

Coordination in Shared Facilities: A New Methodology

John O. Ledyard

California Institute of Technology

Shared facilities are a good example of the difficulties inherent in coordination problems and the benefits to be derived from creative solutions. Traditional methods employed by engineers and others, because they ignore an important aspect of the problem, can yield solutions that appear successful but which significantly underutilize these facilities. This article is intended to be an introduction to the types of problems that can arise and to a new method for systematically studying these problems. The method is illustrated with the results of a study done for NASA, on the coordination of the use of a Space Station, which produced a new computer-assisted institution that outperforms existing institutions.

coordination, shared facility, experimental economics,
group decision support system, Knapsack problem, mechanism design

1. INTRODUCTION

One of the more complex problems in the coordination of the activities of many different individuals is to be found in the management of a shared facility. There are many examples: laboratories shared by scientists, wind tunnels, mainframe computers used simultaneously by many programs, networks that are to handle information flows, natural gas pipeline networks, job shops that must be scheduled, airports, roads, bridges, grazing lands, etc. Consider, for example, the coordination problem found in airport scheduling and the allocation of takeoff/landing slots. The facility is a multidimensional collection of resources shared by many users, each with their own ideas as to how things should happen. The arrival and departure of a single plane requires a slot, baggage handling, a gate, lounge facilities for the check-in of passengers, parking, etc. The arrival and departure of many planes creates a complex coordination problem. Just avoiding chaos is difficult. In addition, weather and mechanical failures can create uncertainties that must be responded to. Finally, the method by which the coordina-

Correspondence and requests for reprints should be sent to John O. Ledyard, Division of the Humanities and Social Sciences, 228-77, California Institute of Technology, Pasadena, CA 91125.

tion is managed can have a profound effect on the decisions of the users (types of aircraft, times of flights, etc.), on the level of service provided, and, ultimately, on the success or failure of the various airlines that use the facility.

Engineers and others have expended considerable effort to discover and create algorithms that can be used to coordinate shared facilities and to schedule their use. Some of these have indeed been able to bring some manner of order to the complex problem. Computer operating systems generally get things done, airlines do manage to coordinate their activities not only at a single airport but also across airports, and traffic (usually) flows across the bridges of New York. But maybe we should ask for more. Are we getting the most we can from these shared facilities? Can we improve the coordination of their use? Can we study the design of new decision support systems in a systematic manner?

This article introduces a new way to study coordination problems, with particular emphasis on the coordination of shared facilities. It is argued that traditional methods (including those of engineers, psychologists, and sociologists) have severe drawbacks, because they have neglected an important aspect of the problem, and that they provide solutions that may underutilize significantly the available resources.¹ A new empirical methodology—experimental economics—with which we can actually test the claims presented herein in a replicable manner will be described. It will also be illustrated how new findings in the theory of games and mechanism design have provided the analytical framework within which new institutions can be created in a systematic manner. Finally, it will be illustrated how the methods of experimental economics can be used to test new institutional designs in the same way new airframes can be tested in a wind tunnel.

2. A VERY SIMPLE SHARED FACILITY

Let us begin with a very simple example of a shared facility: a narrow bridge over which only one automobile can cross at a time. If there is only one car on the road at a time there is no problem, but if two cars, traveling in opposite directions, want to cross simultaneously then there is a problem. The bridge is somewhat isolated and never has more than two cars near it at a time. Each car can take one of two actions—go or wait—yielding four possible outcomes. The situation can be summarized in a simple matrix² (see Table 1) describing the payoffs to each driver in each of the four possible outcomes. The first entry in each square represents the payoff to automobile 1, while the second entry is the payoff to automobile 2.

¹Herein by underutilization does not mean “not used.” What is meant is that, with better coordination, the same resources could produce more output or that fewer resources could produce the same output. That is, the author is interested in more than simply keeping shared facilities busy; he is interested in efficient use.

²Those familiar with game theory will recognize this as the payoff matrix for a game of chicken.

Table 1

		AUTO 1	
		GO	WAIT
AUTO 2	GO	-100, -100 a	-10, 10 b
	WAIT	10, -10 c	-20, -20 d

a. The Obvious Answer

It is easy to see that both drivers would prefer outcomes b and c to outcomes a and d ; however, the driver of auto 1 prefers c , while the driver of auto 2 prefers b . How is the conflict resolved? If each is stubborn and assumes the other will give way, then outcome a will result. If each is courteous and assumes the other will barge right through, then outcome d will occur. In both cases, the outcome that results is entirely different from that which each anticipated in their calculations and clearly worse than what is possible. Some type of coordination seems desirable. The problem to be solved is to find an *institution*³ that guides the drivers to coordinate their actions in a way that yields the highest possible return.

