

The Design of Coordination Mechanisms and Organizational Computing

John O. Ledyard
California Institute of Technology

We provide an introduction to a theory of coordination mechanism design and show how to apply it to an assignment problem. The purpose is to introduce those familiar with organizational computing, but unfamiliar with game theory and economics, to the subject. We also describe briefly how we can test new mechanisms before taking them into the field. Finally, we raise some unresolved research questions.

assignment problem, mechanism design, incentive compatibility,
cooperation, experimental economics

1. INTRODUCTION

I begin with two examples of organizational design problems that I hope will expose some issues of importance to those interested in coordination. The example are highly abstracted from two research problems being studying.

Problem 1. You have just bought a firm with three divisions, each with its own unique management problem. Furthermore, you have chosen three individuals of very different talents to be managers for you and have promised them they can share in the profits generated by the division they manage. Now you must decide who is to manage which division. Each manager best knows their own ability to generate profits in each division and no one else, including you, has that information. You want to assign them so that the total profits of the firm are maximized. You could let them each choose which division they want to

This article relies heavily on the results from an ongoing research project with Mark Olson and David Porter funded by the Flight Projects Office of the Jet Propulsion Laboratory (NASA) through a grant to the Program on Organizational Design at the California Institute of Technology. None of them are responsible for any errors.

Correspondence and requests for reprints should be sent to John O. Ledyard, Division of Humanities and Social Sciences, 228-77, California Institute of Technology, 1201 E. California Blvd., Pasadena, CA 91125.

manage. But what do you do if several choose the same one? Do you flip a coin? You might form a committee of the three and let them vote on an appropriate assignment. But how do you prevent them from wasting a lot of time in these nonproductive meetings? Can you provide a computer algorithm that will solve your coordination problem?

Problem 2. You have been hired by the people of a small island as a consultant. For years the king of the island has decided who gets what and how much. But everyone, including the king, is tired of the political favoritism and the time spent lobbying to protect one's interests. They are also not sure that the goods are getting to those who really can use them. Each thinks they deserve more. They have already considered assigning everyone an equal share of the resources but rejected that process. Equal allocations are not desirable if some islanders hate bananas and others hate coconuts. In a world where different people have different tastes, everyone can be made better off with some trading. An additional problem you face is that those who have been especially favored over the years have claimed they will go to war rather than accept a new system that leaves them significantly worse off. The islanders are widely dispersed and rarely see each other but they can ship goods. They are technically advanced and have a system of networked personal computers. They do not have money nor do they want any: they have made it constitutionally illegal. Can you design an algorithm to solve their problem?

If someone had all the requisite knowledge, access to infinite computing capabilities, and the power to require actions, they could easily solve these problems. The first is an example of an assignment problem: a simple integer programming problem. The difficulty, as we will see in more detail below, is to get the potential managers to correctly provide the necessary information on their abilities. This is called the *incentive compatibility problem*. The second problem is more complex. Even if the information on desired trades is available, computing the appropriate shipping plans, deciding who sends what commodity to whom, can be a major combinatoric problem that takes too much time to solve given the islanders' computing technology. I call this the *computational complexity problem*. Finally, even if one can acquire the information and calculate a desirable solution, one must still make sure everyone is willing to complete the required trades. The islanders must be willing to participate. I call this the *political viability problem*.

Mechanism design, populated mostly by game-theorists and economists, and organizational computing, populated mostly by computer scientists and group psychologists, have developed independently to tackle coordination problems. Mechanism design has matured over the past 20 years by focusing on incentive compatibility and political viability. The analysis has usually been carried out under the working assumption that infinite computing capacity is always available. Any computation required of the individuals or of the system can be instantaneously and correctly completed. Of course, any expert in organizational computing knows this is clearly wrong. For the most part, organizational computing, including decision support and group support systems, has matured by focusing on the potential information processing capabilities of computer systems. The analysis has usually been carried out under the working

assumption that individuals will voluntarily provide all data and take all actions required by the system. People will act like computer algorithm subroutines. Of course, any expert in mechanism design knows this is clearly wrong. It is time to integrate the two research programs so that all realities of the issues are confronted in a sensible way.

