X=10 and Y=6 you would have a redemption value of 400 francs.

Within each market period there will be a total of ___ markets with a fixed capacity of 20 units of X and 20 units of Y in each market to be allocated to participants. Your amount of X and Y and the earnings you will receive will be determined using the following process.

(For the Iterative Groves mechanism we used)

All communication during the market period will conducted through your computer terminal. The experiment will consist of several market periods. Each market period will be composed of trials in which you will submit an order. An order consists of a configuration, a market, and a bid for the market. You can submit an order during a trial by following the instruction prompts on your computer screen. The first prompt will ask you for a configuration of X and Y. You must then enter one of your 9 choices. Next you will be asked if you want to enter a bid for market 1; if you answer yes (y), you will be asked for your bid in francs. If you answer no (n), it will proceed to market 2 and ask you if you would like to bid. Thus, for a trial in a market period <u>you cannot simultaneously have a bid in markets 1 and 2</u>.

<u>EXAMPLE</u>

Enter quantity of X desired: <u>7</u>
Enter quantity of Y desired: <u>12</u>
Do you want to order in market 1? <u>n</u>
Do you want to order in market 2? <u>y</u>
Enter bid in market 2: <u>150</u>
Do you confirm X = 7 Y = 12 B1 = 0 B2 = 150 ? <u>y</u>

Thus, in this example a bid of 150 francs was placed in market 2 for the configuration X=7 and Y=12. The only restriction you have on the bids you submit for a market during a period is that it be greater than or equal to zero. After each participant has placed an order during a trial, a set of <u>provisional</u> configurations will be selected for each market by finding the largest sum of bids submitted for that market for which the sum of the corresponding configurations do not exceed the capacity constraints (X = 20, Y = 20). Each individual will then be given the information as to which participants are provisionally selected in each market and their configurations. In addition, each participant will receive a price for each market. If you are one of the participants that are provisionally selected in a market, your price will be calculated as follows:

[Maximum of the sum of bids (without your bid) submitted for that market for which the corresponding total configurations (without your configuration) does not violate the capacity constraint] -

[Sum of bids (excluding your bid) of the provisional configurations]

If your configuration is not one of the provisional configurations in a market, then you will receive a price which indicates the minimum bid you could have submitted in that market and have had your order be one of the provisional configurations in that market.

<u>EXAMPLE</u>

	Market 2			
Participant	X	Y	Bid	Price
*1	10	10	200	160
2	15	10	160	200
3	10	15	170	350
*4	5	10	150	0
6	11	5	100	200

In this example, we see that participants 1 and 4 have provisional configurations in this market because their combined requests of X=15 and Y=20 do not violate the capacity constraints, and their sum of bids of 350 francs is the largest such sum. The price for participant 1 was calculated by: [310] - [150] = 160 as per the equation above. Similarly, the price for participant 4 was calculated by: [200] - [200] = 0. Notice that in either case the price does not depend on the participant's bid.

For an individual who is not in the provisional configuration in this example, such as participant 3, we see that to be in the provisional set he/she must bid 350 francs since the configuration he/she submitted cannot fit with any other order.

The provisional configurations in the markets will obtain their allocations if the same participants and configurations occur for ___ straight trials (Rule A), or the trials in the market period are exhausted. A total of ___ trials for each market period will be allowed. Market 1 will close after ___ trials, and market 2 after ___ additional trials. That is, if Rule A is not activated after ___ trials the provisional allocations in market 1 will be chosen, and after ___ trials the orders in market 2 will be chosen. After the process stops in a period by either of the conditions a die will be rolled. If the numbers _____ through ____ appear the orders in market 1 will be filled. If the numbers _____ through ____ appear the orders in market 2 will be filled.

Each participant has been given _____ francs of working capital. To determine your costs for the market period, sum up your prices in those markets for which you have obtained an allocation at the market close. If your order is filled, then your earnings will be equal to your redemption value minus your costs. If your order is not filled, you must subtract your costs from your working capital. If you do not obtain an allocation in a market for the period, you will receive zero earnings for the period. You should record your earnings on your Record of Earnings sheet located in the back of your folder. At the beginning of a market period you will be assigned a new Redemption Value Sheet from which to make your decisions. The Redemption Value Sheet will not be the same for all participants. Feel free to earn as much as you can. Are there any questions?

(For the ASUM-Bulletin Board with Queue we used)

When the market opens you will be able to submit an order consisting of a market or the Standby Queue, a configuration, and a bid in francs. Orders will be taken one at a time and posted on the board. You can submit an order by raising your hand and after you are identified, you can submit one order. Your order will be accepted if:

a) It can fit in the available capacity of the market requested, or

b) It can displace existing orders with lower bids, or

c) The Standby Queue is requested.

If place an order in the Standby Queue you must also identify the market for which the order is to be placed on standby. However, you can have only one order in the markets at any one time. Thus, you can have an order in Market 1 or 2 but not both Markets 1 and 2 simultaneously. You can have as many orders as you want in the Standby Queue.

Suppose for example, that the fixed capacity was X = 20 and Y = 20, and there were 2 markets, and the existing orders , none of which are yours, were as follows:

	Market 1			Market 2			Standby Queue			
	X	Y	Bid	X	Y	Bid	X	Y	Bid	Market
	11	9	1100	7	5	500	5	5	400	1
	5	7	1000							
Available Capacity	4	4		13	15					

If you want to submit an order that has quantities X = 4 and Y = 6, you can order space in Market 2 or the Standby Queue and submit any nonnegative bid, or you can order space in Market 1 and displace the X = 5 and Y = 7 order with a bid greater than 1000 francs. Furthermore you can combine your bids with orders in the Standby Queue that were not submitted by you to displace existing orders in Markets 1 and 2 if the entire order can fit and the total bid is greater than the total displaced orders bids. For example, you could have made a bid greater than 600 francs and combined that with the existing 400 franc bid in the Standby Queue and displace the X = 5 and Y = 7 order in Market 1 since you both can fit. In the event that more than one existing order can be displaced by your bid the order with the lowest bid will be the one displaced. <u>If one of your orders in the Standby Queue is combined with another order, then any order you have standing in a market is canceled.</u>

If you have an order in a market you can change it only if you increase your bid. If you increase your bid you can:

a) Move your configuration to another market if you can fit or displace orders with lower bids and/or

b) Change your configuration if it fits. However, if you do not move your configuration to another market you must place a bid higher than the orders you are displacing including your Standing Order.

Your bid change must be greater than _______ francs to be accepted. Once you have an order in a market you cannot withdraw it. However, you can withdraw orders from the standby queue.