One institution would be to adopt the convention that the car traveling north always yields to the driver traveling in a southerly direction. If both drivers use this bridge fairly often and if both travel in each direction then, perhaps, over time, each will receive roughly equivalent payoffs; on average they receive 0, getting -10 when traveling north and 10 when going south. The institution—yield when driving north—solves this simple coordination problem. No other institution can do better.

b. A Complication—Asymmetric Information

The claim that the preceding solution is the best one can do relies on two crucial assumptions: repeated symmetric use by the drivers, and symmetric payoffs to the drivers. Suppose instead that sometimes auto 1 may be carrying an injured person who must get to a hospital as soon as possible. Further assume that in this case the matrix entries are as shown in Table 2.

As before, outcomes b and c are socially better than a and d . Also, as before, 1 prefers c and 2 prefers b . But if the payoffs are capable of meaningful comparison across drivers, then one might argue that c is now better than b since the combined payoff is 10 if c occurs, whereas it is -30 if d is the outcome. The natural question to ask is, "How does the institution we created above perform in this new, somewhat more complex situation?" The answer depends on the

³Herein "institution" is used to describe the mechanisms, organizational rules, decision support systems, customs, etc., used to guide coordination.

Table 2

		AUTO 1	
		GO	WAIT
AUTO	GO	-100, -110 a	-40, 10 b
	WAIT	20, -10 c	-50, -20 d

particular *environment*.⁴ If drivers going to the hospital always drive south, then the institution, northerly drivers yield, produces exactly the desired outcomes. If, however, drivers with injured passengers may be traveling either way, then half the time this institution will produce a less desirable outcome. This example already illustrates that our ability to claim that a particular institutional arrangement is a good one depends on the extent of our information about the environment in which it operates. If we knew that injured people always rode south, we would be finished. Conversely, if we did not have any idea which direction is relevant, then more work remains.

c. Try a Technological Solution

Can we redesign the institution to perform better *if we don't know what the environment is*? A seat-of-the-pants approach is to look for a technological innovation. We could, for example, allow the driver with the injured passenger to signal the situation by sounding a siren and require the driver not sounding a siren to yield the right of way. If both, or neither, sound a siren we can retain the original institution and require the car driving north to yield. A naive examination of this new institution yields the prediction that outcome *c* will always occur when only one car has an injured passenger, while *c* and *b* occur roughly equally often otherwise. Our redesign would seem to work. This is the point at which most analysts stop. However, there is an implicit assumption that must be examined. In particular, it has been assumed that all drivers will follow the rules despite the fact that driver 2 has an incentive to misbehave. To see this, suppose driver 2 always sounds a siren on approach to the bridge *even when automobile 2 carries no injured passenger*.⁵ The drivers will now find themselves back in the situation that existed before the technological innovation of the siren. All signals will be ignored (hearing a siren provides no information). Since the institution to yield when driving north is used when both sirens sound, driver 2 will gain by this deception since, instead of always receiving -10 at each crossing (the amount received by 2 in outcome *c*), an amount of $+10$ is received half the time

⁴The environment includes the data of the problem that are neither under the control of nor known by the institution designer.

⁵Some may argue that no one should, or would, do this. That is probably just as naive and incorrect as the assertion that everyone would. How people would behave in this situation is an empirical issue that can be studied. What is important is that the incentive to misbehave exists and, presumably, some would give in to the temptation.

(when 2 is driving south). Unless there is some way to monitor the cars and to observe whether or not sounding the siren is a valid action, *the incentives of the drivers may lead them to actions that undo the good intentions of the new institution.* Everyone may be worse off than if nothing had been done. For example, if sirens cost money, then the institution to sound a siren if you have an injured passenger will leave everyone worse off than the institution to yield if driving north, because the same outcome occurs in both (c and b each occur half the time) but, in the first case, everyone will have bought a siren. The example illustrates two very important lessons for the institutional designer. First, good intentions in institutional design that ignore the likely behavioral responses of the participants can easily leave the group worse off than no change at all. Second, we may never be able to achieve the best outcome because the necessary information is not at our disposal and those who do know what we need to know may not have the incentive to provide it. If it is not possible to observe whether or not a car has an injured passenger, then there is no mechanism that can yield the desired outcome c . The best one can do is to settle for c or b , each half the time.

d. Add a Bureaucracy

We could try to undo what we have done with the sirens. Rather than just giving up and accepting the outcomes of c and b each half the time, we could hire a policeman to monitor and enforce the rule that only drivers of cars with injured passengers are allowed to sound their sirens. If the penalties are sufficiently strong (e.g., the driver is executed if caught sounding a siren without an injured passenger), then this new institution should encourage all drivers to follow the rules and thus yield the desired outcome c . At this point, we have moved a long way from the original, costless institution "yield if driving north." We have required all drivers to install sirens, and we have hired an enforcement bureaucracy. In those cases when two cars try to cross the bridge in opposite directions and only one contains an injured passenger, we have a combined payoff of 10 ($=20-10$). If both have injured passengers, we only get -20 ($=20-40$); if neither has injured passengers, we get 0 ($=10-10$). If we had left well enough alone and simply used the "yield if going north" institution, we would get -10 ($=-1/2(-30+10)$), -20 , and 0. Thus, by adding a policeman we gain 20 in those cases with one injured passenger and we gain nothing otherwise. Whether this is good or bad depends on the number of times there is just one injured passenger, the cost of each siren, and the cost of the police. It could easily be the case that, although we have apparently achieved the desired outcome, we have had to pay too high a price to enforce the solution. We might have been better off accepting the limits placed on our institutional design by the behavioral responses of the drivers.

e. What Can We Conclude?

