

ORGANIZATION STRUCTURE

Thomas Marschak

This Chapter concerns formal models of organizations that regularly acquire information about a changing environment in order to find actions which are appropriate to the new environment. Some or all members of the organization are specialists. Each of them privately learns something about a particular aspect of the new environment. The organization operates a *mechanism*, which assembles relevant pieces of the specialists' private observations and uses the assembled information to obtain the desired new action. The mechanism has various informational costs and they are measured in a precise way. The research seeks to characterize mechanisms that strike an appropriate balance between informational cost and the performance of the mechanism's final actions. As costs drop, due to improved Information Technology, the properties of good mechanisms, and hence the structure of the organizations that adopt them, may change. The Chapter starts by examining research in which the organization's members reliably follow the mechanism's rules and so incentives are not an issue. It then turns to research in which each member is self-interested and needs an inducement in order to make the informational efforts that the mechanism requires. A number of unmet Research Challenges are identified.

I. INTRODUCTION

This chapter concerns organizations that acquire information about a changing environment in order to take appropriate actions.

The term “organization” is used in many disciplines and is applied to many groups of persons. The term covers government agencies, markets, entire economies, firms, nonprofit institutions, the users of the Internet, the firms in a supply chain, and so on. We shall be concerned with models of organizations where some or all members are specialists: each of them privately learns something about a particular aspect of the organization’s newly changed environment. The organization then adjusts its actions. It seeks new actions that are appropriate to the new environment. It *could* find them by “Direct Revelation” — collecting in one central place *all* the information about the changed environment that the members have privately collected, and using it to obtain a new action. But that would be wasteful: much of the transmitted information would not be needed. Instead, the organization seeks to balance the benefit of appropriate actions against the costs of learning about the current environment, transmitting some of what has been learned, and using the transmitted information to choose new actions. Advances in Information Technology (IT) may reduce those costs, and that may change the structure of successful organizations. *In all the models discussed in this chapter, there is some precise measure of one or more informational costs.*

The members of our modeled organization might be totally self-interested, or they might behave like loyal and selfless team members (or well-programmed robots), choosing their informational efforts and their actions so as to contribute in the best way to a common goal. In either case we consider some appropriate measure of the organization’s “gross” performance, before informational costs are subtracted. For the self-interested traders in a market, gross performance might be closeness to Pareto optimality or perhaps the sum of individual utilities. For a multi-divisional firm whose divisions are loyal and selfless, gross performance might be total profit. That might also be our gross performance measure when each division pursues its own interests. We are concerned with the structure of organizations whose gross performance is high given the informational costs they incur.

The research discussed here approaches the problem in the style that is usual in economic theory: a formal model, with well-defined concepts, is studied, in order to find the conditions under which certain interesting conjectures can be shown to hold. Typically the conjectures arise from certain empirical observations, which suggest loose claims. When it comes to organizations and information, empirical work is hard and scarce, the issues are complex, and the gap between a general empirical claim and a tractable formal model might be very large. Consider for example, the following loose claim which appears to date back almost sixty years, to the time when dramatic advances in Information Technology first loomed on the horizon: *As IT advances, firms (and perhaps other organizations) will become “flatter”: middle management will fade away, and decisions will be made centrally, at the top of a flat hierarchy.* (That paraphrases the conjectures in Leavitt and Whistler, 1958). Or consider a much more recent and quite different conjecture: *As IT advances, firms will find it advantageous to adopt decentralized modes of organization, in which unit managers are given wider authority.*¹ Finally, consider a classic claim in economics, very much older than the preceding two claims: *For organizations*

¹That conjecture is consistent with the following statement in T. Bresnahan, E. Bryjolfsson, and L. Hitt (2002), which performs a careful analysis of IT adoptions and organizational characteristics in a variety of industries: “IT use

whose task is to allocate resources efficiently, the only informationally cheap allocation mechanisms are those that use prices, since the “amount” of information transmitted in such mechanisms is exactly what is needed to find an efficient allocation; in any mechanism that achieves the same thing without prices, the information transmitted exceeds what is needed.

The formal modeler who wants to explore these claims faces major challenges. Terms have to be defined precisely, while still capturing informal usage. (How, for example, should we define “decentralized”?). The informal claim has to be stated in a precise manner. A way of showing that appropriate (and plausible) assumptions imply the formalized claim has to be developed. When the modeler is finished, his demonstrated propositions may be far from those that existing empirical work can support or refute. Even so, the *future* empirical study of organizations might be usefully guided by some of the modeler’s results.

In the formal models which we consider the organization uses a *mechanism* to obtain appropriate new actions once its members have learned about the organization’s newly changed environment. The mechanism uses *messages* and it has informational costs, which are measured in a precise way. We are particularly interested in informationally cheap mechanisms, and we often take the viewpoint of a “designer” whose task is to construct an informationally cheap mechanism that generates actions meeting the organization’s goals. There is a vast literature on mechanisms, but only a small part of it makes informational costs explicit. That is the part with which we shall be concerned. We interpret the term “organization structure” as referring to the mechanism the organization uses, or to certain properties of the mechanism. Some examples of structural questions: Are the mechanism’s message flows arranged in a hierarchical pattern? Are the actions it generates chosen by one central person or are the action choosers dispersed? Is each external environment variable observed by a single person, who is the sole specialist in that variable, or is each variable observed by several persons?²

We discuss past research as well as new paths.³ A number of unmet *Research Challenges* are identified. In Section 2 we consider organizations whose members reliably follow the rules of the mechanism which the designer has constructed, whatever those rules may be. The designer is not concerned with incentive issues. In Section 3 the members become self-interested and the designer has to take incentives into account. In Section 4 we look very briefly at some formal models in which the primitive terms are no longer “external environment” and “organizational action”. Section 5 offers a quick retrospective impression.

2. GOALS, MECHANISMS, AND INFORMATIONAL COSTS: THE “INCENTIVE-FREE” CASE, WHERE INDIVIDUALS OBEY THE DESIGNER’S RULES WITHOUT INDUCEMENT

Consider an organization composed of n persons. The organization confronts a changing environment $e = (e_1, \dots, e_k)$. The set of possible values of each e_j is denoted E_j . The set of possible values

is more likely to be effective in organizations with a higher quality of service output mix, decentralized decision-making, and more skilled workers”. See also T. Bresnahan, E. Brynjolfsson, and L. Hitt, 2000.

²There is a literature which refers to the former case as the “U-form” (unitary form) and the latter case as the “M-form” (multidivisional). See, for example, Harris and Raviv, 2002 and Creti, 2001.

³Some new (unpublished) results are reported in Section 2.7 (on “speak-once-only mechanisms”) and in Section 3.3 (on networks of self-interested decision-makers, who bear the network’s informational costs).

of e is denoted E and is the Cartesian product $E_1 \times \cdots \times E_k$.⁴ In many settings we will have $k = n$, and e_j will describe Person j 's *local environment*, which he privately observes.

At any moment of time, the organization has in force an action a . The action is a vector (a_1, \dots, a_m) . The set of possible values of each a_ℓ is denoted A_ℓ . The set of possible values of $a = (a_1, \dots, a_m)$ is denoted A and is the Cartesian product $A_1 \times \cdots \times A_m$.⁵ In many settings m will equal n and a_ℓ will be the action variable for which Person ℓ has responsibility.

2.1 Two general frameworks for judging the organization's actions

2.1.1 First framework: there is a performance function on $A \times E$

Suppose there is a real-valued function F on $A \times E$ which measures the organization's *performance*, when the current environment is e and the action in force is a . The organization, or its designer, wants $F(a, e)$ to take high values. If e never changes and is known to everyone, then the problem reduces to that of finding, once and for all, an a in A which maximizes F for the perfectly known and unchanging e . We may study that optimization problem, and it may be challenging. But if we do so, we strip away the organizational-design issue, and the question of how information about e might be obtained and who might obtain it. Accordingly, we shall suppose that e varies. We are then in a position to define the organization's task. Its task is to learn something about the current e , so that it can find an appropriate action a . It seeks to strike a good balance between the costs of learning about e and the value of $F(e, a)$ that it achieves when it chooses an action a that is best, given the information about e that it has collected. (Here F has to have properties guaranteeing the existence of a best action). A *mechanism* is used by the organization to fulfill its task. *We defer, until 2.2 below, formal definitions of "mechanism", and of the informational costs associated with a mechanism.*

Note that this framework covers the case where each component of e is a signal, privately observed by some person, about a random variable x , whose probability distribution is common knowledge. The true value of x , which does not become known until after the action has been taken, determines a payoff $W(a, x)$ for every action a . Then we may let $F(a, \bar{e})$ be the conditional expectation $\mathcal{E}(W(a, x) \mid e = \bar{e})$.

2.1.2 Second framework: for each environment e in E , there is a satisfactory (goal-fulfilling) set of actions.

Suppose now that for every e in E , there is a set of *satisfactory* actions (a subset of A), denoted $G(e)$, where G may be called the *goal correspondence*. Any action in $G(e)$ meets the organization's goal for the environment e . The organization, or its designer, asks: is there a mechanism which generates a satisfactory action for every environment e in E , and what are the informational requirements of such a mechanism?

In some settings, any action in $G(e)$ has to meet certain constraints, which are defined by e . In other settings one starts by defining the performance function F of the first framework. Then an

⁴There are settings in which one has to relax the requirement that E be the Cartesian product. For example, each e_j may be a parameter which identifies a production technology in a location j , i.e., it identifies the input/output bundles that are feasible in that location. It may be that because of externalities between the locations, some k -tuples belonging to the Cartesian product $E_1 \times \cdots \times E_k$ cannot occur.

⁵Again, in some settings one would want to relax the requirement that A be the Cartesian product

action a belongs to $G(e)$ if and only if $F(a, e)$ is within some given distance, say δ , of its maximal value. Thus for any fixed \bar{e} in E , we have $G(\bar{e}) = \{a^* \in A : |F(a, \bar{e}) - F(a^*, \bar{e})| \leq \delta \text{ for all } a \in A\}$.

2.1.3 Three examples

First example: an organization that provides health services Suppose each person $i \in \{1, \dots, n\}$ is a diagnostician. Each week he examines a group of T_i patients (who are examined by no one else). For the t th patient, he correctly determines a diagnosis category, say d_t^i . Thus the week's local environment for i is $e_i = (d_1^i, \dots, d_{T_i}^i)$. The organization's action for the week is an assignment of a treatment modality r_t^i to each patient t inspected by diagnostician i . (We do not specify who chooses the modality). The possible modalities comprise the finite collection Q . So the organizational action a belongs to the set

$$A = \left\{ \left((r_1^1, \dots, r_{T_1}^1), \dots, (r_1^n, \dots, r_{T_n}^n) \right) : r_t^i \in Q \text{ for all } (i, t) \right\}.$$

Let $\pi_e(a)$ denote the proportion of patients who, X weeks later, are found to be "significantly improved" given that the observed diagnostic vector is e and the chosen treatment vector is a . For every (e, a) in $E \times A$, the proportion $\pi_e(a)$ is a random variable with a known probability distribution. In our first framework we may have a performance function F , where $F(a, e)$ is a weighted sum of (i) the total cost, denoted $C(a)$, of the treatment vector a , and (ii) the expected value of $\pi_e(a)$.

In the second framework, we have a goal correspondence G , which might, for example, assign to every e the set of actions for which cost does not exceed an upper bound C^* , and at the same time the probability that the proportion of significantly improved patients exceeds a specified lower bound π^* is at least P^* (with $0 < P^* < 1$). Thus

$$G(e) = \left\{ a : C(a) \leq C^*; P^* \leq \text{prob} \{ \pi_e(a) \geq \pi^* \} \right\}.$$

Second example: a multidivisional firm Now person i is the Manager of a production facility which produces product i . The environment component e_i is the current week's cost function for that product: it costs $e_i(q)$ to produce the quantity q . The week's cost function e_i is known with certainty as soon as the week starts. There are n action variables a_1, \dots, a_n , where a_i is the current week's quantity of product i . (We suppose that any positive quantity is feasible). The products are related, and the price at which the quantity a_i of product i is sold depends on the entire vector of quantities $a = (a_1, \dots, a_n)$. The price is denoted $\rho_i(a)$. For the function F , measuring the week's performance (profit), we have

$$F(a, e) = \sum_{i=1}^n [\rho_i(a) \cdot a_i - e_i(a_i)].$$

It may be informationally costly to gather the information about the current e that is needed in order to find an F -maximizing a . If so, it may be appropriate to consider the second framework. The goal correspondence G might specify, for each e , the set of actions for which profit exceeds a lower bound F^* . Thus

$$G(e) = \{a : F(a, e) \geq F^*\}.$$

Third example: an exchange economy This organization's n members are consumers in an L -commodity economy with trade but no production. Person i has a (privately observed) local environment $e_i = (U^i, w^i)$, where U^i is a utility function defined on all vectors with L real nonnegative components and $w^i = (w_1^i, \dots, w_L^i)$ is the vector of Person i 's (nonnegative) initial endowments of the L commodities. An action a specifies trades. It is an nL -tuple $a = \left((a_1^1, \dots, a_L^1), \dots, (a_1^n, \dots, a_L^n) \right)$, where a_ℓ^i may be positive (an addition to i 's endowment of commodity ℓ), negative (a deduction), or zero. Let a^i denote (a_1^i, \dots, a_L^i) . An action a is *feasible for* $e = (e_1, \dots, e_n)$ if $a_\ell^i \geq -w_\ell^i$ for all ℓ and the trades balance, i.e., $\sum_{i=1}^n a^i = 0$. The set of actions that are feasible for e is denoted A_e .

A possible performance function F is as follows:

$$F(e) \text{ is a value of } a \text{ which maximizes } \sum_{i=1}^n U^i(w^i + a^i) \text{ on } A_e.$$

That performance function would be appropriate if the designer of resource-allocating schemes is willing to compare utilities across individuals and to take their sum as an appropriate standard.

On the other hand, it is more conventional to take Pareto optimality and individual rationality as the goal. A trade vector a is *individually rational for* e if

$$(1) \quad U^i(w^i + a^i) \geq U^i(w^i) \text{ for all } i.$$

The trade vector $a \in A_e$ is *Pareto optimal for* e if

$$(2) \quad \left\{ \begin{array}{l} \text{for all } i \text{ and all } \bar{a} \text{ in } A_e, \text{ the following holds:} \\ \text{"}U^i(\bar{a}^i + w^i) > U^i(a^i + w^i)\text{" implies that for some } j \text{ we have "}U^j(\bar{a}^j + w^j) < U^j(a^j + w^j)\text{"}. \end{array} \right.$$

Note that for a given environment $e = ((U^1, w^1), \dots, (U^n, w^n))$, conditions (1) and (2) are restrictions on a alone.

To study the informational requirements of schemes that achieve this goal, the second framework is needed since, for a given e , there may be more than one trade a that is feasible (belongs to A), individually rational, and Pareto optimal. (We would have to restrict the U^i in order to guarantee uniqueness). Formally, the goal correspondence is:

$$G(a, e) = \{a \in A_e : a \text{ satisfies (1) and (2) given } e\}.$$

2.2 How the organization finds its current action when incentives are not an issue

We now consider mechanisms which the organization may repeatedly use in order to find its current action. A mechanism requires the transmission of messages. In our first framework the organization (or its designer) seeks a mechanism which strikes a good balance between the performance measure and the

mechanism’s informational costs. In the second framework it seeks a mechanism which always yields goal-fulfilling actions and is, at the same time, informationally cheap. We start with the assumption that once the mechanism has been chosen, all members of the organization reliably follow its rules. They may be thought of as robots. So we need not worry about designing the mechanism so that each member will *want* to follow the mechanism’s rules. Incentives are introduced in Section 3.

2.2.1 Decentralized many-step broadcast mechanisms which obtain the organization’s action at the final step

Until we reach Section 2.7, where the *designer* decides who shall observe a given environment variable, we will assume that our n -person organization’s current environment e is a vector (e_1, \dots, e_n) , and the local environment e_i is privately observed by Person i . The possible values of e_i comprise the set E_i , and the possible values of e comprise the Cartesian product $E = E_1 \times \dots \times E_n$. There is a set A of possible organizational actions. In our first characterization of a mechanism, we shall suppose that it proceeds in a sequence of steps. At each step each person i broadcasts or announces an *individual message* to everyone. The vector of n individual messages is simply called a *message*. Person i ’s announcement at a given step (his part of the broadcast message) is a function, denoted f_i , of the preceding broadcast message and of e_i . But the variable e_j , for any $j \neq i$, does *not* directly enter the function f_i . The *privacy* of every person is preserved. Others can only learn about the current value of e_i indirectly, through the broadcast message. “Informational decentralization” is an alternative term for privacy preservation.⁶ Suppose that, for a given e , the message \bar{m} has the property that once it is broadcast, the next broadcast message is again \bar{m} . Then \bar{m} is called *an equilibrium message for e* or a *stationary message for e* . When an equilibrium message, say m^* , has been reached, the sequence stops and the organization takes the action $h(m^*)$. The function $h : M \rightarrow A$ is called the *outcome function*.

Formally, let M^i be i ’s *individual message space*, i.e., the set of individual messages that Person i is able to announce. Then the *message space* (the set of possible messages) is $M = M^1 \times \dots \times M^n$. At step t , Person i broadcasts the message $m_i^t = f_i(m^{t-1}, e_i)$, where m^t denotes (m_1^t, \dots, m_n^t) . There is an initial message $m^0(e) = (m_1^0(e_1), \dots, m_n^0(e_n)) \in M$. The message $m^* = (m_1^*, \dots, m_n^*)$ is an equilibrium message for e if, for all i , we have

$$(3) \quad m_i^* = f_i(m^*, e_i).$$

The quadruple $\langle (M^1, \dots, M^n), (m_1^0, \dots, m_n^0), (f_1, \dots, f_n), h \rangle$ is an *n -person privacy-preserving (decentralized) broadcast mechanism on E , with action set A , written in dynamic form*. The term “broadcast” will often be omitted but understood, until we reach section 2.2.9 below, where individually addressed messages are introduced.

In many studies one ignores the message-forming functions f_i . Moreover, one does *not* require that a message have n components, one for each person. Instead it suffices, for the purposes of the study, to define, for each e and each i , the set of messages $\mu_i(e_i)$ (a subset of M) for which the equilibrium

⁶In Section 3.2 the term “decentralized” is given another meaning, related to incentives: in a “decentralized” organization, each person is free to pursue her own self-interest. In particular, she makes a self-interested choice as to whether or not to obey a proposed mechanism’s rules.

condition (3) is satisfied. That set can be specified *without* taking the trouble to specify the individual message spaces and the functions f_i .

Then the *individual-message-correspondence form* of a decentralized (privacy-preserving) mechanism on the environment set E with action set A is a triple $\langle M, (\mu_1, \dots, \mu_n), h \rangle$. Its elements are as follows.