In the event that your order is displaced you can reorder through the process described above. The process will stop when there are no new orders or order changes (increased bids) within _______ seconds of the last order submitted. The orders left standing on the board in Markets 1 and 2 when the process stops are the only orders that can be filled. However, there is a chance that the orders at the market close will not be filled. After the process stops a die will be rolled. The orders in Market 1 will be filled if any of the numbers _____ through ______ appear. The orders in Market 2 will be filled if any of the numbers ______ through ______ appear.

Each participant has been given _______ francs of working capital. If your order is filled your earnings will be equal to your redemption value minus your bid. If your order is not accepted you must subtract your bid from your working capital. If you did not get in a market you will receive zero earnings for the market period. You should record your earnings on your Record of Earnings Sheet located in the back of your folder. Your earnings plus your remaining working capital are yours to keep. At the beginning of a market period you will be assigned a new Redemption Value Sheet from which to make your decisions. The Redemption Value Sheet will not be the same for all participants. Feel free to earn as much as you can. Are there any questions?

(For the AUSM-Bulletin Board we used)

When the market opens you will be able to submit an order consisting of a market, a configuration, and a bid in francs. Orders will be taken one at a time and posted on the board. You can submit an order by raising your hand and after you are identified, you can submit one order. Your order will be accepted if:

a) It can fit in the available capacity of the market requested, or

b) It can displace existing orders with lower bids.

However, you can only have one order standing on the board at any one time. Thus you can have an order in Market 1 or 2 but both Markets 1 and 2 simultaneously.

For example, suppose the fixed capacity was X = 20 and Y = 20, and there were two markets, and the existing orders none of which are yours were as follows:

	Market 1			Market 2		
	X	Y	Bid	X	Y	Bid
	10	9	1100	7	5	500
	5	7	1000			
Available Capacity	4	4		13	15	

If you want to submit an order that has quantities X = 4 and Y = 6 you can either order space in Market 2 and submit any nonnegative bid, or you can order space in Market 1 and displace the X = 5, Y = 7 order with a bid greater than 1000 francs. In the event that more than one existing order can be displaced by your bid the order with the lowest bid will be the one displaced.

If you have an order standing on the board you can change it only if you increase your bid. If you increase your bid you can:

a) Move your configuration to another market if you can fit and/or

b) Change your configuration if it fits. However, if you do not move your configuration to another market you must place a bid higher than than the orders you are displacing including your standing order. Your bid change must be greater than ______ francs to be accepted. Once you have an order in a market you cannot withdraw it.

In the event that your order is displaced you can reorder through the process described above. The process will stop when there are no new orders or order changes (increased bids) within _______ seconds of the last order submitted. The orders left standing on the board in Markets 1 and 2 when the process stops are the only orders that can be filled. However, there is a chance that the orders at the market close will not be filled. After the process stops a die will be rolled. The orders in Market 1 will be filled if any of the numbers _____ through ______ appear. The orders in Market 2 will be filled if any of the numbers ______ through ______ appear.

Each participant has been given _______ francs of working capital. If your order is filled your earnings will be equal to your redemption value minus your bid. If your order is not accepted you must subtract your bid from your working capital. If you did not get in a market you will receive zero earnings for the market period. You should record your earnings on your

Record of Earnings Sheet located in the back of your folder. Your earnings plus your remaining working capital are yours to keep.

At the beginning of a market period you will be assigned a new Redemption Value Sheet from which to make your decisions. The redemption value sheet will not be the same for all participants. Feel free to earn as much as you can. Are there any questions?

(For the Random Process we used)

At the beginning of the market period you will send in your order for an X and Y configuration you would like by submitting an order form, and a ranking of your preferences for Markets 1 and 2. That is, you cannot place a preference ranking for Markets 1 and 2 simultaneously. Order forms can be found in the back of your folder. To submit an order just place your X and Y configuration found on your Redemption Value Sheet on your order form with a ranking of the markets. For example, suppose you want to place an order for a configuration on your redemption value sheet that quantities of X = 9 and Y = 14, and you wanted the market rankings 2 and 1, then you would send in an order form with the following information:

ORDER FORM

		Ranking
X = ________	Y = ________	________
		________

You can submit one order.

After all the orders have been collected we will randomly select the orders and place them in the first available market with capacity available according to the ranking on the order form as they are drawn. After we have exhausted all the orders or the capacity in each market we will determine the orders that are filled by rolling a die twice. If the numbers ______ through ______ appear on the first roll the orders in Market 1 will be filled. If the numbers _____ through _____ appear then the orders in market 2 will be filled.

Your earnings for the market period will be equal to your redemption value if your order is filled, otherwise your earnings will be zero. You should enter earnings for the market period on your Record of Earnings Sheet. Your total earnings over all the market periods are yours to keep.

At the beginning of a market period you will be assigned a new Redemption Value Sheet from which to make your decisions. The redemption value sheet will not be the same for all participants. Feel free to earn as much as you can. Are there any questions?

(For the Random Process with Contingent Contracts we used)

At the beginning of the market period you will send in your order for an X and Y configuration you would like by submitting an order form, and a ranking of your preferences for Markets 1, 2, 3, 1 and 3, and 2 and 3. That is, you cannot place a preference ranking for Markets 1 and 2 simultaneously. Order forms can be found in the back of your folder. To submit an order just place your X and Y configuration found on your Redemption Value Sheet on your order form with a ranking of the markets. For example, suppose you want to place an order for a configuration on your redemption value sheet that quantities of X = 9 and Y = 14, and you wanted the market rankings 2 and 3, 1, 3, 1 and 3, and 2, then you would send in an order form with the following information:

ORDER FORM

X = __________ Y = ___________

Ranking

____
____
____
____
____

You can submit one order.

After all the orders have been collected we will randomly select the orders and place them in the first available market with capacity available according to the ranking on the order form as they are drawn. After we have exhausted all the orders or the capacity in each market we will determine the orders that are filled by rolling a die twice. If the numbers ______ through ______ appear on the first roll the orders in Market 1 will be filled. If the numbers _____ through _____ appear on the first roll and the numbers _____ through _____ appear on the second roll the orders in Markets 1 and 2 will be filled. If the numbers ______ through _____ appear on the first roll and the numbers ______ through _____ appear on the second roll the orders in Market 3 will be filled.

Your earnings for the market period will be equal to your redemption value if your order is filled, otherwise your earnings will be zero. You should enter earnings for the market period on your Record of Earnings Sheet. Your total earnings over all the market periods are yours to keep.

At the beginning of a market period you will be assigned a new Redemption Value Sheet from which to make your decisions. The redemption value sheet will not be the same for all participants. Feel free to earn as much as you can. Are there any questions?