This article is intended as an introduction to the field of mechanism design and its potential for application to problems of interest in organizational computing. The theory has become highly developed, the applications are wide ranging, and there are ways to use economics experiments to provide proof of concept. I will describe some recent basic research results, discuss some recent organizational design applications for which we have found these techniques useful, and then I will list some open research problems that require expertise in organizational computing.

2. EXAMPLE OF MECHANISM DESIGN

Let us return to Problem 1. It is one example of a class known as assignment problems. These are organization problems that arise when a fixed number of heterogeneous objects (such as offices, personal computers, secretaries, or starting times at a golf course) must be allocated (usually on a one-to-one basis) to a fixed number of potential users. Each user values these items differently and only that user knows the true values.¹ The organization problem is to provide a mechanism to determine the allocation in a way that maximizes the organization's goals. Let v_{ij} be the profits that manager i thinks will occur if she manages division i . Assume that her profits will be independent of how the others are assigned. Let x_{ij} equal 1 if i is assigned to manage j . Let x_{ij} be 0 otherwise.² Problem 1 is to choose values of x_{ij} , for $i = 1, \dots, I$ and $j = 1, \dots, J$ with $n = m = 3$, to maximize

$$\sum_{i=1}^I \sum_{j=1}^J v_{ij} x_{ij}$$

subject to

$$\sum_{i=1}^I x_{ij} \leq 1$$

$$\sum_{j=1}^J x_{ij} \leq 1$$

and $x_{ij} \in \{0,1\}$ for all i,j .

Remember, however, only i knows the true value of v_{ij} , so you need to acquire that information before solving this problem. But i will share in the

¹An alternative way to think about this is that each user knows the value to the organization from matching up that user with any particular unit, and only that user knows this. The key is that information about the value to the organization of any assignment is dispersed.

²Another possibility is to let x_{ij} be any number between 0 and 1 and to interpret it as the probability of assigning i to j . We also use this approach in this article.

profits of only her division and so she does not necessarily want to solve your problem. Instead she wants to maximize

$$\sum_{j=1}^J v_{ij} x_{ij}$$

subject to

$$\sum_{j=1}^J x_{ij} \leq 1.$$

To see why this creates a problem, let us consider four mechanisms that someone who has not thought much about it might suggest.

2.1 Mechanism A: *The Obvious Algorithm*

This mechanism is straightforward and based on the assumption that people will provide all information requested. A computer programmer is hired to provide an algorithm for solving the assignment problem as stated earlier. The only data required by the program are the values of $v_i = (v_{i1}, \dots, v_{ij})$ which are known to each i . After the program is written, you simply ask each manager i to provide values of v_i to the computer. Then you run the algorithm, check the result, and assign i to j if and only if $x_{ij} = 1$ in the solution. What would you do if you were individual i ? The answer is easy. You would first figure out for what j the value of v_{ij} is largest. Let us say v_{i1} is the largest value. You now know that you want to be assigned to division 1. How do you accomplish this? You simply make sure that your *reported* value of v_{i1} is larger than any other j 's report about v_{i1} and that your *reported* values of v_{ij} , for $j > 1$, are less than any other's reports about their respective values.³ (You claim exceptional expertise in managing division 1 and predict huge losses if you are placed in charge of division 2 or 3.) Is this a problem? Not if each individual manager is indeed particularly suited to the management of one and only one division and that division is different for each manager. However, if abilities are not so specific and differentiated or if some divisions are simply more difficult to manage and produce lower profits no matter who is the manager, then this mechanism will not necessarily produce the assignment that maximizes the organization's profits. The naive assumption that each manager will report the true value of her v_i will lead to suboptimal results. For example, suppose the profits that will occur to three divisions under each of three managers are as follows:

Division	Manager 1	Manager 2	Manager 3
a	100	60	75
b	200	50	100
c	50	10	10

³One might argue that this type of behavior will be found out when the next profit report is due and, with a serious enough penalty, could be prevented. But it is easy to modify this example slightly and more realistically to have profits determined jointly by the manager and some exogenous, random forces. In this case the manager who reports too high a capability can claim that the lower than expected profits were the result of bad luck and not to her ability. This will not work in the long run but can be successful in the short term. There are plenty of obvious examples.