- M is the message space.
- μ_i is a correspondence from E_i to M (i.e., μ_i is a function from E_i to the subsets of M); the μ^i define a correspondence μ , from E to M , namely $\mu(e) = \mu_1(e_1) \cap \dots \cap \mu_n(e_n)$; a message m in $\mu(e)$ is called an equilibrium message for e .
- μ has the *coverage property*⁷, i.e., for all e in E , the set $\mu(e)$ is not empty (for every e there is at least one equilibrium message).
- h (the outcome function) is a function from M to A .

One can give the broadcast mechanism $\langle M, (\mu_1, \dots, \mu_n), h \rangle$ an interpretation that is sometimes called *the verification scenario*. In this scenario, we imagine a central agent who broadcasts a sequence of trial message announcements. When a message m is announced, each person i responds by saying YES if he finds, using his private knowledge of e_i , that m belongs to the set $\mu_i(e_i)$, and he says NO otherwise. The announcements stop when and only when the announcer has announced an m^* for which all n persons say YES. The organization then takes the action $h(m^*)$. The message m^* lies in the set $\mu(e) = \mu_1(e_1) \cap \dots \cap \mu_n(e_n)$.

In a still further condensed formalism, one does not trouble to identify the individual μ_i , but merely specifies the correspondence μ . Then a mechanism is a triple $\langle M, \mu, h \rangle$. The term “decentralized” or “privacy-preserving” is a restriction on μ . It means that there exist correspondences μ_1, \dots, μ_n such that $\mu(e) = \mu_1(e_1) \cap \dots \cap \mu_n(e_n)$, even though we don’t identify them. Thus it is understood, without being made explicit, that $m \in \mu(e)$ means that Person i has determined that m belongs to $\mu_i(e_i)$, using his own private knowledge of the current e_i to do so.

Now suppose the mechanism designer is given a goal correspondence G from E to A , as in our second framework. Then we say that the mechanism $\langle M, \mu, h \rangle$ *realizes* G if

$$\text{for every } e \text{ in } E, \text{ “} m \in \mu(e)\text{” implies “} h(m) \in G(e)\text{”}.$$

An important observation is that *any goal correspondence can be realized by a Direct Revelation (DR) mechanism*. In a DR mechanism each person i reveals his current e_i , i.e., his announced message m_i belongs to E_i . An action in $G(e)$ is taken once a complete description of the current e is assembled. Formally, we have

- $M = E$.

⁷In an alternative terminology, introduced by Hurwicz (1960) in an early and fundamental discussion of mechanisms, the mechanism is called *decisive on E* if it has the coverage property.

- $\mu_i(e_i) = \{e_i\}$
- $h(m) \in G(e)$.

The only equilibrium message for e is e itself. A DR mechanism is, of course, informationally costly. On the other, hand it has the merit that it can reach an equilibrium message in just one step.⁸ Moreover DR mechanisms play a central role in the incentive literature. That literature often confines itself to DR mechanisms, on the ground that any mechanism can be rewritten as a DR mechanism (the “revelation principle”). That may be a correct claim, but it is not useful if one seeks mechanisms that are informationally cheap.

2.2.2 A three-message example wherein messages may be visualized as rectangles that cover E . The organization has two persons. For person i , the set of local environments is the real interval $E_i = [0, 1]$. Consider the three-message mechanism portrayed in the following figure.

FIGURE 1 HERE

The message space is $M = \{m_1, m_2, m_3\}$, where m_j identifies the rectangle labelled m_j . Then

$$\mu_1(e_1) = \begin{cases} \{m_1, m_3\} & \text{for } 0 \leq e_1 < \frac{1}{2} \\ \{m_1, m_2, m_3\} & \text{for } e_1 = \frac{1}{2} \\ \{m_2, m_3\} & \text{for } \frac{1}{2} < e_1 \leq 1. \end{cases}$$

$$\mu_2(e_2) = \begin{cases} \{m_1, m_2\} & \text{for } 0 \leq e_2 < \frac{3}{4} \\ \{m_1, m_2, m_3\} & \text{for } e_2 = \frac{3}{4} \\ \{m_3\} & \text{for } \frac{3}{4} < e_2 \leq 1. \end{cases}$$

Let the outcome function h be the following:

$$h(m_1) = \frac{5}{8}; h(m_2) = \frac{9}{8}; h(m_3) = \frac{11}{8}.$$

It is easy to verify that the mechanism $\langle M, (\mu_1, \mu_2), h \rangle$ so defined realizes the following goal correspondence:

$$G(e) = \left\{ a : |a - (e_1 + e_2)| \leq \frac{5}{8} \right\}.$$

⁸Define the initial message to be $m^o(e) = (m_1^o(e_1), \dots, m_n^o(e_n)) = (e_1, \dots, e_n)$, and let Person i 's message-forming rule have the property that

$$f_i(m, e_i) = m_i \text{ if and only if } m_i = e_i.$$

Then the initial message is already an equilibrium message. The message formed at Step 1 just repeats it. At Step 1, the action is taken, and it belongs to $G(e)$.

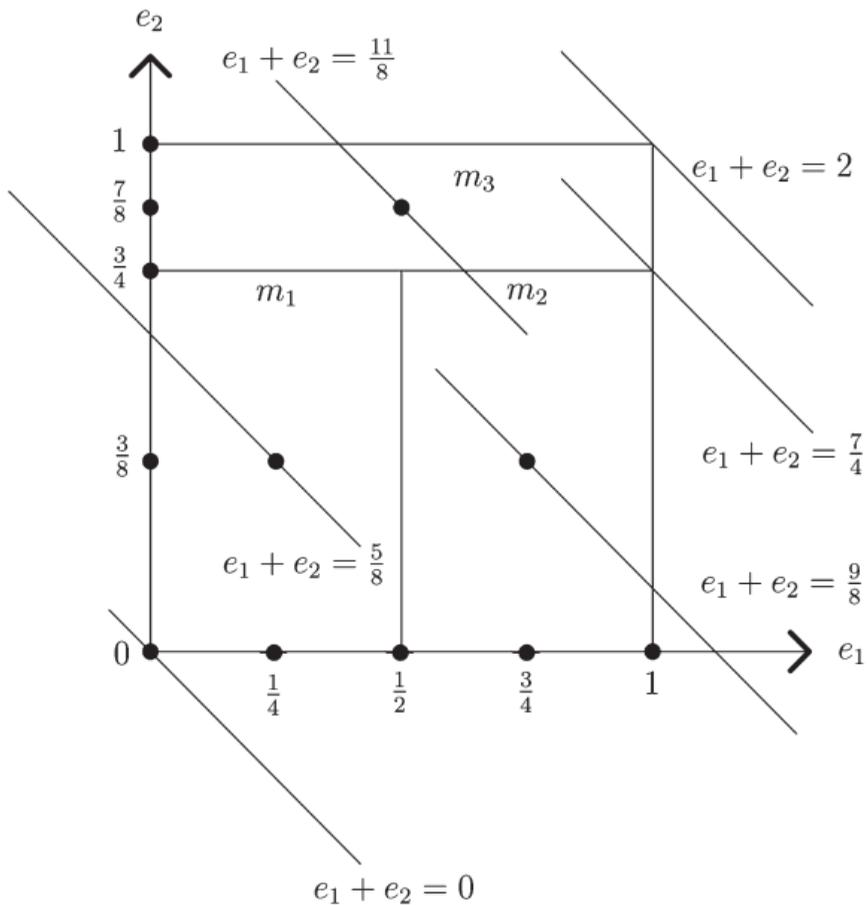


Fig. 1.

That is the case because $h(m_j)$ is the value of $e_1 + e_2$ at the center of the rectangle m_j . Call the center (e_1^j, e_2^j) . The largest value taken by the distance $|(e_1^j + e_2^j) - (e_1 + e_2)|$, over all e in the rectangle m_j , occurs at the “northeast” and “southwest” corners of the rectangle. At those corners the distance equals one quarter of the rectangle’s perimeter. All three rectangles have the perimeter $\frac{5}{2}$.

The above goal correspondence belongs to a class of goal correspondences G_δ , where $\delta > 0$ and $G_\delta(e) = \{a : |a - (e_1 + e_2)| \leq \delta\}$. To interpret this class, go back to our first framework. Suppose that $\hat{a}(e)$ is the organizational action which uniquely maximizes a performance function F , defined on $A \times E$, where the set A is the positive reals and (as in our example) $E = [0, 1] \times [0, 1]$. Suppose that it is very costly for the action taker to learn the exact current value of e (as he would in a DR mechanism). Instead the action taker only learns the rectangle in which e lies. Having learned that the rectangle is m_j , he takes the action $h(m_j)$. As a result there is an “error” $|\hat{a}(e) - h(m_j)|$ at each e in the rectangle m_j . It is straightforward to verify that for *any* correct-action function \hat{a} , no outcome function h achieves a lower maximum error (with respect to \hat{a}), on any given rectangle, than the function which assigns to that rectangle an action that is midway between the minimum of \hat{a} on the rectangle and the maximum.⁹ So a mechanism which minimizes error on each rectangle will use that outcome function, and such a mechanism minimizes error on all of E .

In the Figure-1 example, \hat{a} equals $e_1 + e_2$ and our “midpoint” action is the value of $e_1 + e_2$ at the rectangle’s center. The maximum error of our mechanism (relative to the true $e_1 + e_2$) is $\frac{5}{8}$. It is natural to ask: is our value of δ , namely $\delta = \frac{5}{8}$, the smallest δ such that G_δ can be realized by a three-rectangle mechanism? Is there some other three-message mechanism in which the maximum error (relative to the true value of $e_1 + e_2$) over all e in E is less than $\frac{5}{8}$? The answer turns out to be NO. The argument which establishes that fact has not yet been generalized to the case of k messages. We do not know, in general, the smallest maximum error (relative to the true $e_1 + e_2$) that is achievable by a k -message (k -rectangle) mechanism.¹⁰

2.2.3 The “rectangle” definition of a broadcast mechanism The preceding example suggests that we can define a decentralized (privacy-preserving) mechanism by specifying a covering of $E = E_1 \times \dots \times E_n$, provided that the sets in the covering are *generalized* rectangles, i.e., each is the Cartesian product of its n projections. Let Σ be such a covering of E . Its typical element, denoted σ_m , is a generalized rectangle, i.e., it is a Cartesian product, $\sigma_m = \sigma_m^1 \times \dots \times \sigma_m^n$, where σ_m^i is a subset of E_i . The collection of possible values of the index m is denoted M .

⁹That statement holds as well if the rectangle is a “generalized” one, i.e., it is the Cartesian product of its e_1 -projection and its e_2 -projection and may consist of disjoint pieces.

¹⁰For the general case, one first has to establish that nothing is lost by confining attention to rectangles (such as the three in Figure 1), which are “proper”, rather than being generalized rectangles (each the Cartesian product of its two projections) consisting of disjoint pieces. We next have to argue that we lose nothing by further confining our attention to proper rectangles of *equal* perimeter. (Recall that the maximum error on a rectangle, relative to the true $e_1 + e_2$, equals its quarter-perimeter). That can be shown for our three-rectangle case and certain other cases, but a general argument, for arbitrarily many proper rectangles, is not available. Finally, we have to calculate the smallest obtainable perimeter when we cover E (the unit square) with k proper rectangles of equal perimeter. For our case ($k = 3$) that can indeed be shown to be $\frac{5}{2}$, as in Figure 1. For general k , there is no known formula giving the smallest obtainable common perimeter. There is a conjectured formula, and bounds on the distance between that formula and the unknown true one have been obtained. See Alon and Kleitman, 1986.

To complete our rectangle definition of mechanism, we only need an outcome function h from M to A . Then a broadcast mechanism is defined by the triple $\langle M, \Sigma, h \rangle$, with $\Sigma = \{\sigma_m\}_{m \in M}$. We obtain the $\langle M, (\mu_1, \dots, \mu_n), h \rangle$ specification from the $\langle M, \Sigma, h \rangle$ specification, by letting $\mu_i(e)$ be the set $\{m \in M : e_i \in \sigma_m^i\}$. We obtain the $\langle M, \Sigma, h \rangle$ specification from the $\langle M, (\mu_1, \dots, \mu_n), h \rangle$ specification by letting σ_m^i be the inverse of μ_i , i.e., $\sigma_m^i = \{e_i \in E_i : m \in \mu_i(e_i)\}$.

The verification scenario provides one way to interpret a broadcast mechanism that is specified in the $\langle M, \Sigma, h \rangle$ form. Imagine a central announcer who displays successive rectangles σ_m to all n persons. Each responds with YES if he finds that his privately observed e_i lies in the projection σ_m^i and NO otherwise. When the announcer has found a rectangle σ_m to which all say YES, then he takes the action $h(m)$.

2.2.4 A broadcast mechanism in “agreement function” form In the verification scenario, as we have described it so far, person i responds to an announced message $m \in M$ by inspecting his privately observed e_i and saying YES if he finds that m lies in the set $\mu_i(e)$. In a number of settings it is useful to be more explicit about person i 's procedure, by specifying a *function* which he computes. Let g_i be a function whose domain is $M \times E_i$ and whose range is a finite dimensional Euclidean space. Let person i say YES to the message m if he finds that $g_i(m, e_i) = 0$. We may call g_i person i 's *agreement function*. When $g_i(m, e_i) = 0$, then we may think of person i 's YES announcement as “agreement with” the message m . A mechanism in agreement-function form is a triple $\langle M, (g_1, \dots, g_n), h \rangle$. We obtain the $\langle M, (\mu_1, \dots, \mu_n), h \rangle$ form from the $\langle M, (g_1, \dots, g_n), h \rangle$ form by specifying that $\mu_i(e_i)$ is the set $\{m \in M : g_i(m, e_i) = 0\}$. We obtain the $\langle M, (g_1, \dots, g_n), h \rangle$ form from the $\langle M, (\mu_1, \dots, \mu_n), h \rangle$ form by choosing, for each i , any function g_i which takes the value zero if and only if $m \in \mu_i(e_i)$.

Now suppose that M is the Cartesian product of n individual message spaces M^i , where each M^i is a linear vector space, so that subtracting one value of m_i from another is well defined. Suppose we have written the mechanism in dynamic form, i.e., we have specified a message-forming rule f_i for each person i . Suppose we are interested in the action generated by the mechanism when m is an equilibrium message for e . Then we can rewrite the mechanism in agreement-function form. Let g_i express i 's equilibrium condition for the rule f_i . That is to say, we define

$$g_i(m, e) = f_i(m, e_i) - m_i.$$

When (and only when) $g_i(m, e_i) = 0$, person i 's response to the announced message $m = (m_1, \dots, m_n)$ (in the dynamic version of the mechanism) is to repeat his piece of that announcement. Thus message m is an equilibrium message for e (in the dynamic version) when and only when all n persons i find that $g_i(m, e) = 0$.

2.2.5 A summary We have identified several different ways of specifying a decentralized (privacy-preserving) n -person broadcast mechanism on the environment set $E = E_1 \times \dots \times E_n$, which the n -person organization may use to obtain actions in response to a new environment e in E . The alternative specifications are:

- A mechanism with individual messages in dynamic form. This is a triple $\langle (M^1, \dots, M^n), ((m_1^0, \dots, m_n^0)), (f_1, \dots, f_n), h \rangle$. Here m_i^0 is a function from E_i to M^i and $m_i^0(e_i)$ is i 's initial

individual message when i 's local environment is e_i . In each sequence of announced messages, Person i forms his next individual message by using the message-forming rule f_i .

- A mechanism with individual message correspondences. This is a triple $\langle M, (\mu_1, \dots, \mu_n), h \rangle$, where μ_i is a correspondence from E_i to M .
- A mechanism in which only an equilibrium message correspondence is identified. This is a triple $\langle M, \mu, h \rangle$, where μ is a correspondence from E to M , and it is understood that there are (unspecified) individual message correspondences μ_1, \dots, μ_n such that for all e we have $\mu(e) = \mu_1(e_1) \cap \dots \cap \mu_n(e_n)$.
- A mechanism in rectangle form. This is a triple $\langle M, \{\sigma_m\}_{m \in M}, h \rangle$, where each σ_m is a generalized rectangle in E .
- A mechanism in “agreement function” form. This is a triple $\langle M, (g_1, \dots, g_n), h \rangle$, where g_i is a function from $M \times E_i$ to a (finite-dimensional) Euclidean space.

2.2.6 An example: a price mechanism for an exchange economy. Return now to the n -person L -commodity exchange economy discussed in Section 2.1.3 above. In the classic (Walrasian) mechanism for obtaining individually rational and Pareto optimal allocations, the typical message, broadcast to all n persons, consists of

- a nonnegative price vector $p = (p_2, \dots, p_L)$, with the price of commodity 1 (the numeraire) being one
- a proposed trade vector $a = \left((a_1^1, \dots, a_L^1), \dots, (a_1^n, \dots, a_L^n) \right)$, whose components may be positive, negative, or zero

The proposed trades specified in any message m have the property that $\sum_{i=1}^n a^i = 0$ (where $a^i = (a_1^i, \dots, a_L^i)$), or equivalently, for every i and every commodity ℓ :

$$(\dagger) \quad a_\ell^i = - \sum_{j \neq i} a_\ell^j.$$

The prices and proposed trades in m have the further property that each Person i 's budget balances¹¹, i.e.,

$$(\dagger\dagger) \quad \sum_{\ell=2}^L p_\ell \cdot a_\ell^i = -a_1^i.$$

¹¹That is to say, the amount that i spends on positive additions to his endowment must equal the value of the quantities that he subtracts from his endowment and makes available to others.

Let us write the mechanism in agreement-function form. For person i , the local environment is $e_i = (U^i, w^i)$. Assume U^i to be differentiable and increasing in each of its arguments. The agreement function g_i has $L - 1$ components, denoted g_{i2}, \dots, g_{iL} . Each corresponds to one of the commodities $2, \dots, L$. All of person i 's functions $g_{i\ell}$ equal an arbitrary nonzero number if i 's proposed trade vector is infeasible with regard to some commodity ℓ , i.e., if $a_\ell^i < -w_\ell^i$ (where, as $(\dagger\dagger)$ specifies, $a_1^i = -\sum_{\ell=2}^L p_\ell \cdot a_\ell^i$). For the feasible case, where $a_\ell^i \geq -w_\ell^i$ for all ℓ , consider the bundle that person i holds after the proposed trades have taken place, and consider the ratio of person i 's marginal utility for commodity ℓ (at that post-trade bundle) to his marginal utility for commodity 1. The function $g_{i\ell}$ equals that ratio minus the price of ℓ . For an equilibrium message, each function $g_{i\ell}$ equals zero. That implies that for the prices p , person i 's bundle $(w_1^i + a_1^i, \dots, w_L^i + a_L^i)$ satisfies the first order condition for utility maximization subject to the constraint that the bundle's value not exceed $\sum_{\ell=1}^L p_\ell \cdot w_\ell^i$. Let $U_\ell^i(x_1, \dots, x_L; e_i)$ denote i 's marginal utility for commodity ℓ (the partial derivative of U^i with respect to x_ℓ) when he consumes the bundle (x_1, \dots, x_L) if his utility function is the function U^i specified in e_i .