TABLE 1

State	Description	Supply
1	A launches, B launches	$\bar{y} + \bar{y}$
2	A launches, B delayed	$\bar{y}$
3	A delayed, A launches	$\bar{y}$
4	A delayed, A delayed	0

TABLE 2

Time period	Quantity of X	Quantity of Y	Probability of Availibility
1	20	20	$\rho_1 = 2/3$
2	20	20	$\rho_2 = 1/2$

TABLE 3

	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 4

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

TABLE 5

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	

TABLE 6
Summary of Experiments

Mechanism	Number of Experiments	Total Number of Periods	Subject Pool	Notes
Administrative	5	17	Caltech, PCC*	Rotated Sheets
AUSM	4	20	Caltech, PCC	Priority Contracts Rotated Sheets
AUSM with Queue	4	20	Caltech	Priority Contracts Rotated Sheets
Iterative Vickrey–Groves	4	20	Caltech	Priority Contracts Rotated Sheets
Markets	3	24	Caltech	Complete Contingent Markets, Same Sheets, Subjects Experienced in Multiple Markets
Administrative with Prices	3	13	Caltech	Same Sheets

* PCC = Pasadena City College

TABLE 7
Efficiency by Mechanism

Mechanism	μ	σ	ν	Range
Administrative (A)	63.5	10.0	.16	[39, 76]
Markets (M)	66.4	10.2	.15	[43, 83]
Administrative with Prices >0 (AP)	69.9	8.1	.12	[58, 81]
Iterative Vickrey–Groves (IVG)	77.9	6.8	.09	[60, 91]
AUSM	77.7	4.1	.05	[71, 86]
AUSM with Queue	80.8	4.0	.05	[72, 86]

TABLE 8*
Rank Sum Test (All Periods)

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

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

TABLE 9
Total Revenue Generated

Treatment	μ	σ	ν	Range
AUSM	404.5	48.7	.12	[284, 475]
AUSMQ	475.7	52.0	.11	[375, 560]
IVG	388.4	118.2	.30	[210, 656]

TABLE 10
Revenue by Market

Treatment	Market 1				Market 2			
	μ	σ	ν	Range	μ	σ	ν	Range
AUSM	302.2	52.9	.18	[154, 365]	102.3	22.6	.22	[70, 145]
AUSMQ	353.7	36.2	.10	[300, 425]	122.0	27.5	.23	[75, 185]
IVG	284.1	85.1	.29	[160, 449]	108.1	66.3	.61	[00, 240]

TABLE 11*
Rank Sum and t-Test for Overall Revenue Generated

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

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

TABLE 12*
Rank Sum and t-Test in Priority Markets

	Market 1 (*AUSMQ*)		Market 2 (*AUSMQ*)	
AUSM	$z = 2.25$ $\alpha = .01$	$t = 2.48$ $\alpha = .01$	$z = 3.69$ $\alpha = .00$	$t = 4.45$ $\alpha = .00$

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

TABLE 13
Mean Revenue

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

TABLE 14*
t-Tests for Mean Revenue - Early vs Later Periods

	AUSM 3+	*AUSMQ* 3+	*IVG* 1 and 2	*IVG* 3+
AUSM 1 and 2	$t = 2.65$ $\alpha = .01$	$t = 4.44$ $\alpha = .00$	$t = -.53$ $\alpha = .30$	$t = .93$ $\alpha = .18$
AUSM 3+		$t = -2.68$ $\alpha = .00$	$t = -2.13$ $\alpha = .03$	$t = -.16$ $\alpha = .43$
AUSMQ 1 and 2	$t = -2.29$ $\alpha = .01$	$t = .35$ $\alpha = .36$	$t = -2.80$ $\alpha = .01$	$t = -1.07$ $\alpha = .15$
AUSMQ 3+	$t = 2.68$ $\alpha = .00$		$t = -3.65$ $\alpha = .00$	$t = -1.53$ $\alpha = .07$
IVG 1 and 2				$t = 1.16$

*α indicates the level of significance (one-sided) for the test that the revenue generated by the mechanism in the column equals that in the corresponding row.

TABLE 15

Mean and 95% Confidence Intervals
on Prices in Francs*

Market	Periods 1–4		Periods 5–8	
	μ	(95%)	μ	(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

TABLE 16

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)