There are six possible assignments. The optimal assignment is to place 1 in charge of b , 3 in charge of a , and 2 in charge of c . We denote this by (3,1,2). It yields a profit of 285. Only manager 2 knows the entries in column 2. If the other two are honest and nonmanipulative, 2 can become the manager of division b by reporting values of 10 for a , 210 for b , and 0 for c . The algorithm will then choose (1,2,3) thinking profits will be 320 but they will, instead, be 160. Note that both 1 and 3, as well as the organization, have lost by 2's maneuver. Of course, 1 and 3 can and will retaliate. If there is an upper bound on possible reports, say H , and a lower bound no reports, say L , then in equilibrium 1 and 3 will claim to be able to generate profits of H if they are assigned to division b and L if assigned to a or c . Manager 2 will claim H for division a and L for b and c . The algorithm will assign 2 to a but will have to randomly assign 1 and 3. The profits will be either 270 or 210, with an expected value of 240, which is 84% of the optimal profits of 285.

Can we do better?

Mechanism B: *The Committee*

Suppose we let these managers meet as a committee and choose the assignment themselves. Suppose, in addition, we require only a simple majority, two in the case of three divisions, to be in favor of an assignment to adopt it. What can we expect to happen? Remember again that only each manager knows her own likely profits. Look at the three-division example above. There are six possible assignments but only one of them has the property that there is no alternative a majority prefers. The assignment that will be chosen gives division a to 2, b to 3, and c to 1. Both 2 and 3 will vote in favor of this over any other suggestion. It yields a profit of only 210. Majority rule has allocated the best manager to her worst alternative. The obvious algorithm is better since at least half the time it yields a higher valued assignment. Profits using the committee are barely above the minimum of 175.

Is there any other possibility?

Mechanism C: *The Lottery*

A lottery is a mechanism that is sometimes used to allocate dormitory rooms and rock concert tickets. Individual names are drawn at random and then allowed to select their first choice from the alternatives that are yet unassigned. How would that work here? If manager 1 is selected first, she chooses division b . If manager 3 is selected second, she would choose division a . Manager 2 then has no choice and is assigned to division c . Since the order chosen is random, we can calculate only the expected profits. There are six possible draws. Three yield 210, two yield 270, and one yields the optimum of 285. The expected profits are thus 243. For these parameters, the lottery always does at least as well as majority rule and beats it half the time. For these parameters it is also just a little better than the obvious algorithm since it beats it one-sixth of the time by selecting the optimal assignment instead of one valued at 270. However, different parameters yield different rankings. For example suppose the profits for manager 2 in division a

are 100 (instead of 60). Then, the obvious mechanism is better since it yields a payoff half the time of 310 and half the time of 250 for an expected profit of 280. The lottery yields, for these new parameters, 310 for two of six times, 250 for two of six times, 285 for one of six times, and 210 for one of six times. The expected payoff of 269.

What might an economist say?

Mechanism D: Marketlike Mechanisms

Another mechanism that is sometimes considered is one that tried to mimic the marketplace. It gives a number of tokens to each manager and then lets them use those to bid for an assignment to a particular division. For the assignment problem in which each manager gets only one division, these marketlike mechanisms are equivalent to the obvious algorithm when reports are limited⁴ to being between L and H . In equilibrium, purposive agents will always bid all of their tokens on their most preferred division and zero on the others. This is no different than reporting H for the most preferred slot and L for all others.

At this point we could continue to create mechanisms and evaluate their performance, hoping to find the best one. But there are an infinite number of possibilities, some of which only replicate the performance of others, and a finite amount of time. We also have to find some way to deal with the fact that our evaluation of the three mechanisms has depended on knowing the true values of v . In practice, the designer may only know that the entries in the matrix are bounded between m and M . How should we then evaluate the possibilities? The discipline of mechanism design addresses these issues by providing a systematic approach to organizational design problems. It provides a mathematics of organization design that allows us to analyze structures and to predict performance. Let us turn to that formal framework next.