Using (\dagger) , $(\dagger\dagger)$, we have the following for a message $m = (a, p)$ such that $a_\ell^i \geq -w_\ell^i$ for all ℓ : For $\ell = 2, \dots, L$,

$$g_{i\ell}((a, p), e_i) = \frac{U_\ell^i\left(w_1^i - \sum_{\ell=2}^L p_\ell \cdot a_\ell^i, w_2^i + a_2^i, \dots, w_L^i + a_L^i; e_i\right)}{U_1^i\left(w_1^i - \sum_{\ell=2}^L p_\ell \cdot a_\ell^i, w_2^i + a_2^i, \dots, w_L^i + a_L^i; e_i\right)} - p_\ell.$$

In view of the condition (\dagger) , we can reduce the size of m , by deleting the proposed-trade vector of one person, say person n . In person n 's agreement rule we replace each a_ℓ^n (where $\ell = 2, \dots, L$), with $-\sum_{j=1}^{n-1} a_\ell^j$. We replace the commodity-1 term $-\sum_{\ell=2}^L p_\ell \cdot a_\ell^n$, with the term $\sum_{\ell=2}^L p_\ell \left[\sum_{j=1}^{n-1} a_\ell^j\right]$. Then the message m is a vector of $n(L - 1)$ real message variables, namely $(n - 1)(L - 1)$ trade variables plus $(L - 1)$ prices.

To complete our definition of the mechanism, we have to provide the outcome function h . We let that function be a simple projection operator, i.e., $h(a, p) = a$. If we now assume that each utility function U^i is strictly concave, then the mechanism has the coverage property: for every $e \in E$, there exists an equilibrium message (a, p) . Moreover the allocation a is feasible, Pareto optimal, and individually rational.

Now consider any other mechanism whose equilibrium actions (trades) are also individually rational and Pareto optimal. Under what further restrictions on the rival mechanism can we claim that its message space cannot be "smaller" than that of the mechanism we have just constructed? In particular, if the rival mechanism's messages are again real vectors, when can we claim that those vectors cannot have fewer than $(n)(L - 1)$ components? That is a well-studied question. We shall return to it in Section 2.2.8 below.

2.2.7 Another example: a price mechanism for a firm with managers and a resource allocator¹² In this organization, persons $1, \dots, n - 1$ are Managers and person n is an Allocator. Manager j is in charge of n^j activities. An activity uses resources and it generates profit. There are

¹²This example is discussed in Ishikida and Marschak (1996).

L resources and the quantity e_ℓ^n of resource ℓ is available to the firm. Manager j operates each of his n^j activities, say activity k , at some level x_k^j (a nonnegative real number). We assume $n^j > L, j = 1, \dots, n-1$. When Manager j 's activity-level vector is $x^j = (x_1^j, \dots, x_{n^j}^j)$, he contributes $e_0^j(x^j)$ to the organization's profit and he uses the quantity $e_\ell^j(x^j)$ of resource ℓ , for $\ell = 1, \dots, L$, provided that the Allocator gives him those resource quantities. In each period there are new resource availabilities e_ℓ^n , and these become known to the Allocator. For each Manager j , there are also new profit functions e_0^j and new resource-requirement functions e_ℓ^j , and those become known to Manager j .

So the Allocator's local environment is an L -tuple of resource availabilities, namely $e_n = (e_1^n, \dots, e_L^n)$, and Manager j 's local environment is an $(L+1)$ -tuple of functions, namely $e_j = (e_0^j, e_1^j, \dots, e_L^j)$. The organization's action is an allocation vector $y = ((y^1, \dots, y_L^1), \dots, (y^{n-1}, \dots, y_L^{n-1}))$, where y_ℓ^j is the quantity of resource ℓ allocated to manager j . The action y meets the organization's goal if it permits the managers to choose activity-level vectors that maximize the firm's total profit subject to the current resource availabilities. So the goal correspondence is G , defined by:

$$G(e) = \left\{ y : \text{some nonnegative } (x^1, \dots, x^{n-1}) \text{ satisfies } e_\ell^j(x^j) = y_\ell^j, j = 1, \dots, n-1, \ell = 1, \dots, L, \right. \\ \left. \text{and maximizes } \sum_{j=1}^{n-1} e_0^j(x^j) \text{ subject to } \sum_{j=1}^{n-1} e_\ell^j(x^j) \leq e_\ell^n, \ell = 1, \dots, L \right\}.$$

Now suppose that the sets of possible local environments are as follows. For each Manager j

$$E_j = \{e^j : e_0^j \text{ is strictly concave and differentiable ; } e_\ell^j \text{ is convex and differentiable , } \ell = 1, \dots, L\},$$

while for the Allocator we have $E_n = \mathbb{R}^{L+}$. Assume that for every e in $E = E_1 \times \dots \times E_n$, the set $G(e)$ is nonempty. Then we can construct a mechanism $\langle M, \mu, h \rangle$ which uses prices and realizes G on E . To do so, consider a vector p of nonnegative resource prices, (p_1, \dots, p_L) , and, for each Manager j consider the following local problem:

$$\text{find } x^j \text{ so as to maximize } e_0^j(x^j) - \sum_{\ell=1}^L p_\ell e_\ell^j(x^j) \text{ subject to } x^j \geq 0.$$

Let $S_j(e_j, p)$ denote the set of solutions to that problem.

We use the individual-message-correspondence form $\langle M, (\mu_1, \dots, \mu_n), h \rangle$ to define our mechanism. The message space M is \mathbb{R}^{nL+} . The typical message, broadcast to everyone, is a pair $m = (p, y)$, where p is a price vector and y is a proposed allocation vector. For Manager $j = 1, \dots, n-1$, define

$$\mu_j(e_j) = \{(p, y) \in M : \text{for some } x^j \text{ in } S_j(e_j, p) \text{ and for all } \ell = 1, \dots, L, \\ \text{we have (1) } e_\ell^j(x^j) \leq y_\ell^j \text{ and (2) } p_\ell \cdot (y_\ell^j - e_\ell^j) = 0\}.$$

For the Allocator, define

$$\mu_n(e_n) = \{(p, y) \in M : \sum_{j=1}^{n-1} y_\ell^j \leq e_\ell^n \text{ and } p_\ell \cdot (e_\ell^n - \sum_{j=1}^{n-1} y_\ell^j) = 0, \ell = 1, \dots, L\}.$$

Finally, the outcome function is a projection operator (just as it was in the exchange-economy price mechanism): $h(p, y) = y$. It is quickly verified that if (p, y) is an equilibrium message for e (i.e., $(p, y) \in \mu_1(e_1) \cap \dots \cap \mu_n(e_n)$), then y , together with some activity-level vector (x^1, \dots, x^{n-1}) , satisfies the first requirement in our definition of G . Moreover that activity-level vector satisfies the Kuhn-Tucker conditions associated with the maximization described in the second requirement of our definition of G . The equilibrium vector p is the vector of Lagrange multipliers in those Kuhn-Tucker conditions. So we have $y \in G(e)$. Under our assumptions on E , a Kuhn-Tucker solution exists for every e in E . That means that for every e in E , there exists an equilibrium (p, y) . So our mechanism has the coverage property.

We conclude that our price mechanism indeed realizes our goal correspondence G on the environment set E . Its message space has dimension nL . It is natural to ask: is there another mechanism (with appropriate regularity properties) which also realizes G on E and does so with a message space of dimension less than nL ? The next section deals with questions of that sort.

2.2.8 Using the “uniqueness property”, or “fooling sets”, to obtain a useful lower bound to the message-space size of a “smooth” goal-realizing broadcast mechanism The best known cost measure for a broadcast mechanism is the size of its message space M . Suppose each broadcast message m is a vector of $Q \geq 1$ real numbers m_1, \dots, m_Q , and M is the cube $\{m : A_q \leq m_q \leq B_q, q = 1, \dots, Q\}$, where $A_q < B_q$ for all q . Then the natural size measure for M is its dimension, namely Q . More generally, the dimension of M is a suitable cost measure as long as M is any subset of \mathbb{R}^Q for which “dimension” is defined. (For example, M may be a differentiable manifold). Even more general classes of message space have been studied, and message space size measures have been defined for those classes as well.

Whatever the definition, we typically seek to identify the broadcast mechanisms that realize a given goal correspondence G while using the smallest possible message space. When we do so, however, *smuggling* of many numbers into a single number is a basic difficulty. If, for example, we start with a mechanism in which each message is a vector of $Q > 1$ real numbers, then we may define a new mechanism, which realizes the same goal correspondence as the original mechanism, but has messages comprised of just a single real number. One way to do this is to let the single real number be a decimal which encodes the q th of the Q original numbers as a string composed of our decimal’s q th digit, its $(Q + q)$ th digit, its $(2Q + q)$ th digit and so on. That particular smuggling trick is ruled out if, when writing the mechanism in the message-correspondence form, we require that every set $\mu_i(e_i)$ contain an element $t_i(e_i)$, where t_i is a continuous function. A more elaborate smuggling trick uses the Peano “space-filling-curve” mapping. (See, for example, Apostol, 1957, pp. 396-8). That mapping allows us to recover the Q numbers from a single number in a continuous manner. A requirement stronger than continuity (e.g., differentiability) is therefore needed in order to exclude it. If we want to give a nontrivial meaning to the message-space minimality of a particular mechanism within some interesting class of mechanisms, then *smoothness requirements on the candidate mechanisms are unavoidable*.

Now consider the case of a two-person organization. We are given a goal correspondence G from $E = E_1 \times E_2$ to an action set A . Suppose we suspect that no suitably smooth mechanism can realize G with a message space of dimension less than D . Suppose further that E is a subset of a Euclidean space, and that we have found \bar{E} , a subset of E — called a *test class of environments* — which has

dimension D . Suppose further that G has the *uniqueness property* on \bar{E} . That means that there is no action which is goal-fulfilling for all four corners of the “cube” defined by a pair of distinct environments e^*, e^{**} in \bar{E} , i.e., the four environments $(e_1^*, e_2^*), (e_1^{**}, e_2^{**}), (e_1^*, e_2^{**}), (e_1^{**}, e_2^*)$. Thus for any e^*, e^{**} in \bar{E} , the following holds:

$$\text{If } G(e_1^*, e_2^*) \cap G(e_1^{**}, e_2^{**}) \cap G(e_1^*, e_2^{**}) \cap G(e_1^{**}, e_2^*) \neq \emptyset, \text{ then } e^* = e^{**}.$$

Now note that if a mechanism realizes G on all of E , then, in particular, it realizes G on the test class \bar{E} . Next, recall that a mechanism is privacy preserving: in determining whether or not to agree to a broadcast message m (or determining whether or not m lies in $\mu_i(e_i)$), each person i looks only at his own local environment e_i . Thus if a message \bar{m} is an equilibrium message for both (e_1^*, e_2^*) and (e_1^{**}, e_2^{**}) , then \bar{m} is also an equilibrium message for (e_1^*, e_2^{**}) and (e_1^{**}, e_2^*) . So if the mechanism’s action $h(\bar{m})$ is indeed going to lie in the sets $G(e^*)$ and $G(e^{**})$, as realization of G requires, then $h(\bar{m})$ lies in $G(e_1^*, e_2^{**})$ and $G(e_1^{**}, e_2^*)$ as well. Since G has the uniqueness property on \bar{E} , it follows that $e^* = e^{**}$. So a pair of distinct environments in the D -dimensional test class \bar{E} cannot share the same equilibrium message. Since M must contain at least one equilibrium message for all e in \bar{E} (the coverage requirement), we can claim, informally, that M must be at least as large as \bar{E} .

More precisely, consider any mechanism which realizes G on E and hence on \bar{E} . Consider the restriction of μ to \bar{E} and call that restriction $\tilde{\mu}$. So $\tilde{\mu}$ is a correspondence from \bar{E} to a subset of M , namely the set $\tilde{\mu}(\bar{E}) = \mu(\bar{E})$. Now since \bar{E} has the uniqueness property for G , the *inverse* of the correspondence $\tilde{\mu}$ is a *function*, which we may call $t : \tilde{\mu}(\bar{E}) \rightarrow \bar{E}$. For example t might be the Peano mapping, which assigns a member of \bar{E} to every message in $\tilde{\mu}(\bar{E})$, even though \bar{E} has higher dimension than $\tilde{\mu}(\bar{E})$. A message m in $\tilde{\mu}(\bar{E})$ is an equilibrium message for the environment $t(m) \in \bar{E}$ and for no other environment in \bar{E} . Since every environment in \bar{E} must have some equilibrium message in $\mu(\bar{E})$ (by the coverage property), our function t is *onto* (it is a surjection).

Now suppose our smoothness requirement on the candidate mechanisms is that the function t , the inverse of $\tilde{\mu}$, be differentiable. If M , and hence $\mu(\bar{E}) = \tilde{\mu}(\bar{E})$, had a smaller dimension than \bar{E} , then t would be a one-to-one function from one set onto a second set having higher dimension. That cannot be the case if t is differentiable. (For example, the Peano mapping, while continuous, is not differentiable, so a mechanism in which t is the Peano mapping from $\tilde{\mu}(\bar{E})$ onto the higher-dimensional set \bar{E} violates our smoothness requirement). So we have confirmed our suspicion that D is indeed a lower bound for smooth mechanisms which realize G on all of E . We obtain the same conclusion for other smoothness requirements. Some of them are weaker than our requirement that the function t (the inverse of the message correspondence, restricted to the test class) be differentiable.¹³ Moreover, there is another way to force M to have a dimension at least as large as \bar{E} . We can impose requirements directly on the correspondence $\tilde{\mu}$ rather than on its inverse. Suppose we require $\tilde{\mu}$ to be *locally threaded*. That means that for any neighborhood N in \bar{E} , we can find a continuous function $v : N \rightarrow M$ which is a selection from $\tilde{\mu}$, i.e., $v(e) \in \tilde{\mu}(e)$ for all e in N . The uniqueness property of \bar{E} tells us that for any

¹³In particular, we may use the weaker requirement that t be “Lipschitz-continuous”, i.e., there exists $K > 0$ such that for all m', m'' in $\tilde{\mu}(E)$, we have $\|t(m') - t(m'')\| \leq K \cdot \|m' - m''\|$. (Here the symbol $\|x\|$, for $x = (x_1, \dots, x_\ell)$, denotes $\max\{|x_j| : j \in \{1, \dots, \ell\}\}$). The Peano mapping is not Lipschitz-continuous.

two distinct environments \bar{e}, \bar{e} in N , we have $v(\bar{e}) \neq v(\bar{e})$. It can be shown that this fact rules out a continuous v if M indeed has a smaller dimension than \bar{E} .¹⁴

The technique extends to n -person mechanisms. Let the symbol e^{**}/e_i^* denote the vector obtained from $e^{**} = (e_1^{**}, \dots, e_n^{**})$ when we replace e_i^{**} with e_i^* . Then the statement “the correspondence G from $E = E_1 \times \dots \times E_n$ to an action set A has the uniqueness property on the test class $\bar{E} \subseteq E$ ” means that

$$\text{if } e^*, e^{**} \in \bar{E} \text{ and } G(e^*) \cap \left(\bigcap_{i=1}^n G(e^{**}/e_i^*) \right) \neq \emptyset, \text{ then } e^* = e^{**}.$$

It is interesting to note that the idea just sketched was developed independently by computer scientists and economists.¹⁵ In the computer science field known as “communication complexity”¹⁶, one studies dialogues between n persons that end with one person having enough information to compute a function F of n numbers, each of them privately known by one of the n persons. The dialogue is a sequence of binary strings. The dialogue changes when the privately known numbers change. One wants the worst-case dialogue to be as short as possible. If the function possesses a “fooling set”, then the size of the fooling set provides a lower bound to the length of the worst-case dialogue. In the terminology we have just developed, a fooling set is a set on which F has the uniqueness property.

An illustration of the uniqueness technique: resource allocating mechanisms for a class of exchange economies

Let us return to the n -person L -commodity exchange economies discussed

¹⁴We may want to consider mechanisms whose message space M is not Euclidean but consists, for example, of certain infinite sequences, or of integer k -tuples, or of preference orderings (as when an environment specifies agents’ individual preference orderings and a message identifies a set of possible environments). That has motivated the study of mechanisms with message spaces that are general topological spaces. Then, instead of comparing dimensions, we use a general topological definition of the statement that one message space is “at least as large as” another. If M is Euclidean, then “at least as large as” reduces to “having a dimension no smaller than”. For example, one may define M^* to be at least as large as M^{**} if and only if there is a subspace of M^* which is homeomorphic to M^{**} . For each such topological definition, we seek an associated smoothness requirement on the message correspondence μ used by a G -realizing mechanism, so that a smooth mechanism’s message space is at least as large as a test class \bar{E} having the uniqueness property for G . One such requirement is “spot-threadedness” of μ on the test class. That is a weaker requirement than local threadedness. It means that there is an open set $W \subseteq \bar{E}$ and a function $q : W \rightarrow M$ such that $q(e) \in \mu(e)$ for all e in W . If a G -realizing mechanism obeys that condition, while its message space and the test class \bar{E} are both Hausdorff and \bar{E} is locally compact, then the message space must be at least as large as the test class, where “at least as large as” has the meaning just given. The details are carefully developed in Section II.3 of Hurwicz, 1986, which concerns “a strategy for obtaining minimality results”. See also (among others) Mount and Reiter, 1974; Walker, 1977; Hurwicz and Marschak, 1985.

¹⁵Economic applications of the technique, so as to establish lower bounds to the message-space size required to achieve various resource-allocation goals, include the following papers, in each of which the lower bound is shown to be attainable by a particular mechanism that uses prices: Mount and Reiter (1977); Osana (1978); Sato (1981); Jordan (1982), which shows that the *only* mechanisms that have minimal message space while realizing Pareto-optimality are versions of the competitive mechanism; Chander (1983); Aizpura and Manresa (1995); Calsamiglia and Kirman (1998); Tian (2004); Stoenescu (2004); Osana (2005). On the other hand, the following papers find that realization of the resource-allocation goal requires a message space of infinite dimension: Calsamiglia (1977) (which permits increasing returns in production); Hurwicz and Weinberger (1990), Manresa (1993), and Kaganovitch (2000) (which consider efficient intertemporal resource allocation); and Jordan and Xu, (1999), on expected profit maximization by the managers in a firm.