Although this example is rather simple, it contains most of the interesting aspects of the coordination problem created by shared facilities. First, *we cannot violate physical laws.* In our bridge example, we could not consider solutions that

had both cars cross the bridge simultaneously without crashing. We had to accept the physics of the situation. Second, *we cannot violate behavioral laws*. We cannot ask drivers to take actions that will not be their interest to follow. We had to accept the strategic possibilities as a constraint on what we could do. But, third, within those constraints *institutions matter*. Changing the rules changes the outcome. This means that, if we are clever enough, we can create institutions that will guide the system to desirable outcomes. The example of the bridge is simple enough that we could each probably figure out what the best institutions were without further help. However, if we are not particularly clever, we can easily make things worse than might otherwise naturally occur. Good intentions that ignore the constraints can have costly implications. To see how complicated this can become, let us turn to a more realistic situation for which the appropriate institutions are not as obvious.

3. A MORE COMPLEX EXAMPLE—THE SHARED LABORATORY

A better example of a shared facility comes from a project to design effective institutions to price and allocate resources on NASA's planned Space Station, best thought of as a very fancy research laboratory at a very remote location. As in the previous bridge example, there is a facility of fixed size that is to be used by several participants, simultaneously if possible. The facility, however, is not a simple one-dimensional entity; instead, it is a collection of resources to be shared among many potential users. When a scientist or businessman brings a *payload*⁶ to the Space Station, it will not only occupy space (as cars did on the bridge), it may also need manpower, power, and other resources to fulfill its mission. Generally, the more resources a payload uses, the bigger the returns. However, since each resource is in limited supply, it must be shared among the many potential users of the facility. Furthermore, each payload must use its vector of resources simultaneously—coordination must be managed across all resources.

a. A NAIVE ATTEMPT

As will be evident shortly, the problem to be solved is complex and multi-layered. Therefore, to provide a slow immersion, let us proceed one step at a time. Suppose for now that each payload can be built and operated only in one way. Alternatively, suppose each scientist and commercial researcher has already designed and built their experimental apparatus. They simply await the use of the facilities. We can characterize each payload by the vector of resources required for successful project completion. The problem then is to select a group of payloads that fit within the resource constraints of the station.⁷ The reader

⁶"Payload" is the generic term used to describe the physical equipment shipped to the station to carry out research or production. "Experimental apparatus" is another term that could be used.

⁷Formally, we can write the problem as: Chosen as set C of users to maximize the total return subject to $\sum x_i \leq x$, where x is the vector of all resources available for the distribution and x_i , the vector of required resources for payload i . Note that only C is being chosen.

who has studied optimization will easily recognize this as a Knapsack problem. The station resources represent the limits of the knapsack, whereas the payload requirements represent the size of the items available to be inserted into that knapsack. It is not too difficult to imagine finding at least one configuration of payloads that fits, and may even use all the resources. The real problem is to find that, possibly, unique combination of payloads that produces the best total return, even if all resources are not used. If engineers were provided with a measure of the value created by the successful completion of each payload's mission, they could then solve the coordination problem by maximizing the sum of the values of the included payloads subject to the constraints imposed by the available resources. A naive institution that would appear to solve the coordination problem requests each scientist to provide a description of the resources required for their project, and a number indicating its value, to a computer, which would then use standard algorithms to solve the optimization problem. If the information provided were correct and the computer were of sufficient power, then an optimal outcome could be achieved.

But we have made a mistake in our design! We have taken into account the physical constraints required by resource availabilities, but we have not taken into account the laws of social science. Within the proposed institution, there is absolutely no incentive for the scientists or commercial researchers to provide the correct information to the computer. Assuming that being on the Station is better for each rather than not being included, each potential user has an incentive to overestimate the value of their project and to underestimate its resource requirements. If all do this, the computer program solving the problem using these incorrect data, but assuming they are correct, will select a collection of payloads that, in fact, will violate the physical constraints and will also probably have little relationship to the collection that has the highest total value. When these payloads arrive on the Station, the overestimation of resources will be discovered (and each will claim that an unusual event must have caused them to need more resources than planned) and adjustments will have to be made; some payloads will have to be denied resources. At best, this reallocation will probably occur at random. In effect, although the institution was designed to select an optimal combination of payloads, because of the incentives created by that institution, the actual result will be as if the payloads were simply randomly selected. The laws of behavior will have caused all the good intentions of the institutional designer to be undone.