Figure 1
Tree Structure of Events

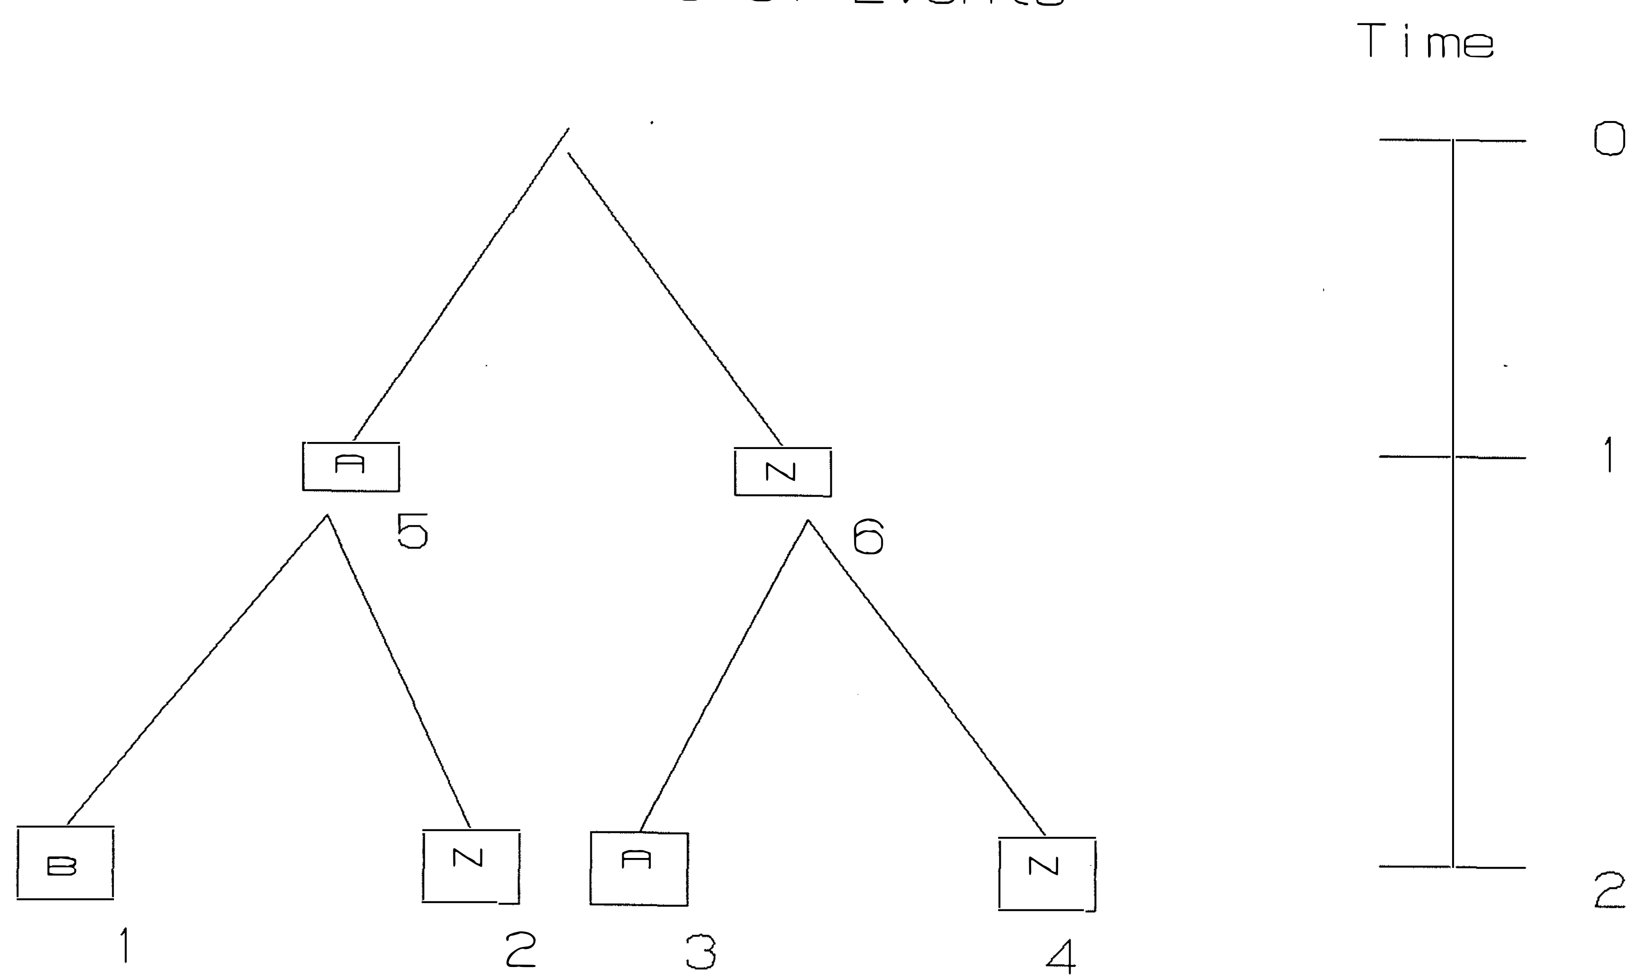

Figure 2

Distribution of Efficiency (30000 Random Combinations)