¹⁶See, for example, Lovász (1990), Karchmer (1989), Kushilevitz and Nisan (1997).

in 2.1.3 and 2.2.6 above. Recall that person i 's local environment is a pair $e_i = (U^i, w^i)$, where U^i is a utility function and $w^i = (w_1^i, \dots, w_L^i)$ is an endowment vector. As before, there is a set E_i of possible local environments e_i . Recall also that the action the economy chooses is a trade nL -tuple $a = (a^1, \dots, a^n)$, where $a^i = (a_1^i, \dots, a_L^i)$. Call the action a interior for e if $a_\ell^i > -w_\ell^i$ for all (i, ℓ) . Slightly modify the goal correspondence defined in 2.1.3 so that it is now defined by:

$$G(e) = \{a : \sum_{i=1}^n a^i = 0 ; \text{ for the economy defined by } e, a \text{ is interior, Pareto-optimal} \\ \text{and individually rational}\}.$$

The price mechanism introduced in 2.2.6 has a message space of dimension $n(L - 1)$ and realizes G on certain environment sets $E = E_1 \times \dots \times E_n$. Can we show $n(L - 1)$ to be a lower bound for any suitably smooth mechanism which realizes G on E ? Yes we can, using the uniqueness technique, provided E contains a suitable test class \bar{E} whose dimension is $n(L - 1)$. In the test class most used in the literature¹⁷, utility functions have the Cobb-Douglas form and endowments are fixed. Specifically, if e is in \bar{E} , then person i 's endowment vector w^i is $(1, 1, \dots, 1)$, and his utility for the bundle (X_1, \dots, X_L) is $\prod_{\ell=1}^L X_\ell^{\alpha_i}$, where the α_i are positive numbers such that $\alpha_1 + \dots + \alpha_L = 1$. Thus each e_i is uniquely determined by $L - 1$ real numbers and so the set \bar{E}_i has dimension $L - 1$. The dimension of \bar{E} is $n(L - 1)$, which equals the dimension of our price mechanism's message space.

The uniqueness property of G on the set \bar{E} is readily shown, using the first order conditions that characterize the unique n interior trade vectors a^i that maximize each person i 's Cobb-Douglas utility subject to the balancing constraint $\sum a^i = 0$. One then has a choice of several smoothness conditions to be imposed on the candidate G -realizing mechanisms. If the set of messages that are equilibrium messages for the environments in \bar{E} has dimension less than $n(L - 1)$, then each of these smoothness conditions rules out a one-to-one mapping from that set of equilibrium messages onto the $n(L - 1)$ -dimensional environment set \bar{E} . But if the mechanism indeed realizes G on \bar{E} , then such a mapping must exist, by the uniqueness property of G on \bar{E} . So if the entire set E contains our test class \bar{E} , we can rule out¹⁸ a smooth broadcast mechanism that realizes G on all of E and has a message space of dimension less than $n(L - 1)$.¹⁹

¹⁷See, among others, Mount and Reiter (1977) and Jordan (1982).

¹⁸For the organization considered in 2.2.7 ($N - 1$ managers and an Allocator), a similar argument (given in Ishikida and Marschak, 1996) establishes that no suitably smooth mechanism can realize the goal function G (defined in 2.2.7) with a message space dimension less than nL (the message-space dimension of the G -realizing price mechanism that we constructed). In our nL -dimensional test class, the k th activity for manager j (who has $n^j < L$ activities) uses only the resource k and earns a profit of $2\alpha_k^j \sqrt{x_k^j}$ when it is operated at level x_k^j . Each member of the test class is defined by an nL -tuple of positive numbers (the numbers α_k^j), so the test class has dimension nL . It is straightforward to show that G has the uniqueness property on that test class.

¹⁹Another interesting setting for dimensionally minimal broadcast mechanisms is the allocation of one or more objects among n persons, each of whom has a private valuation for each object. Consider the case of a single object. Let e_i be i 's valuation for the object and let E_i be the interval $[0, H]$, where $H > 0$. Let the action set be $\{1, \dots, n\}$, where "action i " is the allocation of the object to person i . Let the goal be allocation to a maximal-value person, i.e., the goal correspondence G is defined by $G(e) = \{j : e_j \geq e_i \text{ for all } i\}$. Then one G -realizing broadcast mechanism uses messages $m = (t, J)$, where t is a real number and J is an integer in $\{1, \dots, n\}$. Person i agrees to m if and only if: (1) $i \neq J$ and $e_i \leq t$ or (2) $i = J$ and $e_i = t$. The outcome function h is a projection: $h(t, J) = J$. This mechanism is one way to model a Dutch (descending) auction. Consider a subclass of $E = E_1 \times \dots \times E_n$, namely the "diagonal" class

2.2.9 Mechanisms in which messages are not broadcast but are individually addressed, and responsibility for each action variable is assigned to a selected person

We have assumed thus far that all parts of a given message are broadcast. They are “seen” or “heard” by all of the organization’s n members. It may, of course, be true that while an agreement function g_i has the entire broadcast message m as an argument, the function g_i is sensitive to only a portion of m , namely the portion that i hears. Even though m is broadcast to everyone, we can interpret “hearing only a portion of the broadcast message m ” as ignoring all of m except that portion. Formally, for every person i , we may be able to write every message m as a pair (m_i, m_{-i}) , where m_i is the portion of m that i hears, and to write the agreement function $g_i(m, e_i)$ as $g_i^*(m_i, e_i)$. Similarly we may be able to write the set $\mu_i(e_i)$ as $\{m = (m_i, m_{-i}) \in M : m_i \in \mu_i^*(e_i)\}$, where μ_i^* is a correspondence from E_i to M_i^* and M_i^* is the set of possible values of m_i .

But if we want to permit messages to be individually addressed, and if we want to study the cost born by person i as he hears and processes the messages he receives, and responds to them by sending further messages to certain other persons, then it is more convenient to extend our previous formalism by introducing *network mechanisms*²⁰. That will also have another advantage: it will allow us to be explicit about who is responsible for a given action variable. Our mechanism concept thus far has been silent on this matter.

In defining a network mechanism we may again use the agreement-function form, but an agreement function’s domain and range are now different. We start by letting M denote an n -by- n matrix of sets M_{ij} , where M_{ij} is the set of possible messages that i may send to j . The set M_{ii} on the diagonal of M may be empty, or, if it is not, we may interpret the self-addressed messages in M_{ii} as stored information. Moreover, M_{ik} may be empty for some $k \neq i$. That means that i never sends a message to k . Next let M^i denote the Cartesian product of the sets in the i th row of M , i.e., $M^i = M_{i1} \times \cdots \times M_{in}$. Let M_i denote the Cartesian product of the sets in the i th column of M , i.e., $M_i = M_{1i} \times \cdots \times M_{ni}$. Let $P_i(M)$ denote the Cartesian product of the sets that are in the i th row or the j th column of M . Thus $P_i(M) = M_i \times \{M^i \setminus M_{ii}\}$. We shall say that a message $m_{uv} \in M_{uv}$ is *heard* by i if it is received by i (so that $u \neq i, v = i$), sent by i (so that $u = i, v \neq i$) or stored by i (so that $u = v = i$). Then $P_i(M)$ is the set of the possible message vectors that i can hear.

We shall speak of a *message array* $m \in M$. Its typical component is an individually addressed message $m_{ij} \in M_{ij}$, where m_{ij} is a vector of s_{ij} real numbers; s_{ij} may be zero. The symbol $P_i(m)$ will denote the portion of m that i hears; $P_i(m)$ is an element of the set $P_i(M)$. Let the domain of person i ’s agreement function g_i be the Cartesian product of E_i with the set $P_i(M)$ of possible message vectors that i can hear, and let its range be $\mathbb{R}^{\sum_{j=1}^n s_{ij}}$. The statement “ $g_i(P_i(m), e_i) = 0$ ”, means that person

$\bar{E} = \{e \in E : e_1 = e_2 = \cdots = e_n\}$. It is easily seen that G has the uniqueness property on \bar{E} . But \bar{E} has dimension one, so it does not provide a useful lower bound for mechanisms whose messages are real vectors. The “auction” mechanism, however, uses both real numbers and integers. One has to be careful in choosing a cost measure, and smoothness requirements, for mechanisms of that sort. Much more challenging is the case of several objects, when each person has a valuation for each subset of the set of objects, and each person may be allocated a subset. One seeks a mechanism which finds (at equilibrium) an allocation that maximizes the sum of the valuations. Lower bounds for such mechanisms have been developed by Nisan and Segal, 2005. The mechanisms considered again use both real numbers and integers in their messages. The uniqueness technique, using a counterpart of the “diagonal” test class, plays a central role in that study. The ideas are extended to a much larger class of allocation problems in Segal, 2004.

²⁰They are studied in Marschak and Reichelstein, 1995 and 1998.

i finds the message array m to be acceptable: given his current local environment e_i , and given that he has received the messages m_{1i}, \dots, m_{ni} , he finds it appropriate to send the messages m_{i1}, \dots, m_{in} . The message array m is an *equilibrium message array for the environment* $e = (e_1, \dots, e_n)$ if all persons find it acceptable, i.e., $g_i(P_i(m), e_i) = 0$ for all persons i .²¹

To complete the definition of a network mechanism, we have to specify how the action variables are chosen once an equilibrium message array is found. Let the organization have k action variables, z^1, \dots, z^k ; let Z^j be the set of possible values of z^j , $j = 1, \dots, k$; and let $Z = Z^1 \times \dots \times Z^k$ be the set of possible organizational action k -tuples $z = (z^1, \dots, z^k)$. Partition the index set $\{1, \dots, k\}$ into n sets (some of them may be empty), namely J_1, \dots, J_n . The (possibly empty) set J_i identifies the action variables for which i is responsible. Those are the variables z^r , where $r \in J_i$; they comprise a vector z^{J_i} belonging to the set $Z^{J_i} = \times_{t \in J_i} Z^t$. Person i chooses the value of the action variables that are in his charge as a function of what he has heard. So he uses an outcome function $h_i : P_i(M) \rightarrow Z^{J_i}$.

As before, we want the agreement functions to have the coverage property: for every $e \in E$, there exists a message array m which is an equilibrium array for e . If coverage is satisfied, then a triple $\langle M, (g_1, \dots, g_n), (h_1, \dots, h_n) \rangle$, whose elements we have just defined, is a (*privacy-preserving*) n -person network mechanism on the environment set $E = E_1 \times \dots \times E_n$ with action space $Z = Z^1 \times \dots \times Z^k$. As before, we may be given a goal correspondence $G : E \rightarrow Z$, where the set $G(e)$ consists of the organizational actions z which meet a certain goal when the environment is e . As before, we shall say that a given network mechanism *realizes* G if, for every e in E , the organizational action $(h_1(P_i(m)), \dots, h_n(P_i(m)))$ lies in the set $G(e)$ whenever m is an equilibrium message array for e . Note that every network mechanism has a *communication graph*. Its nodes are the n persons, and there is an edge between i and j if and only if at least one of the sets M_{ij}, M_{ji} is nonempty.

An example: a network “price” mechanism for a three-plant four-person firm Consider a four-person firm. Person 4 markets two products. He obtains revenue from the quantities Q_1 and Q_2 of the products, which are produced, respectively, by person 1 and person 2. Person 4’s privately observed local environment is a *function*, namely the revenue function $e_4(Q_1, Q_2)$. Person 1’s local environment is the cost function $e_1(Q_1)$. For person 2, cost depends not only on product quantity but also on the quantity I of an intermediate material used in production (the material is supplied by person 3). So person 2’s local environment is the cost function $e_2(Q_2, I)$. Person 3 produces the intermediate material; his local environment is the cost function $e_3(I)$. For persons 1 and 3 the local-environment set E_i (where $i = 1$ or 3) is the set of all continuous convex functions from a closed interval $[A_i, B_i]$ (with $0 \leq A_i < B_i$) to the positive real numbers. For person 2, the local-environment set E_2 is the set of all continuous convex functions of two variables, from a set $[A_2, B_2] \times [C, D]$ (with $0 \leq A_2 < B_2, 0 \leq C < D$) to the positive real numbers. For person 4, E_4 is the set of all continuous concave functions from a closed interval $[A_4, B_4]$ (with $0 \leq A_4 < B_4$) to the positive real numbers. The numbers A_i, B_i, C, D stay the same for all the environments $e = (e_1, \dots, e_n)$.

Now consider a network mechanism in which person 4 sends prices u_1 and u_2 to persons 1 and 2, respectively, and 1 and 2 reply with quantities Q_1, Q_2 that they are willing to supply to 4 at those prices. Similarly, person 2 sends an intermediate-material price v to 3, who replies with an

²¹We can also write a network mechanism in dynamic form, in message-correspondence form, or in rectangle form, just as we can for broadcast mechanisms.

intermediate-material quantity I that he is willing to supply to 2 at that price. So there are six message variables: $m_{14} = Q_1, m_{23} = v, m_{24} = Q_2, m_{32} = I, m_{41} = u_1, m_{42} = u_2$. Let each of the sets M_{ij} of possible message-variable values be the nonnegative real numbers. For all other i, j combinations let the set M_{ij} be empty. Ignore the empty sets and note that for the typical message array, say $m = (\bar{Q}_1, \bar{Q}_2, \bar{I}, \bar{u}_1, \bar{u}_2, \bar{v})$, we have $P_4(m) = (u_1, u_2, Q_1, Q_2), P_1(m) = (u_1, Q_1), P_2(m) = (u_2, Q_2, v, I)$, and $P_3(m) = (I, v)$.

Figure 2 portrays the message flows in this six-message-variable mechanism.

FIGURE 2 HERE

The agreement rules of our price mechanism will express the usual conditions for divisional profit maximization. Consider the typical message array $\bar{m} = (\bar{Q}_1, \bar{Q}_2, \bar{I}, \bar{u}_1, \bar{u}_2, \bar{v})$. Person 4 agrees with \bar{m} (he finds that $g_4(P_1(\bar{m}), e_4) = 0$) if and only if his divisional profit $e_4(Q_1, Q_2) - u_1Q_1 - u_2Q_2$ is maximized by (\bar{Q}_1, \bar{Q}_2) . He can determine whether or not that is so *without* knowing two components of the array \bar{m} , namely \bar{I} and \bar{v} . Person 1 agrees if and only if the profit $u_1Q_1 - e_2(Q_1)$ is maximized by \bar{Q}_1 , and does not need to know $\bar{u}_2, \bar{Q}_2, \bar{I}$ or \bar{v} . Person 2 agrees if and only if $u_2e_2(Q_2, I) - vI$ is maximized by (\bar{Q}_2, \bar{I}) , and does not need to know \bar{u}_1 or \bar{Q}_1 . Finally person 3 agrees if and only if $vI - e_3(I)$ is maximized by \bar{I} , and does not need to know $\bar{u}_1, \bar{u}_2, \bar{Q}_1$ or \bar{Q}_2 .

The organization's action variables are Q_1, Q_2 and I . (To minimize notation we use the same symbol for the action variable as for the message variable associated with it). For each action variable, we let the set of possible values (one of the sets Z^k in our general definition) be the nonnegative reals. We have many choices in designing the outcome function. We may, for example, give person 4 responsibility for the action variables Q_1, Q_2 , while Person 3 has responsibility for I , the remaining action variable. Then we write:

$$h_4(P_4(m)) = h_4(Q_1, Q_2, u_1, u_2) = (Q_1, Q_2); h_3(P_3(m)) = h_3(I, v) = I.$$

(The outcome function is simply a projection operator, just as it was in the exchange-economy price mechanism that we considered before introducing network mechanisms).

Under our assumptions on the E_i , our agreement functions have the coverage property. Moreover if $\bar{Q}_1, \bar{Q}_2, \bar{I}$ are the actions of an equilibrium message array, then they maximize the firm's profit. That is to say, our network mechanism realizes the following goal correspondence:

$$G(e) = \{(\bar{Q}_1, \bar{Q}_2, \bar{I}) : (\bar{Q}_1, \bar{Q}_2, \bar{I}) \text{ maximizes } e_4(Q_1, Q_2) - e_1(Q_1) - e_2(Q_2, I) - e_3(I)\}.$$

Note that we may, if we wish, reverse the directions of the flows depicted in Figure 2. We may, for example, let person 1 send a price to person 4, who replies with a quantity. Let the agreement rules stay as they were. Then the set of equilibrium messages for any e does not change and hence the set of actions obtained at equilibrium for any given e does not change. The revised mechanism again realizes G .

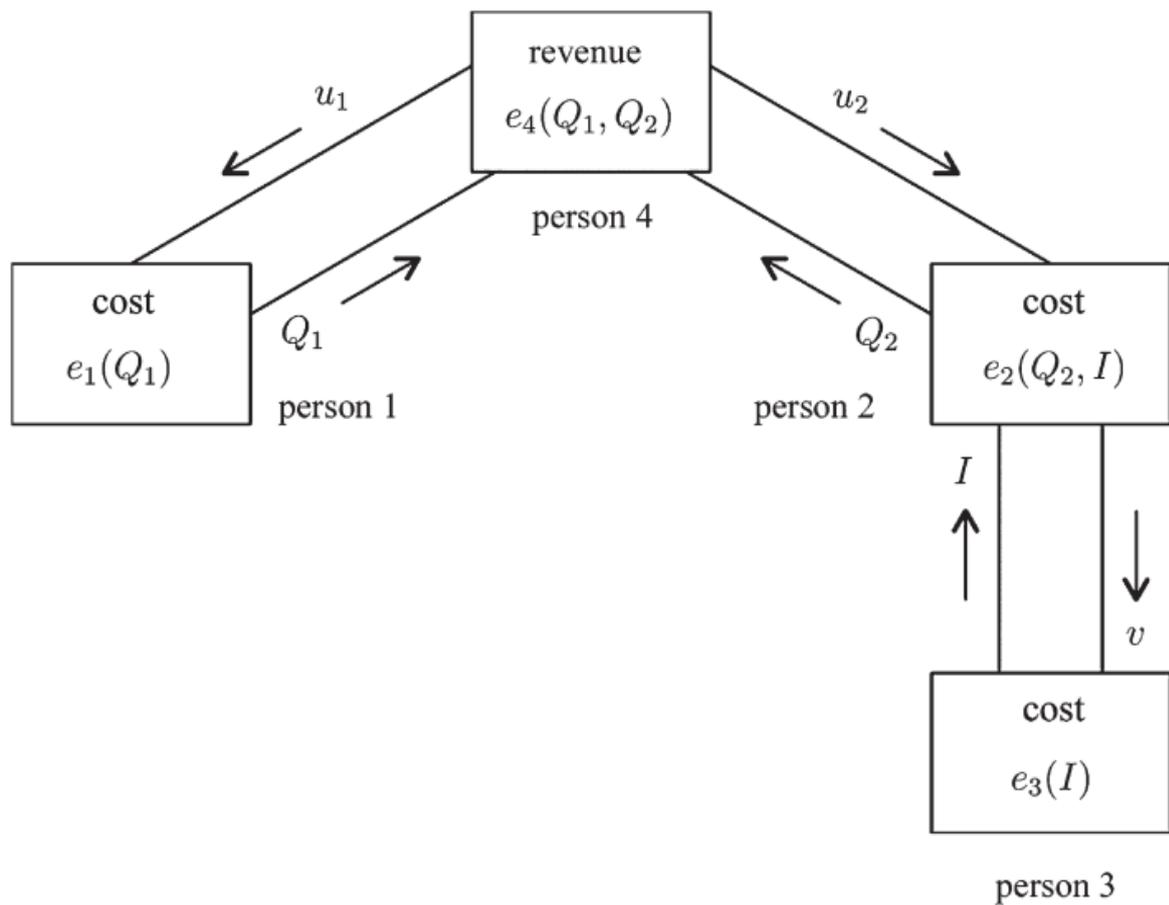


Fig. 2.