A skeptic might argue that, since the managers of the Station will be able to observe the actual use of resources after the payload uses the Station, we should be able to penalize any project that oversteps its original request. To take an extreme position, if a user knew that they would be shot upon discovery that they had underestimated their resource requirements, it is likely they would report honestly and we could then implement the desired outcome. Thus, if the manager can monitor and observe the actual resource use, and if that manager can impose extreme punishments for a faulty report, then we can assume that the reported resource requirements will be correct. If the manager can also acquire independent measurements of the payload value (a feat that is generally

impossible in practice), then the problem can be solved as proposed earlier with some modifications to the enforcement structure.

b. Asymmetric Information

Unfortunately, in most situations, the ability of the managers to acquire independent data is extremely limited or extremely costly. To see why, let us move to the next level of complication. Suppose that not only has our potential user not yet built an experimental apparatus, but also that the payload design can be chosen from among a variety of configurations. Each possible design requires different vectors of resources and each provides results of differing value.⁸ To continue to try to keep the problem simple, let us assume for now that each potential user has only two possible designs. As before, we could ask each potential user to report the data on both designs⁹ to our manager, who would then use the computer to choose a collection of users, and one design for each, to maximize the sum of the values of the experiments sharing the Station resources, subject to the constraints imposed by resource requirements. The problem has increased suddenly in complexity, both with respect to the computer algorithms required and the behavioral implications of the process. Let us look at just one way for the institutional design to go wrong. Suppose that with complete and accurate information, the manager would select design 1 of a particular biologist for inclusion in the laboratory. Suppose, further, that this biologist has a strong desire to see design 2 (a more extensive project) included instead. Finally, suppose that the laboratory manager is not much of a life scientist and has no way of knowing what possible designs the biologist might be able to utilize. In this situation, it may be in the interest of the biologist to overstate greatly the resource requirements¹⁰ of design 1 so that the manager will discard it as a possible choice. It is important to realize that the manager has very little defense against this move. Even the threat of excessive punishments is not enough unless the system is willing to require the scientist to present an operational model of each design so that the stated requirements can actually be checked out.¹¹ Absent these extreme measures, any attempt by the manager to

⁸Consider a simple example. Experimenters on the Station conceivably could choose between providing their own self-contained power source or tapping into Station-provided power. Providing their own power could improve significantly the reliability of supply (the Station may need to cut power in emergencies) and lead to a higher expected value, but it would also increase the mass of the apparatus. The latter is relevant since it increases the resources necessary to raise the payload into orbit.

⁹Note that if we requested data only on one payload, the user might have an incentive to report on the project design that was most likely to result in the inclusion of that experiment on the Station. This may not be the design that a fully informed manager would select.

¹⁰The scientist need not even do this intentionally, with either malice or cunning. In many cases, designs are simply estimates—until the experiment is actually run—so that the biologist need only be pessimistic about the resource use of design 1 and optimistic about the resource requirements of design 2 to achieve the desired ends. Of course, a manipulative user would push the system even further.

¹¹Of course, this method of checking validity is very expensive and similar in spirit to the policeman in the bridge example.

claim that the data for design 1 are incorrect is debatable and subject to appeal. The manager can measure the resources used by the design selected, so the scientist must report accurately on that project, but the manager cannot measure the resources required for designs that "might have been used." The designer of the institution thus faces a fundamental problem: How to get the right information to the manager so that a good system decision can be chosen, while accepting the fact that the individuals who have the necessary information may have an incentive to manipulate.

4. A SYSTEMATIC APPROACH

As we have seen, coordination and allocation of shared facilities is a complex problem. Not only can the engineering issues be difficult, but the situation is complicated as well by the presence of well-meaning but self-interested users who possess the information that the engineer needs to solve the problem. How can we ever hope to study these situations in a way that will allow us to provide better management institutions? Recent developments in game theory, experimental economics, and computer science have created the opportunity for a promising new approach to the study and analysis of shared facility coordination. The important development in game theory has been the discovery that we can axiomatize and describe the class of institutions. That allows us to systematically study institutions as mathematical objects; we need not be dependent solely on historical accident or seat-of-the-pants guesses to discover new institutions. We can, in fact, treat them as we would any other variable of analysis. This discovery allows us to convert the design problem to an optimization problem.¹² The development of experimental economics provides the ability to test new designs in a controlled way. Just as the aeronautical engineer uses a wind tunnel to test new airfoils developed by the theorist using mathematical methods, we can now use experimental methods to test newly conceived institutions. We can actually try out new designs before a full-scale implementation. Finally, developments in computer science have created the engineering capability to actually implement institutions that have never before been possible. However, the best way to understand the potential of this new field of study is to see it in action.