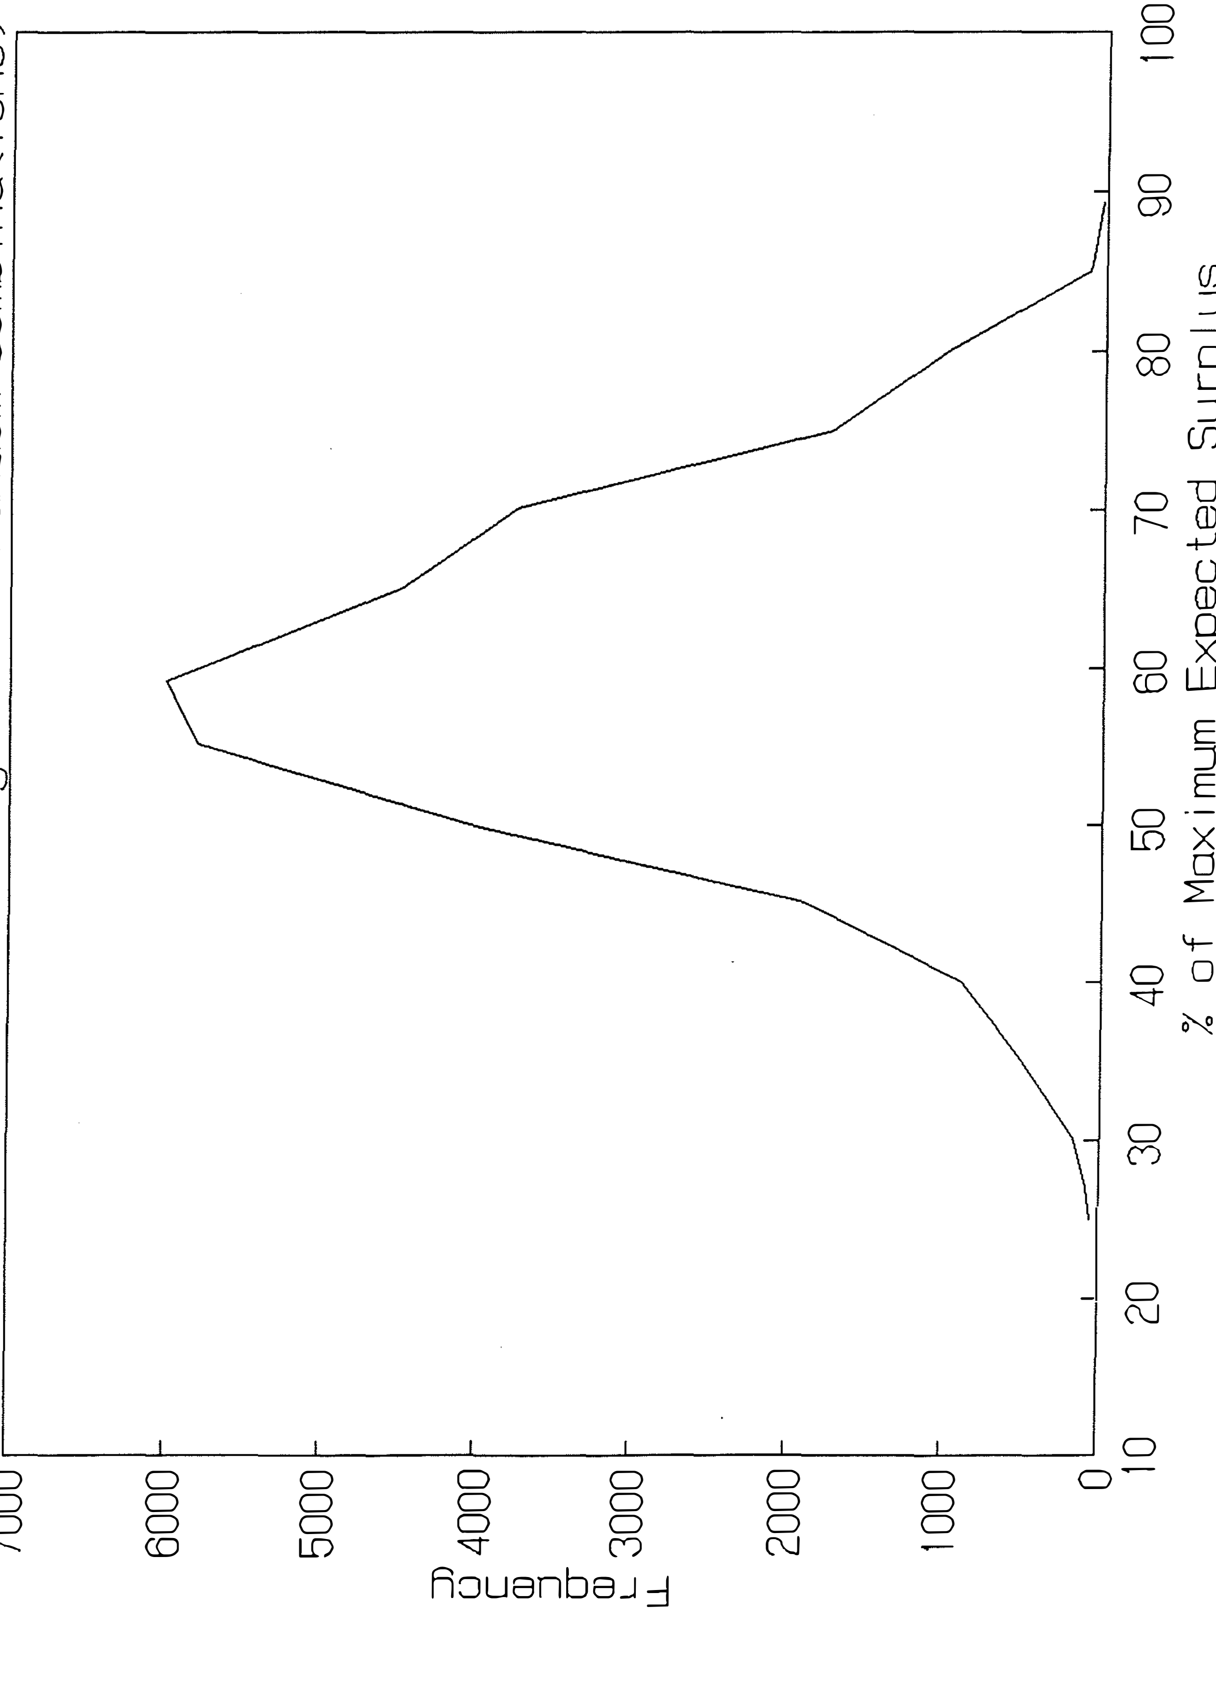

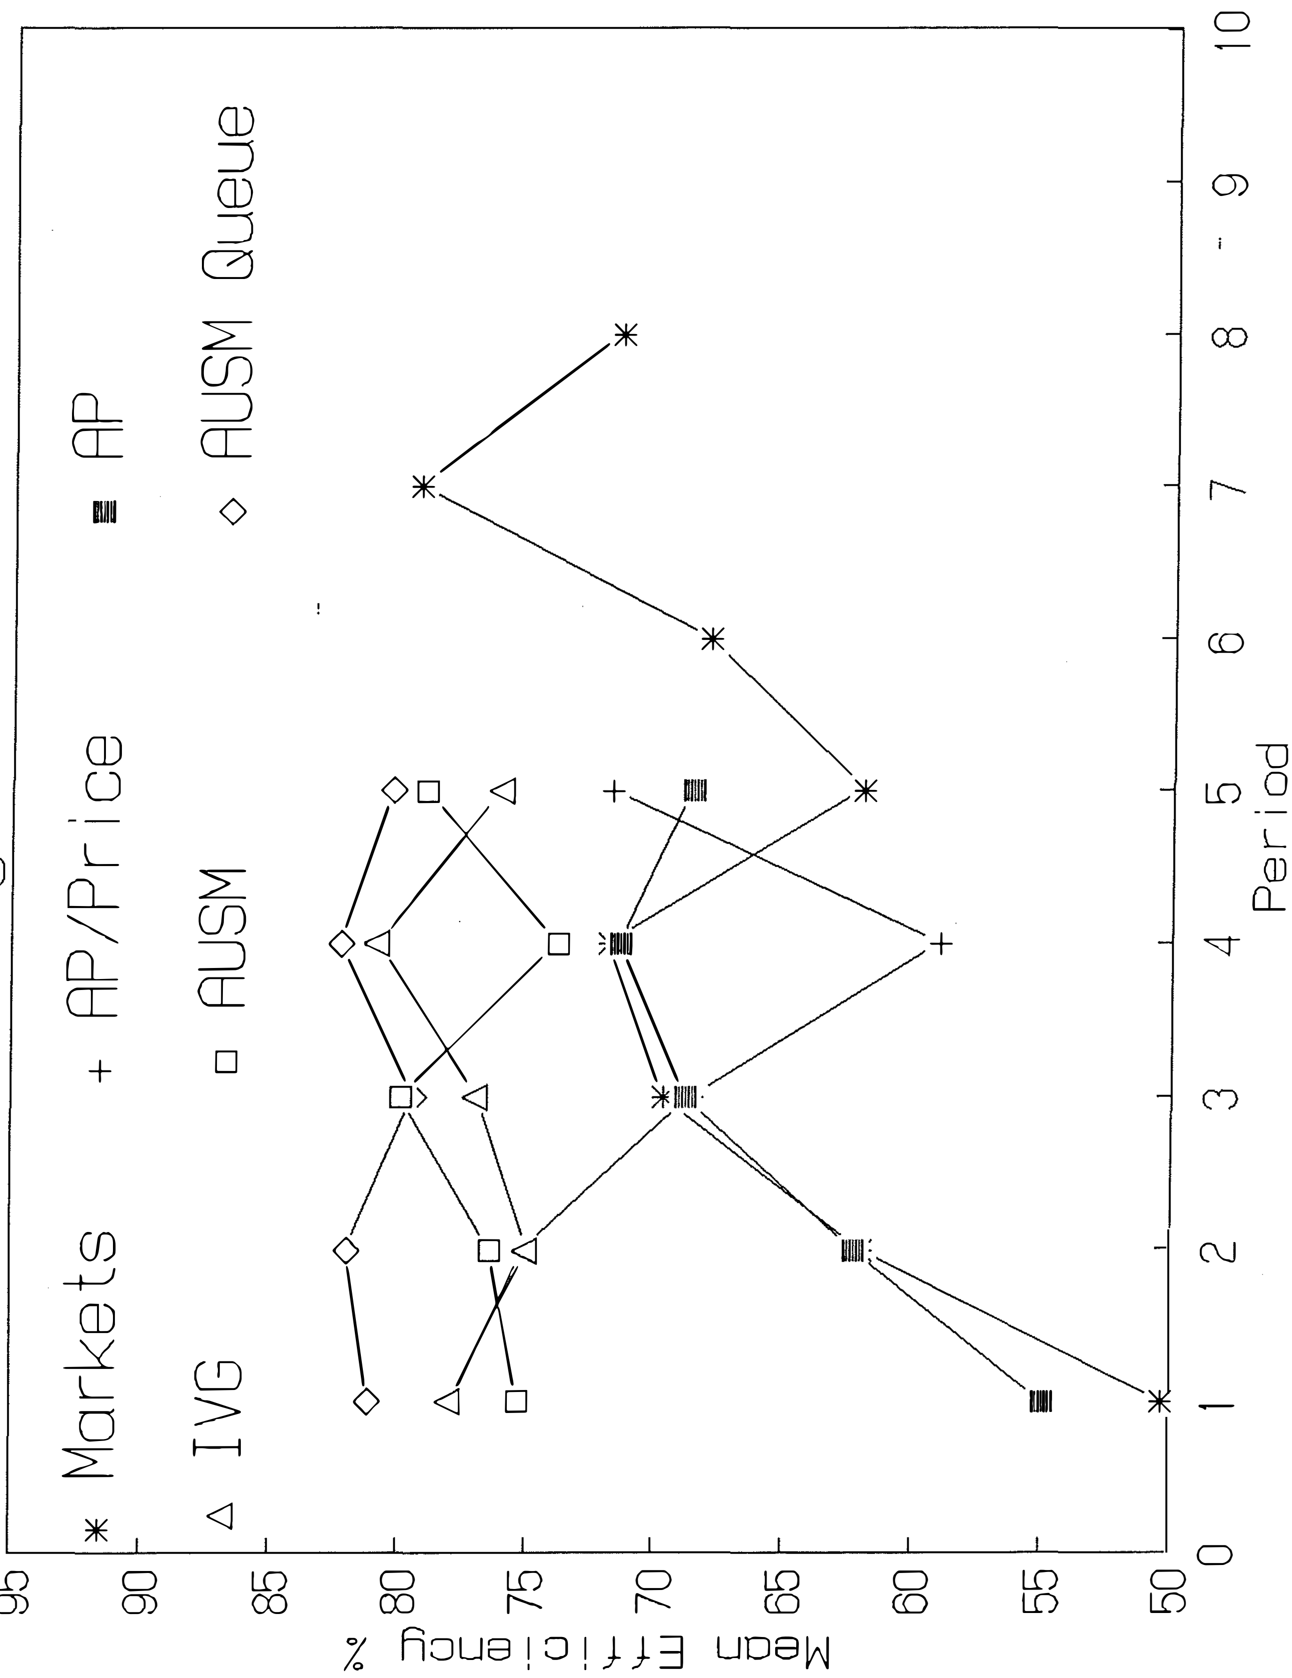

Figure 3
Mean Efficiency Per Period

Figure 4
Histogram of Revenue

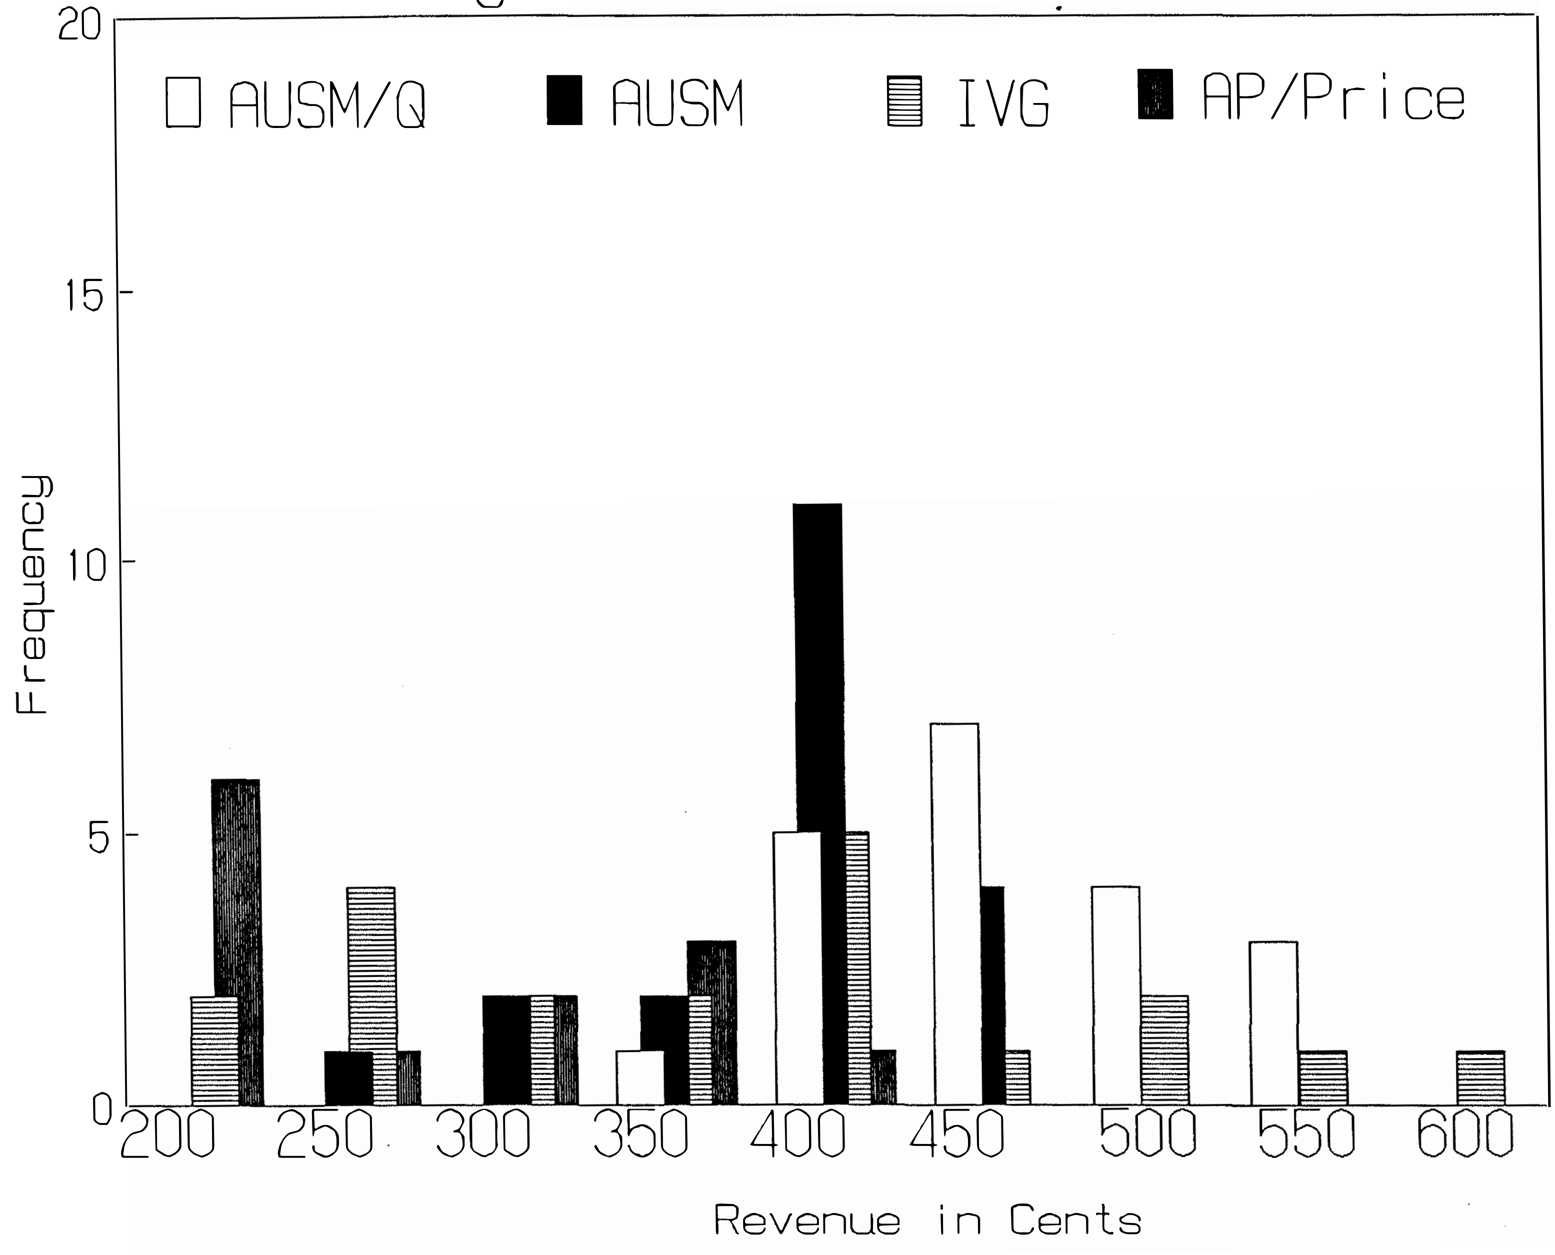

FOOTNOTES

1. This work was partially funded by Caltech and NASA-JPL. 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. This paper represents a major revision of our earlier 1987 paper. All the material on markets is new as are some of the conclusions.
2. See Banks, Ledyard and Porter (1985); and Fox and Quirk (1985) for details.
3. For a more extensive discussion of the Space Station allocation and decision-making problems, see Ledyard (1986).
4. For example, in electric networks, new power plants can be fired up yielding responsive supply, and demand uncertainties are probably uncorrelated - except for weather events such as heat waves - and therefore, 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).
5. Of course this regulation is undoubtedly mostly designed to control pipe-line owners in monopoly positions. But it is also used to coordinate demand-supply imbalances.
6. See Grether, Isaac, and Plott (1981).
7. 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.
8. 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 which 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 which 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 which will be *ex post* optimal *in all possible branches from 5*. For this reason we have chosen *ex ante* Pareto Optimality as the primary performance standard by which we will judge alternative allocation mechanisms.