The costs of a network mechanism One cost measure is simply the dimension of the entire message space M , the set of possible message arrays. Since each m_{ij} is a vector of $s_{ij} \geq 0$ real variables, we have $\dim M = \sum_{i,j} s_{ij}$. If there are no self-addressed messages (i.e., the sets M_{ii} are empty), and if we think of each of the s_{ij} real message variables as requiring a “pipeline” between i and j , then the dimension of M is the total number of pipelines. In the preceding example there are no self-addressed messages and there are six pipelines.

But it is also of considerable interest to study the n *individual* communication burdens. Person i 's burden is the number of variables he hears, i.e., $\sum_{\{j: M_{ij} \neq \emptyset\}} s_{ij}$. That is also the dimension of the set $P_i(M)$, so we may use the symbol $\dim P_i(M)$ for i 's burden. Note that if there are no self-addressed messages, then $\dim M$ equals half the sum of the individual burdens, since that sum counts each pipeline twice. In the example, s_{ij} is either one or zero and the vector of individual burdens, for 1,2,3, and 4, respectively, is (2,4,2,4). It is natural to ask: is there another network mechanism which also realizes G , but does so with fewer than six message variables, and with an individual-burden vector that dominates (2,4,2,4) — i.e., one person's burden is less than in our price mechanism and no person's burden is higher? Once again a smoothness requirement has to be imposed on the candidate mechanisms, to avoid the smuggling of many numbers into one.

We shall define one such smoothness requirement in a general way, for a class of n -person network mechanisms where (as in our three-plant example) (i) each person i 's environment is a real-valued *valuation function* e_i whose arguments are certain action variables, and (ii) the mechanism realizes a goal correspondence, in which the goal-fulfilling action variables maximize the sum of the valuation functions. (That is the case in our three-plant example if we define the valuation function for 1,2, and 3 to be the negative of the cost function, while 4's evaluation function is his revenue function; then the firm's profit is indeed the sum of the four valuation functions).

We start by considering the k action variables z^1, \dots, z^k . Let each set Z^r — the set of possible values of the action variable z^r — be a closed real interval, so that $Z = Z^1 \times \dots \times Z^k$ is a closed cube in \mathbb{R}^k . Next we shall say that the action variable z^r is one of person i 's *concerns* if it enters his function e_i . A given action variable may be the concern of several persons. (Thus, in our example, person 4's concerns are the action variables Q_1 and Q_2 ; person 2's concerns are Q_2 and I). Let A_i be the index set that identifies i 's concerns, i.e., $A_i = \{r \in \{1, \dots, k\} : z^r \text{ enters the function } e_i\}$. Then a given vector z^{A_i} specifies a value for each of i 's concerns, and Z^{A_i} (a closed cube or closed interval) denotes the set of possible values of z^{A_i} . We now fix the concern sets A_i and we consider mechanisms and environment sets for which the following is true:

$$(\dagger) \quad \left\{ \begin{array}{l} \text{The local environment set for person } i \text{ is } E_i = \\ \{e_i : e_i \text{ is a concave function from } \tilde{Z}^{A_i} \text{ to } \mathbb{R}^+\}, \text{ where } \tilde{Z}^{A_i} \text{ is an} \\ \text{open convex set that includes } Z^{A_i}. \end{array} \right.$$

Now consider the goal correspondence Π defined by:

$$\Pi(e) = \{z \in Z : z \text{ maximizes } \sum_{i=1}^n e_i(z^{A_i}) \text{ on } Z\}.$$

Consider a subclass of the environment set E , namely $\bar{E} = \bar{E}_1 \times \cdots \times \bar{E}_n$, where each concave valuation function $e_i \in \bar{E}_i$ takes a separable quadratic form. That is to say, if $e_i \in \bar{E}_i$, then

$$e_i(z^{A_i}) = \sum_{a \in A_i} \left[x_i^a \cdot z^a - \frac{1}{2} y_i^a \cdot (z^a)^2 \right],$$

where x_i^a, y_i^a are numbers privately observed by i . So we may identify each local environment e_i in \bar{E}_i by a point in $\mathbb{R}^{2|A_i|}$. (For a finite set H , we let $|H|$ denote the number of elements it contains). Moreover if $e \in \bar{E}$, then the goal-fulfilling action set $\Pi(e)$ has a single element, since there is a unique maximizer of $\sum_{i=1}^n e_i(z^{A_i})$ on the cube Z . We now let $\Pi(E)$ denote that unique maximizer. Call an element e of \bar{E} *interior* if that unique maximizer is an interior point of Z . Let $t(a)$ denote the set of persons concerned with the action variable z^a , i.e., $t(a) = \{i : a \in A_i\}$. Assume (to avoid trivial cases) that each set $t(a)$ has at least one member. Write the action k -tuple $\Pi(e)$ as $(\Pi^1(e), \dots, \Pi^k(e))$. It is quickly checked that for an interior environment $\bar{e} \in \bar{E}$, we have, for every $a \in \{1, \dots, k\}$

$$\Pi^a(e) = \frac{\sum_{i \in t(a)} x_i^a}{\sum_{i \in t(a)} y_i^a}.$$

We are now ready to define our smoothness requirement.

Consider a network mechanism $\langle M, (g_1, \dots, g_n), (h^1, \dots, h^n) \rangle$ on the environment set defined in (\dagger) . Let each message m_{ij} in a nonempty set M_{ij} be a vector of real numbers. The mechanism is *smooth on the separable quadratic subset \bar{E}* if for some interior environment $\bar{e} \in \bar{E}$, there exists a neighborhood $U(\bar{e})$ and a continuously differentiable function $r : U(\bar{e}) \rightarrow M$, such that

$$\text{for all } e \in U(\bar{e}) \text{ we have } g_1(r(e), e_1) = 0, \dots, g_n(r(e), e_n) = 0.$$

Thus the function r identifies an equilibrium message for each e in the neighborhood $U(\bar{e})$, and that message varies in a continuously differentiable fashion as we move away from \bar{e} . Using a variant of the uniqueness argument sketched in Section 2.2.8 above, we obtain a lower bound on each person's communication burden. The following can be shown:

Proposition A

Suppose $\langle M, (g_1, \dots, g_n), (h^1, \dots, h^n) \rangle$ is a network mechanism on the environment set E defined in (\dagger) , and each message m_{ij} in every nonempty set M_{ij} is a vector of real numbers. If the mechanism is smooth on the separable quadratic subset \bar{E} and realizes the goal correspondence Π on E , then the number of real message variables that each person hears is at least twice the number of his concerns, i.e., for each person i we have $\dim P_i(M) \geq 2|A_i|$.

If we now return to our four-person three-plant example, we see that the concern-set sizes are $(1, 2, 1, 2)$ for persons 1, 2, 3, 4, respectively. But (as already noted) the vector of message variables heard is $(2, 4, 2, 4)$. So Proposition A tells us that no smooth mechanism whose equilibrium messages yield actions that maximize the firm's profit can improve on the Figure-2 price mechanism with regard

to any person’s communication burden. That is a strong result in favor of price mechanisms. Is it confined to situations that do not depart significantly from our example? The answer is not known. In particular, suppose we consider Π -realizing mechanisms in which each individual burden need not be minimal but instead a weaker efficiency requirement is met. Call a mechanism *efficient* in a class if no other mechanism in the class has a lower burden for some person and not higher burdens for the others. The following challenge is unmet.

RESEARCH CHALLENGE # 1: Consider the class of all Π -realizing mechanisms on the above environment set E which are smooth on the separable quadratic subclass \bar{E} . If such a mechanism is efficient in that class, is it always possible to write it so that it becomes a price mechanism, where each message variable can be interpreted as a price or a quantity, and the agreement functions express divisional profit maximization?

Other fundamental results on network mechanisms concern the size of the overall message space M rather than the individual burdens. These results require a stronger condition than the smoothness we have defined. Call the stronger condition *regularity on the separable quadratic subclass \bar{E}* .²² For regular mechanisms, there is a useful lower bound to the size of M . For any Π -realizing mechanism which is regular on \bar{E} we can show that:

$$\dim M \geq 2 \sum_{a \in \{1, \dots, k\}} (|t(a)| - 1).$$

A mechanism is *dimensionally minimal* in a certain class if no other mechanism in the class has a smaller value of $\dim M$. It is of particular interest to know when the communication graph of a dimensionally minimal regular Π -realizing mechanism is *hierarchical*, i.e., the graph is a tree. That may have advantages that are related to incentives (e.g., it may facilitate “control”). But trees may have communication costs that are higher than needed. Using the above lower bound we can characterize the situations where trees turn out to be dimensionally minimal. Before doing so, note that the communication graph of a network mechanism defines a subgraph for every subset of persons. In particular, there is a subgraph for $t(a)$, the set of persons who are concerned with the action variable z^a . If the communication graph of the mechanism is a tree, then that subgraph may or may not be a tree as well. The following has been shown.

Proposition B

There exists a Π -realizing mechanism that is (i) regular on the separable quadratic subset \bar{E} , (ii) dimensionally minimal among all such mechanisms, and (iii) hierarchical, if and only if there is an n -node tree with the property that for every action variable z^a , the tree’s subgraph for the persons concerned with that variable is also a tree.

Propositions A,B, and further propositions that we do not summarize here, all deal with environments that are valuation functions and with a goal correspondence that requires maximization of their sum. The valuation functions have a “public good” property, since a given action variable may

²²For regularity on \bar{E} , we have to add the requirement that the matrix of second partial derivatives of the g_i has full rank at (\bar{m}, \bar{e}) , where $\bar{e} \in \bar{E}$ is the interior environment in our previous smoothness condition and $\bar{m} = r(\bar{e})$; moreover the rank does not change when we vary e in a neighborhood of \bar{e} while keeping the message array constant at \bar{m} .

enter several of them. Are there analogues of these propositions for other classes of environment sets and goal correspondences, where there are no public goods? In particular suppose that person i 's privately observed local environment is a "revenue" function e_i whose argument is a "private" action vector z_i that does *not* enter any other e_j , where $j \neq i$. Suppose, however, that all the action variables have to meet a common resource constraint. We require $z \in C$, where C is some subset of a finite-dimensional Euclidean space. Suppose, in particular, that the set C can be written in the form $\{(z_1, \dots, z_n) : r_1(z_{B_1}) \leq 0, \dots, r_k(z_{B_n}) \leq 0\}$. Here the B_i are index sets and each is a subset of $\{1, \dots, n\}$. The functions r_1, \dots, r_k are fixed and known to all persons. Now we may view person i as being "concerned with" those constraint functions r_t in which his own action vector enters, i.e., those for which $i \in B_t$. Consider the goal correspondence Π^* defined by

$$\Pi^*(e) = \{z : z \text{ maximizes } \sum_{i=1}^n e_i(z) \text{ on the set } C\}.$$

RESEARCH CHALLENGE # 2: Are there propositions characterizing those Π^* -realizing network mechanisms which are regular (in a suitable sense) and are efficient with regard to individual burdens, or are dimensionally minimal? In particular, are there propositions which describe the situations where efficient or minimal mechanisms are hierarchical?

2.3 Finite approximations of mechanisms whose message spaces are continua In most mechanisms studied in the literature, the message space M is a continuum. That is not surprising, since typically the mechanism's final action maximizes some function, in which actions and environments are the arguments. Maximization is most easily studied with the tools of calculus and those tools deal with continua, not with finite sets. Thus in a classic price mechanism, a message specifies prices and proposed quantities. The first order conditions for the required maximization can be expressed in a statement about prices and quantities, and that statement holds at the mechanism's equilibrium messages. The message space is a continuum composed of possible prices and possible quantities.

But continua are not realistic. In practice, one cannot send all the messages in a continuum (e.g., all the points in a real interval). Moreover it may take infinite time to find that message in a continuum which satisfies the required equilibrium conditions. If an organization wants to use a continuum mechanism in a practical way, it has no choice but to *approximate* the continuum mechanism with an appropriate finite mechanism, whose message space is a finite set. The penalty paid for such finite approximation may be an *error*: the actions generated (at equilibrium) by the finite approximation may differ from the goal-fulfilling actions which the original continuum mechanism generates.

If we take the issue seriously, then the following question immediately comes to mind: will the informational advantages of the original continuum mechanism be reflected in its finite approximation? In particular, if we have found (using the tools of smooth mathematics) a continuum mechanism that realizes a given goal correspondence and does so with minimal message space dimension, are finite approximations to that mechanism superior (in an appropriate sense) to finite approximations of a continuum mechanism which also realizes the goal correspondence but has a higher-than-minimal message-space dimension? Does dimension still matter when we turn from continuum mechanisms to their finite approximations?

In fact, for broadcast mechanisms, a theory of finite approximations has been begun²³ and several “dimension still matters” propositions have been established. In these propositions, the environment set is a continuum as well as the message space of the mechanism we are approximating. The view taken is that “nature” is able to choose the organization’s current external environment from some continuum of possible environments, but the continuum message space is “man-made” (for analytic convenience) and we are free to replace it with a finite message space.

The first step is to define a style of approximation. Suppose we are given an n -person broadcast mechanism Λ on an environment set $E = E_1 \times \cdots \times E_n$, with action set A , where both M and each E_i are continua. In particular (as in our introductory discussion of mechanisms in 2.2.1), let M be the Cartesian product $M^{D_1} \times \cdots \times M^{D_n}$. Let each M^{D_i} be a closed D_i -dimensional cube in \mathbb{R}^{D_i} , while each E_i is a closed J_i -dimensional cube in \mathbb{R}^{J_i} . Write the continuum mechanism Λ in agreement-function form: $\Lambda = \langle M, (g_1, \dots, g_n), h \rangle$. Recall that each agreement function g_i has $M \times E_i$ as its domain and M^{D_i} as its range. Thus each g_i has D_i real-valued components, say $g_{i1}, \dots, g_{ik}, \dots, g_{iD_i}$. Let D denote $\sum_{i=1}^n D_i$.

In a finite approximation to Λ , our finite message space, denoted M_ϵ , is the intersection of the D -dimensional continuum message space M with a *mesh- ϵ* lattice of points, which are spaced 2ϵ apart ($\epsilon > 0$). That lattice, denoted S_ϵ^D , is the D -fold Cartesian product of the set

$$S_\epsilon = \{\dots, -2(\ell + 1)\epsilon, -2\ell\epsilon, \dots, -4\epsilon, -2\epsilon, 0, 2\epsilon, 4\epsilon, \dots, 2\ell\epsilon, 2(\ell + 1)\epsilon, \dots\}.$$

Next we replace each agreement function g^i with a new function $g_i^{\eta\epsilon} = (g_{i1}^{\eta\epsilon}, \dots, g_{iD_i}^{\eta\epsilon})$. Each $g_{ik}^{\eta\epsilon}$ is the following two-valued function:

$$\text{for every } m \in M_\epsilon, g_{ik}^{\eta\epsilon}(m, e_i) = \begin{cases} 0 & \text{if } |g_{ik}(m, e_i)| \leq \eta \\ 1 & \text{otherwise;} \end{cases},$$

where $\eta > 0$ is called the *tolerance*. Finally, we have to specify the outcome function of our finite approximation. In the simplest approach, we let the outcome function be the original one, i.e., it is the restriction of h , the outcome function in Λ , to the new finite message space M_ϵ (which is a subset of M). Denoting the new outcome function h^0 , we have $h^0(m) = h(m)$.

Suppose that our new agreement functions satisfy the coverage requirement: i.e., for every $e \in E$, there exists $m \in M_\epsilon$ such that $g_i^{\eta\epsilon}(m, e_i) = 0$, all i . Then the finite mechanism $\Lambda_{\eta\epsilon} = \langle M_\epsilon, (g_1^{\eta\epsilon}, \dots, g_n^{\eta\epsilon}), h^0 \rangle$ is called *the finite exact-outcome approximation of Λ with message mesh ϵ and tolerance η* . To obtain it we have, in effect, rounded off the original functions g_i to a specified accuracy. The accuracy is determined by the the mesh ϵ and the tolerance η . In an alternative approximation of Λ , we do not require the outcome for the message m to be exactly what it was in Λ . Rather we place a mesh- ν lattice on the action set A , so that our finite mechanism’s action set becomes $A \cap S_\nu^\alpha$. We then choose the outcome to be a lattice point that is closest to the action chosen in Λ . Suppose there are α real-valued action variables and that A is contained in a closed cube in \mathbb{R}^α . Then in the *finite rounded-outcome approximation of Λ with message mesh ϵ , action mesh ν and tolerance η* , all elements except the outcome function are the same as those just defined. The outcome function $h^\nu : M_\epsilon \rightarrow A \cap S_\nu^\alpha$ is defined as follows:

²³In Marschak,1987, and Hurwicz and Marschak 2003A, 2003B, 2004.

$$h^\nu(m) = \begin{cases} \text{the element } a = (a_1, \dots, a_\alpha) \text{ of } A \cap S_\nu^\alpha \\ \text{which is closest to } h(m), \text{ where distance} \\ \text{is measured by } \max_{r \in \{1, \dots, \alpha\}} (|h(m) - a_r|) \\ \text{and ties are broken downward.} \end{cases}$$

For both versions, we shall say that our finite mesh- ϵ approximation has the *minimal tolerance property* if η has the smallest value that permits coverage. (Such a smallest value can be shown to exist).

If the mechanism \mathcal{L} is a finite approximation of the continuum mechanism Λ , then for any $e \in E$, we define the *error at e of \mathcal{L}* to be the worst-case distance between the (equilibrium) value of an action variable in the continuum mechanism and its value in the approximation. Let $h^\#$ denote the finite approximation's outcome function. (This is either the exact outcome of the continuum mechanism or it is a rounded-outcome approximation). Since there are α real action variables, the function $h^\#$ has α real-valued components h^a . The error at e of \mathcal{L} is

$$\sup \left\{ |h_j^\#(\bar{m}) - h_j(m)| : m \in M, \bar{m} \in M_\epsilon; g_i(m, e_i) = 0, g_i^{\epsilon\eta}(\bar{m}, e_i) = 0, i = 1, \dots, n; j \in \{1, \dots, \alpha\} \right\}.$$

The *overall error of \mathcal{L}* is $\sup_{e \in E}$ (error at e of \mathcal{L}).

We define the *cost* of the approximation \mathcal{L} to be the number of messages in its finite message space. We impose some regularity conditions²⁴ and obtain the following “dimension still matters” proposition.