This section provides an introduction to the power and capabilities that created this new approach to institutional design by summarizing some of the highlights of a study reported in Ref. 2. The task of the study was twofold. First, we were to evaluate how well existing institutions would perform in the coordination of activities by multiple users sharing a facility of fixed size. Second, if the existing institutions were to be found wanting, we were to design new institutions and to provide supporting evidence as to the extent of the improvements. The two existing institutions studied were examples of the most commonly used organizations: a standard administrative process and a standard

¹²The actual problem may be messy and difficult, but it will, in principle, have a solution.

market process.¹³ Both institutions were found to perform at very low efficiencies (60–65% in seemingly robust experimental tests), where efficiency measures how much total value was achieved by the system relative to the maximal achievable by a fully informed, beneficent manager. Given this low level of performance, new institutions were needed. Several institutions were created using the new theoretical methods of this field of study. The one we will consider herein involves a form of computer-assisted coordination. It was shown that this new institution improves significantly the performance of both markets and standard administrative procedures. In what follows, only a brief description of each of the three institutions and some of the relevant data is given. For more information, see Ref. 1 for the technical report or Ref. 3 for a more policy-oriented presentation.

a. The Environment and Testbed

To begin the study, a class of environments that had many of the features of real-world shared-facility problems but that was still simple enough to allow analysis and study in a controlled manner was created. Our problem here is similar to the engineers who use a wind tunnel to test airplane designs and must choose the forms of turbulence to which they will subject the designs. We chose a variation of the problem presented in Section II. In this "testbed" environment, there is a facility with two resources, x and y , to be shared over two periods by, potentially, six agents. Each agent can design a payload and each is able to select the design from a set of nine alternatives.¹⁴ Each alternative design requires a fixed amount of the two resources to be successful, and each design provides a different value if successful. The possible designs and the value created, if successful, in the first period, are provided in Appendix A. The value of the design, if successful in the second period, is exactly 40% of that in the first period, to incorporate costs of delay and the desire to be first. The facility has the ability to provide 20 units of each resource in each period. Given this information and enough computer time, one can easily compute the allocation of resources that maximizes the sum of the value of successful projects. That allocation is the solution to the coordination problem: who should use the facility, when should they use it, how much of each of the two resources should they use, and what should each design project be? The solution is detailed in Table 3.

The best way to see how difficult it might be to solve this problem is to consider what would result by purely random choices. We made the following calculation 30,000 times. We randomly (using a uniform, equal probability, density) selected one of the six possible users, then randomly selected one of

¹³Communities and markets seem to have been the two dominant forms of institutions spewed out by the historical process of evolution. Each of these forms can have many variations, some of which are known to be important. For example, a committee that has a chairperson with veto power performs significantly differently from one in which simple majorities rule.

¹⁴Each user's set of alternatives is different, and the details of the alternatives are known only to that agent.

APPENDIX A

Valuation Sheet 1

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

Valuation Sheet 2

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

Valuation Sheet 3

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

Valuation Sheet 4

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

Valuation Sheet 5

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

Valuation Sheet 6

Y	7	9	11
X			
7	75	150	175
9	125	175	200
11	150	200	225

Table 3

<i>User</i>	<i>Date</i>	<i>X</i>	<i>Y</i>	<i>Value</i>
1	1	12	9	\$3.25
2	1	3	6	1.25
3	1	5	4	2.00
4	2	8	12	1.10
5	2	12	7	1.00
6	0	0	0	0.00
MAXIMAL TOTAL VALUE				

that user's six possible designs, and then allocated that experiment to day 1 if there were sufficient resources left or to day 2 if day 1 were not possible but day 2 was. We then calculated the total value of this use of resources and compared it with the maximum total value. The percentage arrived at is the *efficiency* of that use. A diagram illustrating the distribution of these efficiencies is presented in Figure 1. One can see that it is relatively easy, by random draws, to achieve a level of 60–70% efficiency for this particular problem.¹⁵ Furthermore, it is very difficult to obtain efficiencies higher than 85%. However, in practice, coordination does not occur randomly but is accomplished through the use of institutional arrangements. Let us see how well various institutions perform in this environment.

b. Administrative Procedures

The first institution tested is based on the salient features of many common administrative processes. In these institutions, administrators first ask for information about the resources required and the value of the projects. They then try to optimize the use of the resources on the basis of that information. Unfortunately, the behavior of the users cause the good intentions of the administrators to be undone. The actual realization of these processes is a first-come, first-served algorithm where the users, who arrive early, get to use the facility if their resource requests do not exceed available resources. When the resource request is too big, the experiment is placed in a queue until the resources become available. The important feature of this type of process is that the users must choose the design they wish to use before they know what will be available when they actually arrive. In our experiments, we modeled this institution by asking all users to choose a design and to submit to the administrator. They did this simultaneously and secretly. The administrator then selects a design at random and allocates it to day 1 if the resources are available, to day 2 if resources are available then but not in day 1, and rejects it if not enough

¹⁵Of course, the particular numbers depend on the data we chose for resource use and payoffs. Other data, however, yield similar results. One particular question that remains to be answered, through extensive replication, is the extent of the robustness of our findings. Early replication suggests that they are indeed robust.