3. ABSTRACT APPROACH TO MECHANISM DESIGN⁵

An abstract description of what is studied in the modern theory of mechanism design includes (1) the set of *environments*, denoted by E with elements e , which describes the problem and the range of possible parameter values including, if necessary, the structure of knowledge and beliefs about the likelihood of any particular environment; (2) the set of *outcomes*, denoted by X with elements x , the choices that the mechanism is to make; (3) the *mechanism*, denoted (M, g) where M is the language of communication with the managers and g is the rule $g(m) = x$ which determines what outcome is selected depending on what the

⁴If there is no limit on credible reports, then there will be no equilibrium for the obvious algorithm. Each manager will try to choose the largest integer and report that as the potential profits for their most preferred division. There is, of course, no largest integer.

⁵This section and the next are somewhat abstract, but provide the theoretical foundations for the more applied discussions later on. The nonmathematical reader can jump to Section 5 with only a small loss in the train of thought.

managers communicate⁶; and (4) the *behavior*, denoted $b(e, g) = m$, a prediction about what the individuals in the organization will do in each environment when confronted with a new mechanism.

The mechanism tells us what is to be done if individuals report m . g can include very complicated iterative processes as well as very simple one-report versions. The behavior rule tells us how these reports will occur if the mechanism is g and the true environment is e . b can include behavior which is independent of other's actions or it can describe the equilibrium of some complicated sequential game. The *performance* of the mechanism in the environment e is then the composition of the design and behavior, or $g(b(e, g))$.

How is this to be evaluated? Generally there is an evaluation function, $U(x, e)$, which tells the designer how to value particular outcomes from the organizational point of view in each environment. In our assignment problem, $U(v, x)$ was the profits yielded for the true values v when the assignment is x . The performance of (M, g) in e given the behavior b is evaluated as $U(g(b(e, g)), e)$. But different mechanisms may be better in different e . The natural Bayesian approach is to adopt prior beliefs about the probability of any e in E represented by a cumulative distribution function $F(e)$, and rank mechanisms by the expected valuation, $\int U(g(b(e, g)), e) dF(e)$.

To summarize, modern mechanism design takes E , X , b , U , and F as given and tries to describe and characterize the solutions to the problem: choose (M, g) to maximize $\int U(g(b(e, g)), e) dF(e)$ subject to the constraints due to individual behavior (called incentive compatibility), the constraints due to computational limitations (called computational complexity), and constraints due to the fact that, generally, new organizations that replace old ones will be unsuccessful unless everyone is better off (called political compatibility).

Formally, the designer's problem is to choose (M, g) to maximize $\int U(x(e), e) dF(e)$ subject to

- | | |
|--------------------------------|--|
| A. Incentive compatibility | [i.e., $x(e) = g(b(e, g))$ for all e in E] |
| B. Computational compatibility | [i.e., g can be computed for all m] |
| C. Political compatibility | [i.e., $v^i(g(b(e)), e) > v^i(g^0(b(e)), e)$ for all i and e in E where v^i is i 's own evaluation of the outcomes x in e and g^0 is the mechanism to be replaced] |

The obvious problems in putting this into practice are

1. What is the full set of feasible mechanisms (M, g) ?
2. What is the right behavior model b ?

⁶In our example of the allocation of managers to divisions, the environment included the matrix of profits, the fact that each manager only knew the values in her own column in that matrix, and the bounds on profits, H and L ; the outcome was an assignment; the obvious algorithm as a mechanism is described by letting $M = \{v\}$ and letting g be the algorithm that chooses the optimal assignment for each v .

3. What are computational compatible mechanisms?
4. Can we describe the set of politically viable mechanisms in a more operationally satisfactory manner?