Proposition C

Consider two regular continuum mechanisms $\Lambda = \langle M, (g_1, \dots, g_n), h \rangle$ and $\Lambda^* = \langle M^*, (g_1^*, \dots, g_n^*), h^* \rangle$. Each is a mechanism on the same environment set, namely the compact set $E = E_1 \times \dots \times E_n$, and each has an action set which is a subset of \mathbb{R}^α . The two message spaces are distinct: M is a D -dimensional subset of \mathbb{R}^D , while M^* is a D^* -dimensional subset of \mathbb{R}^{D^*} , where $D^* > D$. Let \mathcal{L} be a mesh- ϵ rounded-outcome minimal-tolerance approximation of Λ , and let \mathcal{L}^* be a mesh- $\bar{\epsilon}$ rounded-outcome minimal-tolerance approximation of Λ^* . Suppose that \mathcal{L}^* **costs no more than \mathcal{L}** . Then if ϵ is sufficiently small, the overall error of \mathcal{L}^* exceeds the overall error of \mathcal{L} .

Note that the proposition does not require us to specify goal correspondences realized by the two continuum mechanisms. But it can certainly be applied to two continuum mechanisms which realize the same goal correspondence. In that case it tells us, informally speaking, that if we want to come close to achieving our goal, then the low-dimensional mechanism Λ is a better candidate for finite approximation than the high-dimensional mechanism Λ . We achieve lower overall error, for a given “budget” if we approximate Λ than if we approximate Λ^* . So dimension indeed continues to matter.²⁵

It remains open whether or not there are similar propositions for network mechanisms.

²⁴In a regular mechanism each function g_{ik} is continuously differentiable. Moreover there exists a number $\delta > 0$ such that for all $\delta_{ik} \in [-\delta, \delta]$ and for all e in E , there is a unique message m satisfying all the equations $g_{ik}(m, e_i) = \delta_{ik}$.

²⁵Another proposition lets us be “kind” to the high-dimension mechanism by permitting its approximation to have the original exact outcomes, while we are “harsh” to the low-dimension mechanism by requiring its outcomes to be rounded off. Even so, it is better to approximate the low-dimension mechanism. This proposition, however, requires the high-dimension mechanism to have a “projection” property: each message is a pair (a, q) and the outcome function is $h(a, q) = a$.

RESEARCH CHALLENGE #3: If two continuum network mechanisms realize the same goal correspondence and have suitable regularity properties, but the second has a higher number of individually addressed message variables, is it better to approximate the first? (Do we achieve a smaller overall error, for a given “budget”, when we do so?)

Returning to the case of broadcast mechanisms, a very difficult question remains unaddressed. In constructing finite broadcast mechanisms, we have confined our attention to finite mechanisms which approximate regular continuum broadcast mechanisms. Do we ignore certain efficient finite mechanisms when we do so?

RESEARCH CHALLENGE # 4: Given a goal correspondence G , can there be a finite broadcast mechanism which is NOT an approximation of any regular G -realizing continuum broadcast mechanism but makes better use of a given “budget” (achieves lower overall error with respect to G while not using more messages) than any such approximation?

2.4 The dynamics of a mechanism Return now to the dynamic form of a broadcast mechanism, which started our discussion in 2.2.1. The mechanism is a quadruple $\langle (M^1, \dots, M^n), (m_1^0, \dots, m_n^0), (f_1, \dots, f_n), h \rangle$. It defines a difference-equation system, namely:

$$m^t = f((m_1^{t-1}, \dots, m_n^{t-1}), e_i), i = 1, \dots, n,$$

with an initial message $m^0(e) = (m_1^0(e_1), \dots, m_n^0(e_n))$. We have been interested thus far in the achievements of the mechanism once it has reached an equilibrium message. A difficult missing piece in the story has to do with the stability properties of the difference-equation system. We would like the action taken once an equilibrium message is reached to meet a specified goal, but we would also like the difference-equation system to display some sort of convergence to the system’s equilibria. *Do the mechanism’s informational requirements (e.g., its message-space size) grow if we require stability as well as goal realization?* One can construct examples where the answer is Yes.²⁶

Some progress has been made on this question when the difference equations are replaced by differential equations. In particular, Jordan (1995) developed the following new mechanism concept.²⁷ For each person i there is a real-valued “control message” c_i , whose possible values comprise a set C_i . There is also a broadcast “state message” $m = (m^1, \dots, m^q)$, with q real components, which is continually adjusted as a function of both $c = (c_1, \dots, c_n)$ and m . It is not required that $q = n$. We

²⁶A two-person example, due to Reiter and discussed in Hurwicz, 1986, is as follows. Person i ’s local environment is a real-number pair $e_i = (e_i^1, e_i^2)$. The action set is \mathbb{R} . For each e the goal specifies a unique action for all e such that $e_1^1 \neq e_2^1$, namely $F(e) = \frac{e_1^1 e_2^2 - e_2^1 e_1^2}{e_1^1 - e_2^1}$. If we do not require stability, we can realize the goal with a two-message-variable mechanism. The typical message is a pair $m = (m_1, m_2)$. The mechanism’s difference-equation system is $m_1^t = 2m_1^{t-1} - e_1^1 m_2^{t-1} - e_2^2$, $m_2^t = m_1^{t-1} + m_2^{t-1} - e_2^1 m_2^{t-1} - e_2^2$. The outcome function is a projection: $h(m) = m_1$.

But this system fails to satisfy the following local stability requirement: for m_0 sufficiently close to an equilibrium value of m , the system should converge to that equilibrium. Moreover, if we seek any two-message-variable difference equation system which realizes F at equilibrium and uses the projection outcome function, we find that if the functions f_1, f_2 have continuous partial derivatives, then the system is not locally stable. On the other hand, a locally stable F -realizing mechanism with *four* message variables can be constructed.

²⁷See also Jordan, 1986; Mount and Reiter, 1987; Saari and Simon, 1978.

have the differential-equation system

$$c_i = f_i(m, e_i), i = 1, \dots, n; \frac{dm^j}{dt} = \alpha_j(c, m), j = 1, \dots, q,$$

where t is a time point. Note that the function α_j does not have e_i as an argument. The interpretation is that each person i continually adjusts his own control variable c_i in a privacy-preserving manner. We do not specify who adjusts a given component of m itself, but that adjustment does not require direct knowledge of the privately observed e_i s. In the case where $q = n$ and i has responsibility for the message variable m^i , we have a complete privacy-preserving scenario: person i observes the entire broadcast message m , adjusts his own part of m , and chooses his control variable as a function of the broadcast message and his local environment.

The general question is then as follows. Suppose we are given a particular message correspondence μ from $E = E_1 \times \dots \times E_n$ to the state-message space M (the set of possible values of m). We are interested in this correspondence because it realizes some goal (i.e., if $m \in \mu(e)$, then there is an action $h(m)$ which lies in a set $G(e)$ of goal-fulfilling actions), but the goal itself is not part of the research question. Instead we ask: how large do the sets C_i have to be if the equilibria of the differential-equation system always lie in the set $\mu(e)$ and the system has a local stability property for every e ? How large, in other words, is the extra informational cost of stability? Several general results characterize the required size of the C_i .

One would like applications of these results to classic questions, notably the informational merits of price mechanisms. Consider, once again, the n -person L -commodity exchange economy and the Pareto-optimal message correspondence. If we construct a privacy-preserving dynamic mechanism which uses prices, has suitable regularity properties as well as local stability, and achieves Pareto-optimality at the equilibrium message, then is its total message space (i.e., $M \times C$) dimensionally minimal among all dynamic mechanisms with those properties? Much remains to be learned about this question, but for certain reasonable classes of economies, and certain versions of the dynamic price mechanism the answer is Yes. For one such class, Jordan (1995) studies dynamic mechanisms in which the message m specifies current trades, the control variable c_i is a vector of i 's "demand prices" (marginal utilities), and the adjustment rules for m (i.e., the functions α_j) adjust the trades so that they are Pareto-improving. It is shown that if we delete the stability requirement for such mechanisms, then a lower bound on the dimension of each C is $n(L-1)$. It is then shown that stability can be achieved without increasing the dimension of C beyond $n(L-1)$. In other types of dynamic mechanism, the control variables are trades as well as prices. It turns out that if such a mechanism is formulated so that privacy is preserved and local stability is achieved, then C has to be very large and the stabilized price mechanism may no longer be minimal among all such mechanisms.

Note that for a *finite* broadcast mechanism we have an upper bound on the time required to achieve equilibrium —namely the time required to announce all the messages in the finite set M , in some pre-determined sequence. If M is large, that upper bound is of little interest. We may then want to choose the sequence with care, perhaps by approximating (in some suitable way) the rules of a locally stable mechanism in which the sets M and C are continua. Such approximation remains to be studied.

Note also that if we truncate a difference-equation broadcast mechanism $\langle (M^1, \dots, M^n), (m_1^0, \dots, m_n^0), (f_1, \dots, f_n), h \rangle$ after T steps, then we have defined a new privacy-preserving broadcast mechanism in

which every possible broadcast message describes a possible T -step “conversation”, namely a proposed sequence of announcements m^0, m^1, \dots, m^T . Person i agrees to the proposed conversation if he finds — given his e_i , given the proposed sequence of announcements by others, and given his function f_i — that the proposed sequence $m_i^0, m_i^1, \dots, m_i^T$ is exactly what he would announce. So, using the uniqueness technique discussed in 2.2.8, one could study the goal correspondence realized by the T -step truncation and could ask whether there are broadcast mechanisms which realize the same goal correspondence using a smaller message space.

2.5 Constructing an informationally efficient mechanism The informationally efficient (or cheap) mechanisms that have appeared in the literature are, in a sense, historical accidents. They are mechanisms which allocate resources so as to satisfy a goal correspondence that expresses Pareto-optimality, or perhaps profit maximization. Typically each message proposes prices and quantities, and agreement to the message by all persons means that its prices and quantities correspond to the required first-order maximization conditions. With some exaggeration one might say that the literature started by asking: “what is the precise information-related question to which ‘price mechanisms’ (or ‘competitive mechanisms’) is the answer?” That was a natural challenge, given the long history of sweeping but never rigorously defended claims about the price (or competitive) mechanism’s informational merits.

But what if prices had not yet been discovered? Imagine looking for low-cost mechanisms among all those that realize a goal correspondence. If the goal required Pareto optimality, then a search for such mechanisms would eventually discover mechanisms that use prices. How might such a search proceed?

Two new books, one by Hurwicz and Reiter (2006), and the other by Williams (forthcoming), deal with this fundamental puzzle. It would be futile to attempt any kind of summary here. But we can roughly visualize one of the main issues in the Hurwicz/Reiter agenda by going back to the “rectangle” definition of a mechanism in 2.2.3 and the two-person, three-message example in Figure 1 of 2.2.2. Suppose that our environment sets are $E_i = [0, 1], i = 1, 2$ and that we have an action variable a that takes three values, namely u, v, w . Suppose we have not yet constructed a mechanism which yields (at equilibrium) a value of the action for every $e \in E = E_1 \times E_2$. Relabel the three rectangles m_1, m_2, m_3 in Figure 1 as U, V, W respectively. Let those rectangles (which overlap at boundary points) define the goal correspondence G that we want to realize. Thus

$$G(e) = \{a : a = u \text{ if } e \in U; a = v \text{ if } e \in V; a = w \text{ if } e \in W\}.$$

We may call U the *level set of the correspondence G for the action u* , and similarly for V and W . Formally the level set corresponding to the action a is $G^{-1}(a) = \{e \in E : a \in G(e)\}$. Consider the mechanisms which realize G and suppose that we write all of them in rectangle form, so that there is a generalized rectangle σ_m for every message m , and m is an equilibrium message for all the environments in σ_m . (Recall that a generalized rectangle is a set of environments e that is the Cartesian product of its E_1 -projection and its E_2 -projection). There are many such mechanisms, but they all use more generalized rectangles (messages) than we need *except* the three-message mechanism defined in 2.2.2. In that mechanism there are just three rectangles σ_m , namely the same three rectangles U, V, W that define the goal correspondence. An inefficient G -realizing mechanism might, for example, needlessly add a fourth message, by dividing the rectangle $\sigma_{m_3} = W$ into two further generalized rectangles.

The three-message mechanism defined in 2.2.2 is efficient: *it covers each level set of our goal correspondence with rectangles σ_m in such a way that the total number of rectangles σ_m is minimized.* To illustrate further, suppose we modify our example by allowing just the two actions u and w and letting w be goal-fulfilling for all environments that lie in W or V . (The action u remains goal-fulfilling for all $e \in U$). Now the level set for the action w is no longer a (generalized) rectangle; it is now the union of W and V . Nevertheless, an efficient goal-fulfilling mechanism has to cover that level set (as well as the level set corresponding to u) with generalized rectangles σ_m and it has to do so in a minimal way. The efficient mechanism will again require three messages. It will cover the level set corresponding to the action w with *two* generalized rectangles, namely the rectangles W and V .

So the search for an efficient goal-realizing privacy-preserving mechanism requires us to inspect the level sets, to find a way of covering each of them with generalized rectangles, to find a way of indexing (labelling) each of our generalized rectangles with an index m , and to do all this while keeping the size of M (the set of values if the index m) as “small” as possible. The basic set-theoretic properties of such a mechanism-designing algorithm are worked out in the Hurwicz/Reiter book. The algorithm yields efficient mechanisms whatever the goal correspondence may be, whether the action set and the environment sets are finite or are continua.

The book of Williams also deals with a mechanism-designing algorithm, but from a very different point of view. Smoothness requirements are imposed on the goal and on the candidate mechanisms. Tools of differential topology are used rather than purely set-theoretic tools. Some of the results imply that the agreement functions of an efficient goal-realizing mechanism can be found by solving an appropriate system of partial differential equations.

Once these two books are understood, they may open a massive research agenda for the designers of practical computer-friendly algorithms that construct mechanisms (protocols). It remains to be seen, for example, whether the general results in these books will eventually allow a computer to generate protocols that yield the minimal-length dialogues studied in the computer-science communication-complexity literature (briefly discussed above in 2.2.8). At present that literature finds bounds on the length of the dialogues but does not tell us how to construct the minimal-length dialogues themselves.

2.6 Finding a best action rule (outcome function) once a mechanism has conveyed information about the environment to each person: the methods of the Theory of Teams The central problem studied in the Theory of Teams (J. Marschak and R. Radner, 1972) is the choice of a rule that tells each member of an organization — called a team — what action to choose, given certain information about the organization’s randomly changing environment. The rule has to maximize the expected value of a payoff function whose arguments are the environment and the team action. Mechanisms, as we have defined them, do not appear in the statement of the central problem studied in the Theory of Teams, but they are part of the story which implicitly *precedes* that problem.²⁸

Here is one version of the central n -person team problem. The team has to chose an action, namely a vector $a = (a_1, \dots, a_n)$, where a_i is the responsibility of person i . Let A_i denote the set of possible values of a_i and assume that every n -tuple a in $A_1 \times \dots \times A_n$ is a possible team action. The team

²⁸In Chapter 8 of *The Economic Theory of Teams*, there is a discussion of “networks”, with a number of examples. One may interpret the network concept developed in that chapter as a mechanism in our sense.

earns a (real-valued) payoff $H(a, e)$, where $e = (e_1, \dots, e_n)$ is a vector of random local-environment variables. We study a given *information structure*, specifying what each person knows about a given e . The set of possible values of e is denoted E . Let η_i be a function from E to a signal set Y_i . Let E be a subset of a (finite-dimensional) Euclidean space, and similarly for each Y_i and each A_i . A probability distribution P on E is given. In the information structure $\eta = (\eta_1, \dots, \eta_n)$, person i observes the signal $y_i = \eta_i(e)$ when the environment is e .

The signal $\eta_i(e)$ might be a vector, and e_i might be one of its components. In our discussion of privacy-preserving broadcast mechanisms, each e_i was automatically known to one person, namely person i . We may interpret $\eta_i(e)$ as the information about the current e that person i possesses once the mechanism has terminated. Then $\eta_i(e)$ indeed includes e_i , but it also describes the information about the other e_j that is revealed to i by the mechanism's terminal message.²⁹

For a given information structure η , we consider the possible *team action rules* $\alpha = (\alpha_1, \dots, \alpha_n)$, where α_i is a function from Y_i to A_i . Thus the team action is $(\alpha_1(y_1), \dots, \alpha_n(y_n))$ when the signal vector is $y = (y_1, \dots, y_n)$.³⁰ An action rule $\hat{\alpha}$ is *team-best for the payoff function H and the information structure η* if $\mathcal{E}H\left((\hat{\alpha}_1(\eta_1(e)), \dots, \hat{\alpha}_n(\eta_n(e))), e \right) \geq \mathcal{E}H\left((\alpha_1(\eta_1(e)), \dots, \alpha_n(\eta_n(e))), e \right)$ for all action rules α , where \mathcal{E} again denotes expectation. A necessary condition for an action rule $\hat{\alpha}$ to be team-best for H, η is that it be *person-by-person-satisfactory (pbps) for H, η* . That means that for every i and every y , the action $\hat{\alpha}_i(y_i)$ maximizes the conditional expected value $\mathcal{E}\left[H\left((\hat{\alpha}_1(\eta_1(e)), \dots, \hat{\alpha}_{i-1}(\eta_{i-1}(e)), a_i, \hat{\alpha}_{i+1}(\eta_{i+1}(e)), \dots, \hat{\alpha}_n(\eta_n(e))), e \right) \middle| \eta_i(e) = y_i \right]$ on the set A_i . If H is differentiable and strictly concave in a for each e , then the pbps condition is sufficient as well as necessary.

Consider the case of a team with the linear-quadratic team payoff function $W(a, e) = 2a'e - a'Qa$. Here A_i is the real line. The random variables e_1, \dots, e_n also take values in the real line and they are independently distributed with finite second moments; $Q = ((q_{ij}))$ is an n -by- n symmetric positive definite matrix. The function W is differentiable and strictly concave in a for each e . Accordingly the pbps condition is both necessary and sufficient. A best team action rule is linear. Its coefficients can be found by solving a system of linear equations.

That permits the explorations of information structures with interesting organizational properties. (For some explorations, it is also helpful to assume that each e_i is normally distributed). For example in a “management-by-exception” information structure, each person $i \neq n$ knows only his own e_i . But person n is a manager who learns the value of every e_j whenever that e_j lies in a specified “exceptional” region. A best team action rule α will take advantage of the manager's exceptional information. In a variant of this structure, an “emergency conference” of all n persons is called whenever some person j observes an exceptional value of e_j . When that happens, all persons learn that exceptional value. We can vary the exceptional regions and in each case we can compute the structure's “gross” performance,

²⁹Thus we can express an information structure as a message correspondence μ , where m lies in $\mu(e)$ if and only if $m = (\eta_1(e), \dots, \eta_n(e))$.