Figure 1.

resources are available on either day. On first glance, it might seem that this is the same process used to generate Figure 1, the random assignments, but that would be incorrect. While it is true that the final allocation is decided randomly, each user has the option of purposely selecting a design. One might expect intelligent users to be able to select designs that yield outcomes better than those generated by purely random selection. Their performance of this administrative process is summarized¹⁶ in Figure 2. Early performance is bad. (Fifty-five percent is worse than one would expect from purely random behavior.) Users, apparently not anticipating the full consequences of their actions, would select designs that required a lot of resources and, therefore, increased the probability that they would never get to use the facility. After about three repetitions, the users were able to coordinate implicitly their designs to the extent that they boosted the efficiency of the process to 70%—better than random but not by much. The main problem is that users like 1, 4, or 5, in response to the uncertainty as to whether or not they will be able to receive resources, chose “medium or small” sized designs so that they would “fit” into the available resources. The process caused these users to choose designs that inefficiently utilized the resources. A naive observer would see all the resources being used and all the users being served, but would miss the effect the institution has on the choice of project. That effect created a 30% efficiency loss.

¹⁶Each data point represents the average, across trials, efficiency achieved; that is, actual value divided by maximal value. (In a field trial, we could not measure this because we would not know the payoffs. Here we know them.) In each trial, we repeated the experiment for several periods to try to uncover whether there were learning effects or stability problems with the institutions. As one can see, it did require some learning for the administrative process to outperform even random selection.




Figure 2.

c. Market Processes

The second institution tested was a market organization that consistently has produced efficiencies in the 90% range in laboratory testing in simple environments. The institution is called the "Double-Oral Auction," which is an open-outcry market similar in structure to the commodity pits of the Chicago Mercantile Exchange or the posts of the New York Stock Exchange.¹⁷ For our facility, we implemented this institution by first dividing the resources of the facility among all potential participants.¹⁸ They are then free to buy or sell these units from each other at any mutually agreeable price. All bids and offers are made publicly, and any active bid or offer can be accepted at any time by any participant. Typically, the model that best predicts behavior in this institution is the Law on Demand and Supply. That is, it is usually the case that prices converge to those that equate demand and supply, and the allocation is that predicted by the competitive market model. The data for our environment are contained in Figure 3, where they are superimposed on the data for the administrative institution. The surprising finding is that the Double Auction market performed essentially identically to the administrative organization.¹⁹ In retrospect, the reason is rela-

¹⁷At Caltech we have a computer-based version of this market that allows a somewhat more orderly exchange than one normally sees in the loud and vigorous action of the pits. The Caltech Double Auction system is PC-based. It is described in more detail in Ref. 4.

¹⁸The efficiencies attained by the Double Auction seem to be independent of this initial allocation.

¹⁹"Surprising" is used because, as economists, we tend to expect markets to perform at very high levels of efficiency unless there are significant externalities. In the shared facility testbed we created, there were no externalities, and so we did not expect the very low level of performance we encountered.

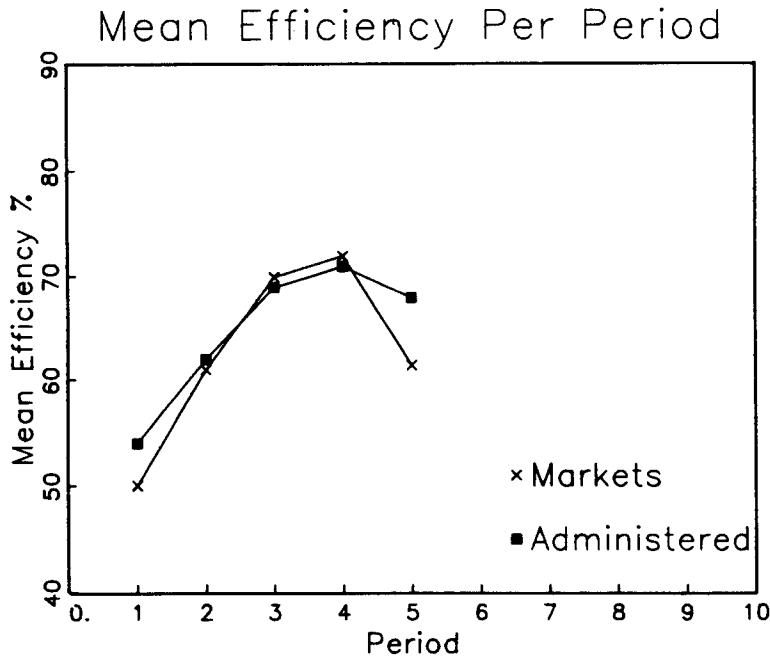


Figure 3.