4. WHAT IS THE SET OF ALL INCENTIVE COMPATIBLE MECHANISMS?

A very helpful shortcut has been developed that allows us to begin to get a handle on the set of incentive compatible mechanisms, those that satisfy model A. It does not allow us to describe all of them, but it does give us an upper bound on the possibilities for performance. This shortcut relies on some game-theoretic principles of behavior and is known as the *revelation principle* [1-3]. It states: Given the appropriate behavior and given full information processing and computing capabilities, the performance of any mechanism can be replicated by a *direct incentive compatible* mechanism.

A direct mechanism is one in which all individuals simply report their information about the environment and g picks one x for each reported e . An incentive compatible direct mechanism is one in which each individual finds it in their own interest to report the truth. That is, $b(e, g) = e$ for all e . The revelation principle then provides that given b and any mechanism (M, g) with performance $g(b(e, g))$, there exists a direct mechanism (E, h) such that $e = b(e, h)$ and $h(e) = g(b(e, g))$ for all e in E . A plausibility argument goes as follows. If behavior is purposive, then $b(e, g)$ will have the property that each individual is responding as best as is possible for them given the private information they have. Now consider the following direct mechanism: have each agent report her information to the designer and have the designer then compute the strategy, for that individual, using the original mechanism. If the individual incorrectly reports the information then the designer will compute something other than what was originally the best for the true situation. So it is in the interest of the agent to report honestly: the direct mechanism $h(e) = g(b(e, g))$ is incentive compatible and obviously duplicates the performance of g given b .

Using the revelation principle, we can restate the design problem as follows: Choose (E, h) to

$$\text{maximize } \int U(h(e), e) dF(e)$$

subject to

- A. $e = b(e, h)$ for all e in E
- B. Computational compatibility
- C. Political compatibility

With this device we have significantly narrowed the domain over which we must search for the desired coordination mechanism. In fact, for many problems, such as the assignment problem, we have converted our search into an analytic and solvable maximization problem. However, we have done so at the

cost of ignoring some potential frictions. If the set of direct mechanisms includes some that are not computationally compatible then our shortcut will rule them out although there may have been some less direct computational compatible mechanisms which would have yielded the same outcome. So if we impose computational compatibility in our problem, after using the revelation principle, we may miss desirable mechanisms. Second, in stating the revelation principle we assumed that "the appropriate behavior and full information processing and computing capabilities" were present. This really requires purposive game-theoretic agents who act as if they have extremely general and effective computing capabilities. It is really an open empirical question as to when this is valid. We discuss the empirical issues in Section 6.

Now that we have some of the theoretical foundations displayed, let us return to our example, Problem 1, and see what these principles let us do.

5. ASSIGNMENT PROBLEM—THEORY

We have discovered in the last two sections that, under the appropriate empirical conditions on individual behavior, we can systematically search for a best mechanism to solve the assignment of managers to divisions by narrowing our domain of search without loss of generality. In particular, if the assignment algorithm can be solved for all possible matrices of values, then the solution to Problem 1 is equivalent to the following:

1. Solve the optimal assignment problem for the range of possible profit matrices using the obvious algorithm with some additional constraints, called incentive compatibility conditions, which ensure that each manager will report correctly.
2. Ask each manager to report the values for her column of the matrix.
3. Assign the managers according to the solution to item 1 above.

The theory establishes that the best mechanism possible (given the incentive, computational, and political constraints on the designer) will be found this way. Formally one solves for $x(v)$ for all possible v to

$$\max \int \sum_{i=1}^I \sum_{j=1}^J v_{ij} x_{ij}(v) dF(v)$$

subject to

$$\sum_{j=1}^J x_{ij}(v) = 1 \text{ for all } i$$

$$\sum_{i=1}^I x_{ij}(v) = 1 \text{ for all } j$$

and

$$\sum_{i=1}^I v_{ij} x_{ij}(v) \geq \sum_{i=1}^I v_{ij} x_{ij}(v/w_i)$$

for all v, w_i where (v/w_i) replaces v_i with w_i . The last inequalities are the incentive compatibility constraints.⁷ The mechanism is then simply: ask each manager i for her v_i and choose the assignment $x(v)$. This still leaves us a difficult maximization problem to solve but we can use more theory. In particular, Mark Olson [4] has found that

Theorem (Olson)

For the assignment problem, $x(v)$ satisfies the incentive compatibility constraints if and only if x depends at most on the ordinal properties of v .