9. Even so, it is easy to construct examples of market-like institutions and game structures in which only pure strategy equilibria exist and in which the equilibria allocations are not optimal.
10. The rules impose a bid-ask improvement rule and provide for the public broadcast of the standing bid-ask spread.
11. 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

$$U^*(a^i, x^i, b^i) \equiv g^i[M^i(a^i) - b^i](\Sigma_e \gamma^i(a^i, x_e^i, e))$$

$$+ [1 - (\Sigma_e \gamma^i(a^i, x_e^i, e))]\, g^i(-b^i) \;,$$

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 VII contains data which suggests that assuming risk-neutral subjects was not at too much variance with reality.

12. 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.
13. The Caltech computerized version was used which allows for up to 19 simultaneous markets. For more information on this software see Johnson et. al. (1988).
14. 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 a lot of common information.
15. The calculation is tedious but can be done on a PC.
16. We do not know whether the vector of efficient payloads is the only equilibrium. We also do not know whether or not there are other equilibrium joint payload designs $(x^1, \ldots, x^6)$ that yield higher efficiencies than this equilibrium.
17. 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 and 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 (see Harsanyi (1967–8) and Harsanyi (1980)) for these points—then one could adapt the approach of D'Apresmont 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.
18. This point has been made by Hyeck (1945) Marshak (1972) and others.
19. For example, see Smith (1980) or Ferejohn, Forsythe, and Noll (1979).
20. The rules used were developed in response to early testing. Initially we did not close the markets sequentially; however, the process was pushed by the subjects to the last trial in most instances (even with 40 trials) and 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 which required individuals to better *their* previous bids in the market they were ordering. This rule caused individuals to be cautious in their bidding or "locked" them into larger projects yielding low efficiency levels and so was eliminated.
21. One modification of the Iterative Vickrey--Groves mechanism, proposed by seminar participants at New York University, which might stabilize its performance and improve its efficiency, would be to allow the algorithm which 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 ∞ iterations) one might achieve the full Vickrey–Groves scheme. This should be followed up.