tively simple: there is no market equilibrium for the testbed environment we created. That is, there are no prices at which demand equals supply and so market prices fluctuate chaotically and, therefore, do not provide the signals necessary to coordinate the activities of the participants in this shared facility.²⁰ The final allocation of resources appears almost random so the numbers are similar to those for the, also random, administrative process.

d. AUSM—A New Design

Having tested forms of the two most common institutional structures for coordinating the activities of the shared facility and found them wanting, we could have taken two alternative paths to try to improve performance; try to fine-tune these; or try something totally different. The failure of the tested institutions was so fundamental that fine-tuning was unlikely to produce any significant changes. Instead, we decided to try to design a new institution. Although we took several approaches, only the most successful are reported here—an auctionlike institution we call AUSM (Adaptive User Selection Mechanism). This institution is a computer-assisted iterative procedure designed to allow multiple users to coordinate their design choices and facility use. It is an institution that would be relatively difficult to implement without the availability of computers, since it uses an optimization algorithm to make choices based on information submitted by the potential users of the facility. The general structure of the

²⁰The technical reason for the nonexistence of an equilibrium lies in the presence of two nonconvexities in the structure of the environment. The details of this structure and an analysis of the behavior of prices and participants can be found in Ref. 2.

Table 4

User	Period 1			User	X	Y	Nominal Value
	X	Y	Nominal Value				
1	10	9	1100	2	7	5	500
3	5	7	1000				
UNUSED	5	7			13	15	
TOTAL VALUE			2100				500

process is relatively easy to understand.²¹ At any point during the process, there is a publicly available *proposal for use of the facility*. This proposal lists who is tentatively scheduled as a user, what resources they can use, and a *nominal value*²² for that package of resources. The list might appear as shown in Table 4.

The idea is that, if no further action is taken, this proposal is to be the allocation of the facility's resources and each user will pay to the facility an amount equal to the nominal value. Further actions are new proposals that are to be considered for inclusion on the list. The rule for adding a new user's proposal is fairly straightforward: A new proposal is included in the proposed use if it will raise the nominal value of the total use. For example, suppose you are user 7 and you submit a proposal for $x = 4$, $y = 6$ in period 2 at a nominal value of 1. Since those resources are available and since adding your proposal to the provisional use plan will increase the total nominal value (by 1), your proposal will be accepted. On the other hand, suppose you submitted a proposal for $x = 7$ and $y' = 6$ for period 1 with a nominal value of 500. That proposal would be rejected because, to provide that level of resources would require removing user 3 from the list with a change in the total nominal value of -500 , a reduction. If the proposal had been submitted with a nominal value greater than 1,000, say 1,050, it would have been accepted and user 3 would have been bumped from the list. The new proposal for use of the facility would have been as shown in Table 5.

Of course, user 3 now has the right to present a new proposal, which potentially can bump either you or some other user. The process proceeds until there are no more proposals forthcoming within a preannounced time limit (e.g., within 1 week of the last proposal) or until a preannounced closing time (e.g., 3 months from the start of the process) is reached. It is possible to manage this process via a computerized bulletin board with remote users.²³ It is also possible to provide a "what if" facility so that users can ask about the consequences of particular bids. After all, the purpose of the procedure is to encour-

²¹The full details of AUSM and instructions for the experiments can be found in Ref. 1. We have also developed a demo-program for IBM-PCs illustrating AUSM and providing an interactive tutor; this is available on request for a nominal fee.

²²I use the phrase "nominal value" to indicate that this need has no relationship to the actual value of the payload.

²³It need not be managed as the usual "art auction" with everyone in the same room being egged on by a professional auctioneer. The purpose of the process is not to encourage competition but to encourage coordination.

Table 5

User	PERIOD 1			User	X	Y	Nominal Value
	X	Y	Nominal Value				
1	10	9	1100	2	7	5	500
7	7	6	1050				
UNUSED	3	5			13	15	
TOTAL VALUE			2150				500

age coordination of the designs and resource use through the provision of information about the various mutual possibilities. AUSM is designed to elicit that information in a way that it is in the interest of each user to provide it correctly. How well does AUSM do this? The data on efficiency for our experiments with AUSM are presented in Figure 4. The results from markets and administrative institutions are included for contrast. There are two major findings. First, we have designed an institution that performs at much higher efficiencies than did the standard institutions. The 80% average efficiency achieved by AUSM is a remarkable improvement over the average efficiencies of 63% for the administrative institution and 66% for the market institution, especially in the light of Figure 1 (the distribution of efficiencies from random selection). The second finding is even more remarkable. Look only at the data from the first periods of each experiment. One finds markets at 50%, the administrative process at 55%, and AUSM at 80% efficiency. One must conclude that not only does AUSM perform better on average, but it achieves its level of performance very quickly and with relatively little prior information required. The goal of providing improved performance through the design of a new institution has been attained. It is an open question whether we can do still better.