That is, any incentive compatible mechanism need only, indeed can only, ask the managers, i , for a ranking of the profits expected from each division, j . This means that the obvious algorithm and the marketlike mechanism are not incentive compatible, which we already knew from our analysis in Section 2. However, they are, as the revelation principle states, equivalent to another incentive compatible direct mechanism. Conversely, the lottery is incentive compatible since it is equivalent to asking each manager for a ranking and then selecting one at a time, randomly, and assigning each to their top-ranked alternative from those divisions remaining.^{8,9} Using Olson's result we can eliminate the best inequalities in our problem and look only at $x(v)$ that depend on rankings of v_i .

6. PROOF OF CONCEPT OR HOW CAN WE BE SURE?

How can we be sure of a design based only on theoretical analysis? One way is to create an organizational design test bed similar in spirit to the methods aircraft designers use with their wind tunnels and havel architects with their towing tanks. At Caltech, our organization design test bed is the Experimental Economics Laboratory. This is not the place for an extensive description of experimental economics methodology, but perhaps I can give you enough an introduction to pique your interest.¹⁰

What is important to understand is that economic experiments are a cross between feasibility studies, using simulated participants or real subjects (does the software work?), and very complicated field studies. The key aspect is that we can control the incentives of the participants in the laboratory through the technique of value inducement (see [8]), we can control the structure and timing of information that participants have about the environment in the laboratory, and we can control the mechanism. In the field, only the mechanism is directly

⁷These constraints require, in the language of game theory, that the true value of v_i be a dominant strategy response to the request for information. That is, manager i must want to respond with the truth no matter what the other managers report.

⁸Olson and others call this the Serial Dictator mechanism for obvious reasons.

⁹The committee is also incentive compatible if for the true environment, there is a single assignment which a majority prefers to all others (a condorcet winner). Otherwise, incentive compatibility may be absent.

¹⁰Two recent books on the subject are *Handbook of Experimental Economics* [5] and *Experimental Economics* [6]. For a description of this methodology applied to the assignment problem consult [7].

controllable: in feasibility studies no attempt is made to control incentives. A bonus we get from control is that we can measure performance characteristics that cannot be measured directly in the field.

How does this work? Let me illustrate with the assignment problem. Subjects are recruited to participate in an economic experiment, testing an organization design for the assignment problem, in which they are paid based on their decisions.¹¹

Each subject is given a description of the generic group decision problem, in this case that I subjects are to be allocated to J slots. Each subject is also given a description of what that subject will be paid, in dollars, at the end of the experiment, for each possible outcome of the group decision process. For example, subject 1 may be told she will be paid \$10 if she is assigned to slot 1, \$20 if she is assigned to slot 2, and \$1 if she is not assigned to any slot. Each subject knows only her own payoffs. In this way we have provided the subject with the strong financial incentive to act as if she is in an environment where $v_i (= 10, 20)$ and the status quo—relevant for political viability—is 1. We now can confront the collection of subjects with different mechanisms and see what happens. Usually a number of trials are completed with the same subjects, payoffs, and mechanism to make sure the results are reasonably stable.

It is important to recognize two key advantages of this methodology: (1) we have the right model of the human participants and (2) we can measure such things as the values obtained. The first is important because studies that use robots (simulated humans) to compute behavior in group systems are then highly dependent on getting that model correct. The second is important because we can then compare the coordination mechanism on grounds other than ease of use, speed of interaction, etc. In particular, we can measure how well the mechanism actually performs in achieving the goals set for it. For the assignment problem, as it was originally stated, the goal was to maximize the sum of the values to the individuals of the final assignment. Since we, as experimenters, have chosen those values and know the final assignment, we can exactly measure the value of that assignment. In field studies or in lab studies where the subjects are not paid on the basis of the outcomes, there is absolutely no way to do that.