³⁰Appropriate measurability assumptions have to be made when E is not finite. They guarantee that (i) P implies a probability distribution on the set $\alpha_i(\eta_i(E))$ for every i , and (ii) $H\left((\alpha_1(\eta_1(e)), \dots, \alpha_n(\eta_n(e))), e \right)$ has a finite expected value.

i.e., expected team payoff when a best action rule is used. The gross performance of one interesting structure can be compared to that of another. The *cost* of each structure, however, needs to be measured in some consistent way if we are to characterize structures whose net performance is high.

Unfortunately, it is difficult to obtain similar explicit results about gross performance once we leave the linear-quadratic case. It is difficult even though we remain in the class of payoff functions W that are strictly concave in a for each e , so that the pbps condition is sufficient as well as necessary. We are, after all, performing a search in function space, and that is difficult unless the functions can be parametrized in some convenient way. Nevertheless it would seem plausible that for some class of concave payoff functions, algorithms could be constructed that come close to yielding best team action rules.

RESEARCH CHALLENGE # 5: Construct an algorithm that yields best (or nearly best) team action rules for a wide class of concave functions W (containing the linear-quadratic function and others as well), and does so for a wide class of information structures and probability distributions on the environment variables

2.7 Designing an organization “from scratch”: choosing its members, what each observes, and the speak-once-only mechanism that they use³¹ So far we have assumed (without comment) that the n members of our organization are already in place, and that we have no choice as to the external variables (the local-environment variables) which a given member observes. That is natural if our organization is an economy and its members are a given collection of consumers and producers. There is a natural privately observed external variable for each consumer; it describes her individual preferences and perhaps her individual resource endowment. For a producer, the natural external variable describes his current technology. Similarly, when modeling a firm with several divisions, it is natural to let the members of the organization be the division managers, who are already in place. For each manager, the natural privately observed external variables are those that characterize his production technology.

Once we leave such settings, we may want to enrich our modeling toolkit. We may want to take the view that the organization does not yet exist but is being designed. The designer has a clean slate. He is given external variables e_1, \dots, e_k , and a set E which contains their possible values. He is given a set A of possible organizational actions. He is given a goal, which identifies at least one appropriate action in A for each external environment $e = (e_1, \dots, e_k)$ in E . But he can *choose* the following:

- The size and composition of the collection of persons who make up the organization.
- The identity of the person who will observe each external variable; some external variables may be observed by more than one person; some persons may not observe any external variable.
- The *speak-once-only mechanism* which the organization uses to find a new action when the environment changes.

³¹Some of the ideas in this section grew out of conversations with Ilya Segal. Some of the results in Section 2.7.6 are due to Jeff Phan. The discussion of delay in 2.7.10, as well as the result in Section 2.7.7, are largely due to Dennis Courtney.

2.7.1 Speak-once-only mechanisms: introductory comments In a speak-once-only mechanism, the newly changed external variables are observed by their designated observers. Each observer sends forward to others a message based on those observations and then stays silent. The recipients of those messages send forward messages to others, based on the messages they have received and on their external observations (if they are designated external-variable observers), and then stay silent; the recipients of *those* messages send forward still further messages and then stay silent. And so on. When all sending has stopped, one member, called *the action taker*, takes an action based on the messages he has received (and on his own external observations if he is a designated external-variable observer). The mechanism realizes a given goal if the action taker’s action is always goal-fulfilling for the environment e which initiates the process.

Our speak-once-only requirement is certainly restrictive. For full generality, we would let the designer choose a mechanism in which each sender sends messages to others at each of a sequence of steps. But our clean-slate assumption, where the designer chooses the size and composition of the organization, the observing assignments, *and* the mechanism used to generate new actions presents a formidable challenge. Confining attention to speak-once-only mechanisms is a reasonable compromise, if one wants to start learning something about the structure of the designer’s chosen organization.

To help motivate a research agenda on speak-once-only mechanisms consider the following question:³²

What might be inefficient about a one-person mechanism, wherein a single person observes all the external variables and then finds a goal-fulfilling action?

A first step in making sense out of this question (and more complicated ones) is to choose some cost measures for a speak-once-only mechanism, so that “efficiency” (and inefficiency) can be defined. First we assume that each external variable e_k is real-valued, and that every message in the mechanism is a vector of real numbers. For simplicity we assume that for every external vector e , there is a unique goal-fulfilling action $F(e)$ where $F(e)$ is a real number. Consider the following three cost measures for a speak-once-only mechanism:

- The number of persons.
- Each person’s *burden*, defined as the number of real variables observed or received.
- The mechanism’s *delay*, i.e., the total elapsed time until the action-taker has computed F , given that (i) no one sends message variables or computes F until he has finished doing all the receiving and observing that the mechanism requires of him, and (ii) it takes one time unit for a person to observe or receive one real variable, but no extra time is required for the sender of a message to compute and send it or for the action taker to compute F .

³²One can ask a version of this question, even if one specifies that n , the number of persons, is greater than one and cannot be changed by the designer, and that each external variable is observed by one and only one person. A mechanism in which one person does essentially all the work, would be one in which, say, Person 1 collects full information about e from all the others and thereupon finds the goal-fulfilling action.

2.7.2 The approaches of Radner/Van Zandt and Mount/Reiter Papers by Radner (1993), Radner and Van Zandt (1995), and Van Zandt (1997, 1998, 2003A, 2003B) study a class of speak-once-only mechanisms. The messages transmitted are composed of real numbers, and there is an action taker who acquires enough information from others to compute a real-valued goal function F of external real variables e_1, \dots, e_k , each of which is observed by someone. The cost measures are number of persons and delay.³³ But the function F has to have the form $F = H(H_1(e_1) * H_2(e_2) * \dots * H_k(e_k))$, where $*$ is an associative operation. Moreover, any person i , who receives, say, the real numbers v_1, \dots, v_r , is only able to compute $v_1 * v_2 \dots * v_r$, where $*$ is the same associative operation. *It takes one time unit to perform the operation $*$ and no time is required for communication*, so the Radner/Van Zandt work may be viewed as a model of organizational computing. It can be shown that if F is differentiable in the k variables, then no generality is lost if the operation $*$ is constrained to be addition. That is to say, any differentiable F having the above associative form can be rewritten so that it becomes a function of the *sum* of certain terms $\tilde{H}_i(e_i)$. Non-differentiable goal functions with the associative property include $\max(e_1, \dots, e_k)$, where e_1, \dots, e_k are real numbers.

In the book by Mount and Reiter (2002)³⁴, the goal function $F(e_1, \dots, e_k)$ can be any analytic function. The problem is to assemble a *minimal-delay F -computing network* of processors (persons). Each of them receives inputs from other processors; each of those inputs consists of at most d real numbers. In *one time unit* a processor computes a certain analytic function of the numbers received and sends the result to one or more other processors. The class of functions that the processors are able to compute is a primitive of the model. If F itself has d real arguments, and is one of the available functions, then the problem is trivial, since a single processor can then compute F itself in one time unit. Instead, one seeks networks (directed graphs) that compute F with minimal delay when the functions available for each processor belong to an interesting class. The network is not required to be a tree, so cycles are permitted. But it is shown (under weak assumptions) that no generality is lost if one confines attention to trees. A tree has no cycles and so it is, in our terminology, an F -computing speak-once-only mechanism.

While similar models of networks of processors (automata) have been studied by others, the Mount/Reiter research has a major novelty: the possible values taken by the input numbers, by F , and by each processor's function, can be continua rather than finite sets. The research reported in the book does not explicitly seek efficient combinations of number of processors, individual burdens, and delay. Instead it seeks to characterize the networks (trees) that are capable of computing certain

³³While Radner (1993) and Radner and Van Zandt (1995) consider a “one-shot” situation, where each environment vector $e = (e_1, \dots, e_k)$ is processed before a new e appears, papers by Van Zandt (1999, 2003a, 2003b) go on to study the much more challenging situation where a new e arrives before the previous one has been completely processed, and the successive e s follow a stochastic process. The current computed action, which is a function of the previous e , is then somewhat obsolete. The penalty due to obsolescence is studied. In particular, Van Zandt (2003b, sketched also in Van Zandt 1998, Section 3.3) studies the performance of an organization which repeatedly allocates resources and thereby earns, in each time period, a payoff that is a quadratic function of the allocation and of that period's environment vector. But the information used by the allocator reflects an earlier period's environment vector, since it takes time for that information to reach him. Results are obtained by exploiting the fact that the mean of k variables is a sum and hence it can be computed by a sequence of associative operations. Early discussions of obsolescence as one of the costs of a mechanism appear in Marschak, (1959) and (1972).

³⁴An easily accessible summary of some of the book's ideas is given in Mount and Reiter, 1998.

functions F with a given delay, when the processor functions obey key conditions like twice differentiability. Some propositions, for example, concern the number of processors to which each processor sends. There are also results about the relation between the delay for a continuum-valued F and the delay for each of a sequence of finite-valued functions that approximates F . It turns out that the former delay is the limit of the latter sequence of delays.

The Mount-Reiter model is a highly innovative way to study the complexity of a given goal function F . Since it does not treat number of processors (persons) and individual burdens as explicit costs, it does not easily lend itself to the study of some of the efficient-organization questions that we shall now consider. Note that if each processor's function is required to be *addition*, and if number of processors *is* a cost element, then the Mount/Reiter model becomes the Van Zandt/Radner model. Note also that in both models one finds, in many interesting cases, that there is an upper limit to the useful number of persons (processors). Going above that number does not further decrease delay.

2.7.3 A formal definition of a speak-once-only mechanism To define a speak-once-only mechanism, a directed graph has to be specified. It will be useful to modify the conventional terminology of directed graphs to fit our context. A directed graph is defined by a set of nodes and a set of ordered pairs whose elements belong to the set of nodes. We shall say that the first node in the pair *sends to* the second and the second *receives from* the first. There is an *arc* between them. Then node i 's *received-from set*, denoted R_i is the set of nodes from which i receives, while i 's *sent-to set*, denoted S_i is the set of nodes to which i sends. We shall call a node j a *leaf* if it sends but does not receive. We call a node r a *root* if it receives but does not send. (Usually the terms leaf and root are reserved for trees, but our graph need not be a tree).

An n -person *speak-once-only mechanism on the local-environment sets* $E_1 \subseteq \mathbb{R}^{D_1}, \dots, E_k \subseteq \mathbb{R}^{D_k}$, with action space $A \subseteq \mathbb{R}^\alpha$ is a pair $\Lambda = \langle G, (\vec{\rho}_{k+1}, \dots, \vec{\rho}_n) \rangle$. Here G is a graph and $\vec{\rho}_{k+1}, \dots, \vec{\rho}_n$ are vectors of *sending functions*. If $j \in \{1, \dots, k\}$, then e_j is an environment variable; its possible values comprise the set $E_j \subseteq \mathbb{R}^{D_j}$. We specify that:

- G is a connected directed graph with nodes denoted $1, \dots, k, k+1, \dots, n$, where nodes $1, \dots, k$, with $k < n$, are leaves (they correspond to the environment variables), node n (the action taker) is the only root, there are no directed cycles, and there is at least one directed path from every node to the root.
- for $i = k+1, \dots, n-1$, the vector $\vec{\rho}_i$ has one sending function, $\rho_{i\ell}$, for every ℓ in the sent-to set S_i . Person i sends $D_{i\ell}$ real variables to person ℓ in S_i , so the range of $\rho_{i\ell}$ is $\mathbb{R}^{D_{i\ell}}$. The domain of $\rho_{i\ell}$ is the set of possible values of the variables i receives or observes, i.e., the domain is a subset of $\mathbb{R}^{\sum_{t \in R_i} D_{ti}}$, where $D_{ti} \equiv D_t$ if $t \in \{1, \dots, k\}$.
- The remaining vector of functions is $\vec{\rho}_n$. It has a single component, denoted simply ρ_n , which yields the organization's action. The range of ρ_n is $A \subseteq \mathbb{R}^\alpha$ and the domain is a subset of $\mathbb{R}^{\sum_{t \in R_n} D_{tn}}$.
- The function ρ_n , and every function ρ_{ij} , where $i \in \{k+1, \dots, n-1\}, j \in S_i$, satisfies a smoothness requirement, e.g., it is differentiable at all points of its domain.

Thus any person i in $\{k+1, \dots, n\}$ whose received-from set R_i contains $t \in \{1, \dots, k\}$ is an *observer of the environment variable* e_t . In addition, person i receives a message, say m_{ji} , from every $j \in R_i$ with $j > k$. The message m_{ji} is a vector with D_{ji} real components. Person i sends a message, say m_{ij} , to every person $j \in S_i$; that message is determined by the function ρ_i . The action taker n receives message variables, say m_{jn} , from certain persons $j > k$. His choice of the action a in A is determined by the function ρ_n . For person i in $\{k+1, \dots, n\}$ we shall define i 's *individual burden in the mechanism* Λ to be $\sum_{t \in R_i} D_{ti}$.

2.7.4 An illustration, which leads to three questions To illustrate, consider four real and positive local-environment variables. It will be notationally convenient to call them w, x, y, z . Suppose we want to realize the following goal function F^* :

$$F^*(w, x, y, z) = \left(w + \frac{1}{w} + x + \frac{1}{x} + y + \frac{1}{y} + z + \frac{1}{z} \right) + (wxyz).$$

Here is a two-person mechanism that realizes F^* .

FIGURE 3 HERE

Each person's burden is 3 and the delay³⁵ is 5.

We can ask the following questions about the Figure-3 mechanism and other possible speak-once-only mechanisms which also realize F^* :

(1) **If we reduce the number of persons to just one, must the burden and the delay of an F^* -realizing mechanism rise?** Our one-person mechanism would be a five-node tree, with four nodes corresponding to w, x, y and z , and the sole person at the root. Clearly the burden of the sole person will be four and the delay will be four. We actually reduce delay when we move to the one-person mechanism, since no time is used for the processing of messages received from other persons.³⁶

(2) **Is there a two-person F^* -realizing mechanism in which neither person has a burden more than two and delay is not more than four?** The mechanism would improve on the Figure-3 mechanism with regard to delay and it would improve on the one-person tree with regard to burden. Such a mechanism would exist if it were possible to *aggregate* three of the four external variables into one, i.e., if it were possible to write F^* in the form $G(H(w, x, y), z)$, where G, H are real-valued

³⁵Suppose a new (w, x, y, z) occurs. Three time periods then pass. At that point, Person 1 has completed his observing (processing) of his three assigned variables w, x, y and Person 2 has completed his observing of z . Now two more time periods pass. During the first of them, Person 2 receives (processes) the first of the two message variables which Person 1 sends, and during the second, Person 2 receives (processes) the second message variable. At the end of the fifth period, Person 2 has received (from Person 1) the messages he needs, and is able to take the action $F^*(w, x, y, z)$.

³⁶That illustrates a deficiency of our definition of delay, where the computing of the final action takes no extra time, once the action-taker has collected the information needed to do so. On the other hand, if we insisted on measuring computing time, then we would need a detailed model of computing, such as those studied in the Van Zandt/ Radner and Mount/Reiter work.

$$\text{action: } \mathbf{F}^* = w + x + y + z + wxyz + \left(\frac{1}{w} + \frac{1}{x} + \frac{1}{y} + \frac{1}{z}\right)$$

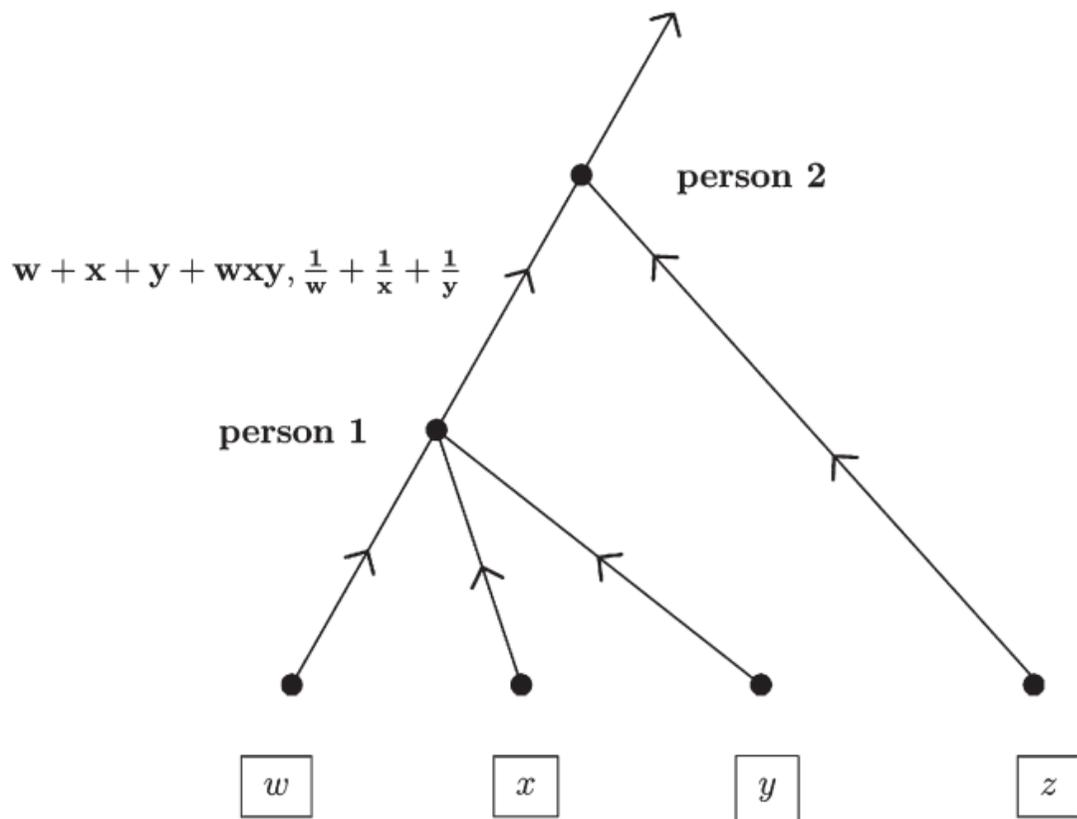


Fig. 3.

functions (with suitable smoothness properties). It is natural to call the function H an *aggregator*. The single number $H(w, x, y)$ contains all the information about w, x, y that is needed in order to compute the goal function. Consider a mechanism in which Person 1 observes w, x, y and then sends the number $H(w, x, y)$ to Person 2, the action taker, who observes z and is able to compute the action $F^*(w, x, y, z)$ once he receives the number H . After three periods, Person 1 is done, and the action taker knows z . It takes just one more period for the action taker to learn H . So delay is indeed four. The key issue is whether the functions G, H exist.

(3) Is there a two-person F^* -realizing mechanism in which neither person has a burden more than 3 and one person has a burden less than 3?

Additional questions about F^* can be posed. In trying to answer them, it is clear that we have to restrict the F^* -realizing mechanisms that we are permitting. Once again the issue of *smoothness* is inescapable. If we permitted the sort of smuggling of many numbers into one number that we have already discussed (in 2.2.8), then no message need ever contain more than one number, and so we would obtain trivial and uninteresting answers to our questions. A workable definition of a smooth mechanism is that its sending functions ρ_{ij} be differentiable on their domains. Weaker and stronger definitions can be explored as well.