5. CONCLUSIONS AND OBSERVATIONS

There are several important ideas that one should consider carefully. First, *institutions matter*. As we change the rules by which we coordinate behavior, we can change the outcomes that will result within any given environment. This means that the possibility exists to design institutions to achieve desirable goals. Second, *the laws of social science should not be ignored*. Any attempted solution to a coordination problem that simply asks users to present the correct information to an administrator (or a computer) who will then calculate the optimal solution generally is doomed to failure. Users will manipulate the information provided in an attempt to achieve more than honesty will, and the calculations based on this flawed information, although picking out the optimum for this information, can easily produce bad outcomes (sometimes worse than if no intervention were made) relative to the true data. This can occur in a way that cannot be measured objectively in the field. In many applications, designing a new computational algorithm is only the beginning; implementation requires that the design be

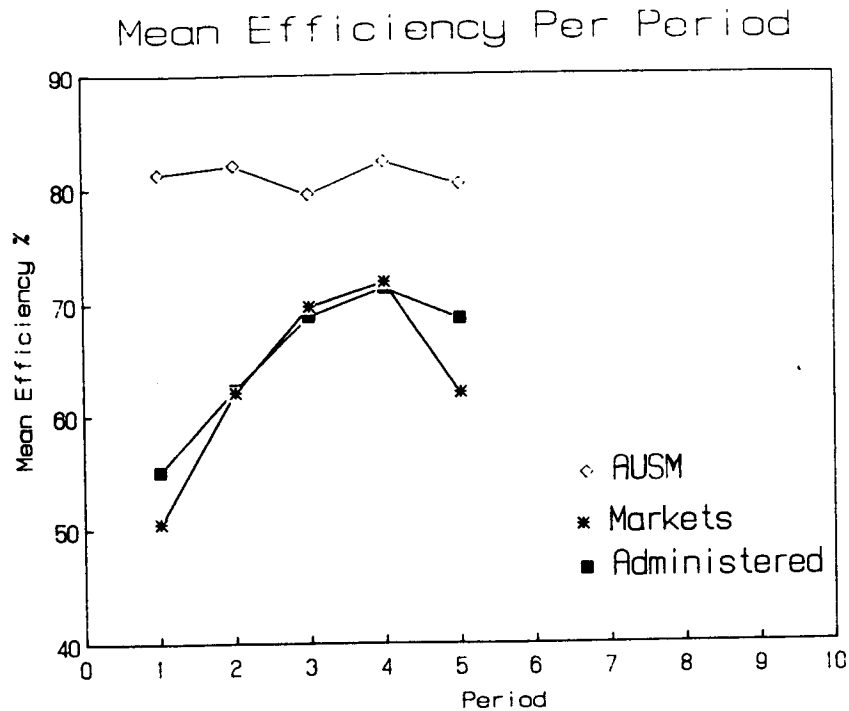


Figure 4.

sensitive to the incentives created at the interface between the algorithm and the users of the system. Third, *new methods in game theory and experimental economics provide a way to study coordination problems systematically*. Game theory provides the theoretical framework within which to design new institutions subject to the constraints imposed by informational and behavioral constraints. Experimental economics provides the testbeds on which prototypes of new institutions can be tried and modified before full implementation. Replicated testing of new and old designs is possible. Arguments and indirect observations can be replaced with hard data. Performance of institutions in classes of environments can be predicted and measured. We can test designs against each other. Finally, *there is an institution that coordinated users of a shared facility better than either markets or standard administrative procedures*. The computer-assisted group support system we call AUSM has been shown to increase the efficiency of use of a shared facility by 33% (from 60% efficiency to 80% efficiency) in test of prototypes. It remains an open question whether there are any better institutions, but the methodology to answer that question now exists. We need no longer depend solely on seat-of-the-pants solutions for complex coordination problems.

REFERENCES

1. Banks, J., Ledyard, J., and Porter, D. The design of mechanisms to allocate space station resources. California Institute of Technology, Social Science Working Paper No. 648, June 1987.
2. Banks, J., Ledyard, J., and Porter, D. Allocating uncertain and unresponsive resources. *Rand J.*

- Econ. 20* (spring 1989), 1-25. Also, California Institute of Technology, Social Science Working Paper No. 680, Sept. 1988.
- 3. Space Station Freedom Resource Allocation and Pricing Policy: Options and Recommendations. Jet Propulsion Laboratory, Pasadena, Calif., D-5684, Sept. 1989.
- 4. Johnson, A., Lee, H.Y., and Plott, C. Multiple unit double auction user's manual. California Institute of Technology, Social Science Working Paper No. 676, June 1988, Revised Sept. 1989.