Two disadvantages of this methodology are (1) it is expensive and (2) we may be missing certain scale effects. The former depends on how much control the experimenter wants. Paying subjects \$2,000 and \$1,000 instead of \$20 and \$10 certainly increases the likelihood that their induced valuations are indeed their true valuations for the experiment. Paying them \$0.02 and \$0.01, instead, would lead to the likelihood that they would react in this environment to other incentives that we cannot measure, or even identify. The latter problem occurs even in engineering studies using wind tunnels and towing tanks. There is no guarantee that an America U.S. Cup racing sailboat design that performs at a

¹¹At the Caltech lab this room has 20 IBM computers linked by a token ring. Other labs may have a different configuration of machines. In one set of experiments we have used a collection of computers around the globe connected through Internet.

very high level in the tank, where the model is one-tenth scale, will not just capsize when built to full scale.

One measure of particular interest to mechanism designers is efficiency; that is, the valuation achieved by a mechanism relative to the maximum achievable. The total value received measures the amount of the benefits the decision process has created to be distributed. Efficiency measures how much of the maximum benefits have been *discovered* by the mechanism. So all other things being equal we would want to design mechanisms that produce high efficiencies. The problem, of course, is that there is no reason to expect the efficiency of a mechanism to be constant across environments. So to be sure about the relative performance one must check it out under a number of possible situations. Let us look at one study that did this for the assignment problem.

Olson and Porter [7] tested two of the mechanisms we have described in Section 2 in two very different environments called the high contention environment and the low contention environment. The difference in the two environments was the degree to which several agents wanted to be assigned to the same slot. If each agent's most preferred slot was different from all the others', there is 0 contention. If all agents have exactly the same ranking for all slots, there is a lot of contention, say 1. Olson and Porter used an environment with six slots and six agents with low contention, but not 0, and an environment with six slots and eight agents with high contention, but not 1. The two mechanisms considered were the lottery and a marketlike mechanism. Each mechanism was tested in each environment using two different subject groups for 20 trials each. The average amount of money paid each subject for 45-60 minutes' effort was around \$14.

The efficiency achieved on average during the final 15 trials in each experimental session is detailed in Table 1.

What is noticeable is that the marketlike mechanism really does do better than the lottery in both situations. Olson and Porter do not explain why this occurs, but it appears to be because the lottery allocates primarily on the basis of one's first choice slot while the marketlike mechanism allows those who do not get their first choice to sometimes get their second choice. The theory, as it stands, would not have uncovered this phenomenon. It remains an open question whether one can find any other mechanism that would produce higher efficiencies.¹²

¹²If one were willing to allow subjects to use their own money to bid for allocations (then, for example, someone who will be paid \$20 could bid up to \$18.99 and still be better off than getting no slot), then one can easily improve the efficiency of the allocations achieved. Olson and Porter show that one kind of auction attains 98% and 99% efficiencies in low and high contention environments. The problem with these mechanisms is that the individuals do not get to distribute the higher values attained, since they pay some of it away just to get an assignment. In both environments they end up individually worse off using cash to bid than they do using chits to bid even though efficiency is higher with cash. In the high contention environment the subjects are 300% better off when cash is not used. The implications for political viability ought to be obvious. Because of the high contention the sellers of the slots (in this case, the experimenters) are able to command very high prices and extract most of the benefits. In low contention environments this is not possible and subjects come out about the same with or without cash. (Recommendations to sell off scarce parking places at universities generally falter for these reasons.)

Table 1
Average Efficiencies

	Low Contention, %	High Contention, %
Lottery mechanism	92.1	81.4
Marketlike mechanism	94	89.4

7. OPEN RESEARCH ISSUES

Remember the research program is to find a systematic methodology to solve the optimal mechanism design problem: choose (M, g) to maximize total organizational benefits subject to incentive compatibility, computational compatibility, and political compatibility. There are at least three issues that need more research.