22. Of course, if enough contingent bids could be submitted in a sealed-bid, it could mimic an iterative process, but in an informationally more complex manner.
23. If omitted users could replace just the marginal units of those users in the potential allocation, then it would not be costly to bump 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, yielding an entirely different mechanism.
24. The queue was the only major design change we made in response to early testing. There were, however, two other minor, but significant, variations in design which we chose in response to experimental testing. Originally, we had a stopping rule that the auction would run for T minutes and the allocation at T would be final. It had undesired effects. In the pilot experiments, very little bidding occurred until $T - \varepsilon$ when a flurry of bids were presented. Allocations were essentially random. This was easily solved 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 change bids up or down 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 $T - \varepsilon$. We fixed this by revising the rules, by which bidding could be done, in a way which made each bid a potentially binding contract. Further, an explicit improvement rule for bids was added. 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.
25. Recall that our participants were experienced with multiple markets and computerized bid and ask procedures.
26. See Daughety and Forsythe (1987) for an initial attempt to use experimental methods to study this issue.
27. In that more elaborate testbed, we induced more continuous demands and the existence of the competitive equilibrium. AUSM produced efficiencies at the 80 percent level while markets were 89 percent efficient (80 percent in early periods) and the administrative process was only 60 percent efficient.
28. We would change the model in Section II as follows: Let B^i be i's budget as determined by NASA and Congress and $W^i(Z^i)$ be the utility for the scientific returns Z^i. Let $Z^i = f(x^i, a^i)$ be the scientific return from project a^i if resources x^i are used, and let $c^i(a^i)$ be the cost of design and construction of project a^i. Now let $U^i(a^i, x^i, b^i) = MAX\ W^i(Z^i)$ subject to $b^i + c^i(a^i) \leq B^i, Z^i = f(x^i, a^i)$, and $a^i \in A^i$. This change in U^i could be easily accommodated in the experimental design.
29. One pilot experiment with AUSM modified to allow users to buy and sell (by "buying" negative amounts) never got off the ground. No feasible "trial allocation" ever occurred. If this result is replicated, other mechanisms will be needed for situations with significant supply decisions.

REFERENCES

Banks, J. S., J. O. Ledyard, and D. Porter. "Pricing Evolution, and Design Planning of Space Station under Uncertainty." JPL Economic Research Series No. 22, Pasadena, CA., (August 1985).

________________. "The Design of Mechanisms to Allocate Space Station Resources." Social Science Working Paper No. 648, California Institute of Technology, Pasadena, CA., (June 1987).

Chao, H., and R. Wilson. "Priority Service: Pricing, Investment and Market Organization." *American Economic Review* 77, (December 1987): 899–916.

Cox, J. C., B. Roberson, and V. L. Smith. "Theory and Behavior of Single Object Auctions." *Research in Experimental Economics* 2, (1982): 1–44, JAI Press, Inc., London.

D'Aspremont, C., and L. A. Gerard-Varet. "On Bayesian Incentive Compatible Mechanisms." In *Aggregation and Revelation of Preferences*, edited by J.-J. Laffont. North Holland, Amsterdam: (1979): 269–288.

Daughety, A., and R. Forsythe. "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, New York: John Wiley & Sons Inc., (1959).

Ferejohn, J., R. Forsythe, and R. Noll. "An Experimental Analysis of Decision Making Procedures for Discrete Public Goods: A Case Study of a Problem in Institutional Design." In *Research in Experimental Economics*, edited by V. Smith. Greenwich: (1979).

Fox, G., and J. Quirk. "Uncertainty and Input-Output Analysis." JPL Economic Research Series No. 23, Pasadena, CA., (October 1985).

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

Grether, D., M. Isaac, and C. Plott. "The Allocation of Landing Rights by Unanimity by Competitors." *American Economic Review* 71, (May 1981): 166–171.

Groves, T. "Incentives in Teams." *Econometrica* 41, (1973): 617–663.

Harris, M., and A. Raviv. "A Theory of Monopoly Pricing Schemes with Demand Uncertainty." *American Economic Review* 71, (June 1981): 347–65.

Harsanyi, John C. "Games with Incomplete Information Played by 'Bayesian' Players." *Management Science* 14, (1967–8): 159–182, 320–334, 486–502.

Harsanyi, John C. "Bayesian Probability in Game Theory." In *Science Belief and Behavior: Essays in Honor of R. B. Braithwaite*, edited by D. H. Mellor. Cambridge University Press, (1980): 192–196.

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

Johnson, A., H. Y. Lee, and C. R. Plott. "The Multiple Unit Double Auction Users Manual." Caltech Working Paper No. 676, Pasadena, CA., (January 1988).

Kirkpatrick, S., C. D. Gelatt, Jr., and M. P. Vecchi. "Optimization by Simulated Annealing." *Science* 220, (May 1983): 671–220.

Konijn, H. "Distribution-Free and Other Prediction Intervals." *The American Statistician* 41, (February 1987): 11–15.

Koopmans, T. C., and M. J. Beckmann. "Assignment Problems and the Location of Economic Activities." *Econometrica* 25, (January 1957): 53–76.

Ledyard, J. O. "The Economics of Space Station." In *Economics and Technology in U.S. Space Policy Proceedings of a Symposium*, edited by Molly K. Macanley. Washington, D.C.: Resources for the Future, Inc., (1987): 127–170.

Lehmann, E. *Non-Parametrics: Statistical Methods Based on Ranks.* San Francisco: Holden Day, (1975).

Marshak, T. "Computation in Organizations: The Comparison of Price Mechanisms and Other Adjustment Processes." In *Decision and Organization*, edited by C. B. McGuire and R. Radner. North Holland, Amsterdam: (1972).

Milgrom, P., and R. Weber. "A Theory of Auctions and Competitive Bidding." *Econometrica* 50, (1982): 1089–1122.

NASA. *Space Transportation System: User Handbook.* Washington, D.C.: (May 1982).

________________. *Space Station Operations Task Force: Panel 3 Report.* Washington, D.C.: (December 1987).

Pitblado, R. "Financial Incentives for Efficient Airline Scheduling and Service Prioritization." Working Paper, GSB University of Rochester, New York: (July 1987).

Plott, C. R. "Industrial Organization Theory and Experimental Economics." *Journal of Economic Literature* 20, (December 1982): 1485–1528.

Plott, C. R., and D. P. Porter. "An Experiment with Space Station Pricing Policies." Mimeo, Pasadena, CA., (1988).

Reiter, S., and G. R. Sherman. "Allocating Indivisible Resources Affording External Economies or Diseconomies." *International Economic Review* 3, (January 1962): 108–135.

Reiter, S. "A System for Managing Job-Shop Production." *The Journal of Business of the University of Chicago*, (July 1966): 371–393.

Reitman, D. "Pricing, Quality and Priority Service in Congested Markets." GSB, Stanford University, (November 1985).

Ressenti, S. J., V. L. Smith, and R. L. Bulfin. "A Combinatorial Auction Mechanism for Airport Time Slot Allocation." *The Bell Journal of Economics* 13, (1982): 402–417.

Smith, V. L. "Experimental Economics: Induced Value Theory." *American Economic Review Proceedings* 66, No. 2, (May 1976): 274–279.

________________. "Experiments with a Decentralized Mechanism for Public Goods Decision." *American Economic Review* 70, (September 1980): 584–600.

________________. "Microeconomic Systems as an Experimental Science." *American Economic Review* 72, No. 5, (December 1982): 923–955.

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