1. What mechanism g actually satisfies the constraints? We saw that the revelation principle has the potential to simplify this question with respect to incentive compatibility. To do the same for computational compatibility seems to be a task that is ideally suited for those who study organizational computing. Involved are issues of heuristics versus full optimizing algorithms, memory constraints, presentation effects, etc.
2. What is behavior b ? Is the revelation principle valid as an empirical issue? What is the form of individuals who have bounded computational capabilities? How does this interact with the computational aspects of the mechanism? How complex must a problem be before game theory loses its predictive ability?
3. How do we deal with political compatibility? In many cases, the designer does not know how participants value the current mechanism that is to be replaced by the new design. Thus, a mechanism must be designed to either reveal that information or to guide the transition to a new mechanism that satisfies the political viability constraint. Those in organizational computing that worry about why new and exciting group decision support systems are not being adopted very rapidly should recognize the importance of this transition problem. Old vested interests will fight adoption unless they can be assured they will be at least as well off as they are now. Increases in efficiency are not sufficient to lead to new adoptions; the distribution of those gains is as important.

There is now a framework within which to carry out our very ambitious, but exciting, program: to create a methodology and technology that will make the design of organizations as straightforward as the design of buildings. The research and applications are just beginning. Organizational computing, game theory, and experimental economics will all play fundamental roles, but I suspect the biggest gains will accrue to those whose research includes more than one of these approaches.

REFERENCES

- [1] R.B. Myerson, "Mechanism Design by an Informed Principal," *Econometrica*, Vol. 51, pp. 1767-1797, 1983.
- [2] A. Gibbard, "Manipulation of Voting Schemes: A General Result," *Econometrica*, Vol. 41, pp. 587-602, 1973.
- [3] M. Harris and A. Raviv, "Allocation Mechanisms and the Design of Auctions," *Econometrica*, Vol. 49, pp. 1477-1499, 1981.
- [4] M. Olson, "Dominant and Nash Strategy Mechanisms for the Assignment Problem," Social Science Working Paper 770, California Institute of Technology, Pasadena, 1991.
- [5] J. Kagel and A. Roth, eds., *Handbook of Experimental Economics*. Princeton, NJ: Princeton University Press, 1993.
- [6] D.D. Davis and C.A. Holt, *Experimental Economics*. Princeton, NJ: Princeton University Press, 1993.
- [7] M. Olson and D. Porter, "Assignments with and without Transfers," *Economic Theory*, Nov. 1993.
- [8] V. Smith, "Experimental Economics: Induced Value Theory," *American Economic Review*, Vol. 66, pp. 274-279, 1976.

Journal of Organizational Computing

Editor: Andrew Whinston, *University of Texas at Austin*

The **Journal of Organizational Computing** is a quarterly journal that publishes original research articles concerned with impacts of computer and communication technology on organizational design, operations, and performance. It serves as a forum for stimulating and disseminating research into the implications of these technologies for organizational structure and dynamics, the technological advances needed to keep pace with organizational changes, and emerging tech-

nological possibilities for improving organizational productivity.

Theoretical, experimental, and survey research are all appropriate to the **Journal**. The **Journal's** major editorial areas are computer-supported cooperative work (CSCW), groupware, computer modeling of organizations, investigations of organizational systems issues such as computer-aided coordination and organization learning, economics of organizational computing, and behavioral studies of organizational computing.

If you have enjoyed reading the **Journal of Organizational Computing**, please pass this subscription order on to your library or a colleague.

Please enter my 1993 subscription for the **Journal of Organizational Computing**. I understand that journal subscriptions must be prepaid, and that there is a \$20.00 charge for postage and handling outside the U.S. and Canada.

Name _____
 Address _____
 City/State/ZIP _____

This is a:
 Personal Subscription/\$50.00
 Institutional Subscription/\$115.00

Please return this order form to the Subscription Department.



ABLEX Publishing Corporation
 355 Chestnut Street
 Norwood, New Jersey 07647
 (201) 767-8450 / FAX (201) 767-6717