What is known about efficient speak-once-only mechanisms and what might one hope to learn? To organize the remarks we now make, we shall use the above illustration, and its accompanying three questions. For each question, we consider the tools available to answer it, as well as the answer itself and its possible generalization.

2.7.5 General remarks suggested by Question (1) For a goal function F that is sensitive to all k external variables, it is obviously true that a one person F -realizing mechanism has a tree structure, the sole person's burden is k , and the delay is also k . A non-tree F -realizing mechanism may improve on the one-person mechanism with regard to delay if the observing of the external variables is split among several persons who do their observing simultaneously, provided that the reduction in observing time exceeds the message-reading time.

2.7.6 General remarks suggested by Question (2) Let us relabel the external variables as $x_1, \dots, x_m, y_1, \dots, y_n$. Given a real-valued differentiable goal function $F(x_1, \dots, x_m, y_1, \dots, y_n)$, the general aggregation question is as follows:

Is there some neighborhood on which we can *aggregate the m variables x_1, \dots, x_m into $r < m$ variables*, i.e., do there exist a neighborhood $U \subseteq \mathbb{R}^{m+n}$, and $r < m$ real-valued functions G, H_1, \dots, H_r which are differentiable on U , such that on U we have

$$(+) \quad F(x_1, \dots, x_m, y_1, \dots, y_n) = G(H_1(x_1, \dots, x_m), \dots, H_r(x_1, \dots, x_m), y_1, \dots, y_n)?$$

An important contribution to answering this question is a theorem of Abelson (1980). Let F_i denote the partial derivative of F with respect to x_i . Abelson's theorem is as follows:

There exists a neighborhood $U \subseteq \mathbb{R}^{m+n}$ and differentiable functions G, H_1, \dots, H_r (where $r < m$) which satisfy (+) on U if and only if at most r of the functions F_i are linearly independent on U .

Checking the linear independence of the F_i , is not, in general, a straightforward matter. But a technique closely related to the Abelson theorem provides answers to the aggregation question in certain cases. To introduce the technique, let us use the following notation:

- For an n -tuple of nonnegative integers $\alpha = (\alpha_1, \dots, \alpha_n)$, and for $y = (y_1, \dots, y_n) \in \mathbb{R}^n$, let the symbol y^α denote the array of symbols $y_1^{\alpha_1} \dots y_n^{\alpha_n}$. (This array will be used to identify partial derivatives of varying orders). The symbol $|\alpha|$ denotes the sum $\alpha_1 + \dots + \alpha_n$.
- The symbol $D(F, x_t, y^\alpha)$ denotes the following partial derivative of order $1 + |\alpha|$:

$$\frac{\partial^{1+|\alpha|} F}{\partial x_t \partial y_1^{\alpha_1} \dots \partial y_n^{\alpha_n}}.$$

- The symbol $D(G, z)$ denotes the partial derivative $\frac{\partial G}{\partial z}$.

Now let us associate with the goal function F the *Hessian*

$$\mathcal{H}(F) \equiv \begin{pmatrix} D(F, x_1, y_1) & \dots & D(F, x_1, y_n) \\ \vdots & \ddots & \vdots \\ D(F, x_m, y_1) & \dots & D(F, x_m, y_n) \end{pmatrix},$$

and the *bordered Hessian*

$$\mathcal{BH}(F) \equiv \left(\begin{array}{c|c} D(F, x_1) & \\ \vdots & \mathcal{H}(F) \\ D(F, x_m) & \end{array} \right).$$

For easy reference, note that

The m rows of \mathcal{BH} may be indexed by the m variables x_1, \dots, x_m that are being aggregated. The first of the $n+1$ columns may be indexed by the function F and the remaining columns by the non-aggregated variables y_1, \dots, y_n .

We now state two propositions about aggregation.³⁷ We call them Proposition (*) and (**).

Proposition (*) provides bordered-Hessian conditions that are necessary for the existence of functions G, H_1, \dots, H_r satisfying (+).

Proposition³⁸ (*)

³⁷They are found in Mount and Reiter, 1996 and in Appendix B of the book by Mount and Reiter.

³⁸The proof of Proposition (*) is as follows:

If (+) holds on U , then everywhere on U we have

$$D(F, x_i) = \sum_{k=1}^r D(G, H_k) D(H_k, x_i)$$

Let F be a \mathcal{C}^3 function. If there exist \mathcal{C}^2 functions G, H_1, \dots, H_r such that (+) holds on some neighborhood U , then

$$\text{rank } \mathcal{BH}(F) \leq r \text{ at every point of } U.$$

Now we turn to Proposition (**), which provides conditions that are sufficient for the existence of functions G, H_1, \dots, H_r satisfying (+). In stating Proposition (**), we consider open subsets of \mathbb{R}^{m+n} of the form $W \times V$ where $W \subseteq \mathbb{R}^m, V \subseteq \mathbb{R}^n$. We denote points in $W \times V$ by (p, q) , where $p \in W, q \in V$.

Proposition ()**

Let F be a \mathcal{C}^{k+1} function, $k \geq 2$. Suppose that $\text{rank } \mathcal{BH}(F) \leq r$ everywhere on $W \times V \subseteq \mathbb{R}^{m+n}$ and $\text{rank } \mathcal{H}(F) = r$ everywhere on W . Then there exists a neighborhood $U \subseteq W \times V$ and \mathcal{C}^k functions G, H_1, \dots, H_r for which (+) holds.

We can successfully apply Proposition (*) to the goal function $F^* = w + x + y + z + wxyz + (\frac{1}{w} + \frac{1}{x} + \frac{1}{y} + \frac{1}{z})$ considered in Question 2. Any three variables, say w, x, y , can be aggregated into two real variables, using the \mathcal{C}^2 functions $H_1 = x + \frac{1}{x} + y + \frac{1}{y} + z + \frac{1}{z}, H_2 = xyz, G = H_1 + w + \frac{1}{w} + wH_2$. Can they be aggregated into just one real variable, i.e., do there exist differentiable functions G, H such that on every neighborhood in \mathbb{R}^+ we have $F = G(H(w, x, y), z)$? To check this, the relevant \mathcal{BH} is

$$\begin{pmatrix} 1 - 1/w^2 + xyz & xy \\ 1 - 1/x^2 + wyz & wy \\ 1 - 1/y^2 + wxz & wx \end{pmatrix}.$$

That has rank two on any neighborhood in which (for example) all four variables are neither zero nor one, and no two variables take the same value. Proposition (*) tells us that if the rank exceeds one on a neighborhood, then on that neighborhood we cannot aggregate three variables into one. There is no

and

$$D(F, x_i y_j) = \sum_{k=1}^r D(G, H_k y_j) D(H_k, x_i),$$

since the H_k are independent of the variables y_1, \dots, y_n . We can therefore write $\mathcal{BH}(F)$ as

$$\begin{pmatrix} \sum_{k=1}^r D(G, H_k) D(H_k, x_1) & \sum_{k=1}^r D(G, H_k, y_1) D(H_k, x_1) & \dots & \sum_{k=1}^r D(G, H_k, y_n) D(H_k, x_1) \\ \vdots & \vdots & \ddots & \vdots \\ \sum_{k=1}^r D(G, H_k) D(H_k, x_m) & \sum_{k=1}^r D(G, H_k, y_1) D(H_k, x_m) & \dots & \sum_{k=1}^r D(G, H_k, y_n) D(H_k, x_m) \end{pmatrix}.$$

So everywhere on U , each column of $\mathcal{BH}(F)$ is a linear combination of the r column vectors

$$\begin{pmatrix} D(H_1, x_1) \\ \vdots \\ D(H_1, x_m) \end{pmatrix}, \dots, \begin{pmatrix} D(H_r, x_1) \\ \vdots \\ D(H_r, x_m) \end{pmatrix}.$$

That implies that the rank of $\mathcal{BH}(F)$ is at most r at each point of U . ■

two-person F^* -realizing mechanism with differentiable sending functions in which neither person has a burden more than two and delay is not more than four.

On the other hand, consider the following goal function F^{**} , identical to F^* except that we now raise the term $\frac{1}{w} + \frac{1}{x} + \frac{1}{y} + \frac{1}{z}$ to the power 2. Thus

$$F^{**}(w, x, y, z) = (w + x + y + z) + (wxyz) + \left(\frac{1}{w} + \frac{1}{x} + \frac{1}{y} + \frac{1}{z} \right)^2.$$

Suppose we now ask whether we can aggregate three variables into two (as we could for F^*), i.e., do there exist differentiable functions G, H_1, H_2 such that on every neighborhood in \mathbb{R}^{4+} we have $F^{**} = G(H_1(w, x, y), H_2(w, x, y), z)$? Again \mathcal{BH} has three rows (indexed by the three variables being aggregated) and two columns (indexed by F and by the remaining variable). Proposition (*) says that if the proposed aggregation is possible then \mathcal{BH} has rank at most 2. But that is the case whether or not the proposed aggregation is possible. So Proposition (*) cannot be used to rule out the proposed aggregation. Proposition (**) says that the proposed aggregation is indeed possible if the rank of \mathcal{BH} is at most 2 and at some point the rank of \mathcal{H} is exactly 2. The latter condition cannot be satisfied, since \mathcal{H} has just one column. So Proposition (**) does not tell us that the proposed aggregation is possible.

More generally:³⁹

The necessary condition of Proposition () has no force (it is automatically satisfied) if*

$$m < n + 1$$

*(and hence $r < n + 1$). The sufficient condition of Proposition (**) cannot be satisfied if*

$$n < r.$$

We can, however, generalize Proposition (*). The following generalization is more widely applicable than Proposition (*).

Proposition (*')

Let F be a \mathcal{C}^{k+1} function, $k \geq 2$. Fix an integer $\ell \geq 0$. Using our definition of expressions of the form $D(F, x_j, y^\alpha)$, consider the matrix

$$M = \begin{pmatrix} D(F, x_1, y^{\alpha(1)}) & \dots & D(F, x_1, y^{\alpha(\ell)}) \\ \vdots & \ddots & \vdots \\ D(F, x_m, y^{\alpha(1)}) & \dots & D(F, x_m, y^{\alpha(\ell)}) \end{pmatrix},$$

³⁹The limitations of Propositions (*), (**) (and the extended Proposition (*') which follows) might lead one to explore algebraic approaches. In particular, one might seek counterparts of the Abelson Theorem in which linear independence is replaced by algebraic independence. T functions are algebraically dependent if there is a polynomial, with the T functions as its arguments, which takes the value zero at all points in the functions' domains. The T functions are algebraically independent if there is no such polynomial. Such a counterpart of the Abelson Theorem might hold for rational goal functions F , i.e., F is the quotient of two polynomials.

where, for all $t \in \{1, \dots, \ell\}$, the symbol $\alpha^{(t)}$ denotes a t -tuple of nonnegative integers with $|\alpha^{(t)}| < k$. If there exist \mathcal{C}^k functions G, H_1, \dots, H_r such that on some neighborhood

$$F = G(H_1(x_1, \dots, x_m), \dots, H_r(x_1, \dots, x_m), y_1, \dots, y_n),$$

then for all ℓ the matrix

$$M^* \equiv \left(\begin{array}{c|c} D(F, x_1) & \\ \vdots & \\ D(F, x_m) & M \end{array} \right)$$

has rank at most r on that neighborhood.⁴⁰

While Propositions (*) and (**) did not allow us to resolve the aggregation question for the function $F^{**} = (w + x + y + z) + (wxyz) + \left(\frac{1}{w} + \frac{1}{x} + \frac{1}{y} + \frac{1}{z}\right)^2$, we can now do so, using Proposition (*'). To fit the notation of Proposition (*'), we first relabel the four variables as x_1, x_2, x_3, y . Our function is :

$$F^{**} = (x_1 + x_2 + x_3 + y) + (x_1 x_2 x_3 y) + \left(\frac{1}{x_1} + \frac{1}{x_2} + \frac{1}{x_3} + \frac{1}{y}\right)^2.$$

We ask whether there exist real-valued \mathcal{C}^2 functions G, H_1, H_2 such that $F^{**} = G(H_1(x_1, x_2, x_3), H_2(x_1, x_2, x_3), y)$ on all neighborhoods in \mathbb{R}^{4+} .

To apply Proposition (*'), we first have to choose the partial derivatives that will make up the matrix M^* . It is clear that if we are going to rule out the existence of a given number of aggregators, then we want to make M (and hence M^*) as large as possible without repeating columns or introducing redundant columns like those consisting entirely of zeros. It turns out that the following matrix M^* suffices to resolve the aggregation question for our function F^{**} :

$$M^* = \left(\begin{array}{ccc} D(F, x_1) & D(F, x_1, y_1) & D(F, x_1, y_1^2) \\ D(F, x_2) & D(F, x_2, y_1) & D(F, x_2, y_1^2) \\ D(F, x_3) & D(F, x_3, y_1) & D(F, x_3, y_1^2) \end{array} \right).$$

If we now choose a suitable neighborhood U and perform a Gauss-Jordan reduction on M^* for any point in U , we obtain the identity matrix. So M^* has rank 3 and Proposition (*') tells us that the proposed aggregation cannot occur on U . The neighborhood U must be chosen so that we do not divide by zero during the reduction process. To do so, it suffices to pick a U that does not intersect any of the zero sets of the numerators and denominators of any of the entries of M^* at any stage of the Gauss-Jordan process. Since those numerators and denominators are always polynomials in x_1, x_2, x_3, y_1 , it is possible to find such a U .⁴¹

⁴⁰The proof is essentially the same as the proof of Proposition (*). The existence of functions G, H_1, \dots, H_r on some neighborhood U implies that each column of M^* is a linear combination of the column vectors

$$\left(\begin{array}{c} D(H_1, x_1) \\ \vdots \\ D(H_1, x_m) \end{array} \right), \dots, \left(\begin{array}{c} D(H_r, x_1) \\ \vdots \\ D(H_r, x_m) \end{array} \right).$$

That implies that the rank of M^* is at most r on U . ■

⁴¹Precisely stated, our result is as follows:

While Propositions (*), (**), (*)' are useful for a variety of goal functions, a general difficulty is the absence of a *usable necessary AND sufficient condition* for a proposed aggregation. The Abelson condition is both necessary and sufficient, but for many goal functions it is very difficult to verify.

2.7.7 A general puzzle suggested by Question (3): Can Question (3), and similar questions, be answered using the Abelson condition?

Consider again the two-person mechanism in Figure 3, which realizes the goal function $F^* = \left(w + \frac{1}{w} + x + \frac{1}{x} + y + \frac{1}{y} + z + \frac{1}{z}\right) + (wxyz)$. For a general discussion of two-person F^* -realizing mechanisms it will be useful to start by relabeling the two persons. Let us call the action taker A and the other person B. In the Figure-3 mechanism, Person B observes three environment variables and has a burden of 3, while Person A, the action-taker, also has a burden of 3. (He observes one environment variable and he receives two numbers from B). Turning to arbitrary two-person mechanisms, including mechanisms wherein some environment variable is observed by more than one person, we shall argue that

(+) F^* cannot be realized by a two-person mechanism in which the maximum burden is 3 and one person has a burden less than 3.

Suppose a two-person mechanism realizes F^* and has maximum burden less than or equal to 3. Since the values of all four variables are needed in order to compute F^* , we can claim:

(++) the mechanism's graph must have a path from each environment variable to A.

Now suppose that *B observes exactly two variables*, say x and y . Then, by (++), A must observe the other two, z and w . A cannot observe both x and y (in addition to z and w) since the maximum burden is less than 3. A cannot observe just one of x, y (in addition to z, w), for if he did, then his burden would exceed 3, since (by ++) he must also receive at least one number from B. If A's burden is less than 3, then B must send just one number. So F may be written $F(x, y, z, w) = A(B(x, y), z, w)$, where A and B (in a slight abuse of notation) denote scalar-valued functions computed by persons A and B. Consulting Abelson's theorem (for the case $r = 1$), we see that for any fixed values of x, y , say x_0, y_0 ,

the functions $F_x(x_0, y_0, z, w) = 1 - \frac{1}{x_0^2} + y_0 w z$, $F_y(x_0, y_0, z, w) = 1 - \frac{1}{y_0^2} + x_0 w z$ must be linearly dependent,

where the symbol F_j again denotes the partial derivative of F with respect to the variable j . But that is the case if and only if the vectors $(1 - \frac{1}{x_0^2}, y_0)$, $(1 - \frac{1}{y_0^2}, x_0)$ are linearly dependent in a neighborhood in \mathbb{R}^2 , which is, in turn, the case if and only if $x_0 - \frac{1}{x_0} - y_0 + \frac{1}{y_0} = 0$ in a neighborhood in \mathbb{R}^2 . That, however, is not true. We conclude that F cannot be realized by a two-person mechanism with a minimal burden of 3 in which B observes exactly two variables.

There exists a finite number of polynomials f_1, \dots, f_k in the variables x_1, x_2, x_3, y_1 such that on any open subset U of \mathbb{R}^4 not intersecting the zero sets of f_1, \dots, f_k , there do not exist C^2 functions G, H_1, H_2 such that

$$F^{**}(x_1, x_2, x_3, y) = G(H_1(x_1, x_2, x_3), H_2(x_1, x_2, x_3), y) \text{ at every point of } U.$$

On the other hand, we can easily rule out the possibility that B observes exactly one variable. If he does, then (by $(++)$) A must observe the other three. But then A has reached the maximum burden of 3 and is not able to receive any variable from B .

So B must observe exactly three variables, say x, y, z . Then the only way that $(+)$ could be violated (i.e., the only way to improve on the mechanism of Figure 3) would be to give A a burden less than three. By $(++)$, A must observe w and since A needs to know more than w in order to compute F , A must have a burden of at least two. Since A cannot compute F knowing only two environment variables, he must receive one or more numbers from B . To keep A 's burden below 3, he must receive only a single number from B . If that were so, we could write

$$F(x, y, z, w) = A(B(x, y, z), w),$$

where both of the functions A, B are scalar-valued. But this is again ruled out by the Abelson criterion (for $r = 1$), since the three functions

$$F_x(x_0, y_0, z_0, w) = 1 - \frac{1}{x_0^2} + y_0 z_0 w, F_y(x_0, y_0, z_0, w) = 1 - \frac{1}{y_0^2} + x_0 z_0 w, F_z(x_0, y_0, z_0, w) = 1 - \frac{1}{z_0^2} + x_0 y_0 w$$

are linearly independent on every neighborhood in \mathbb{R} . So we have established $(+)$.

Looking carefully at the preceding argument, one sees that it generalizes so as to yield the following Proposition. The statement “ s of the n variables in F variables may be smoothly aggregated into $r < s$ ” means that F may be written as a continuously differentiable function whose arguments are r continuously differentiable scalar-valued functions of those s variables, together with the remaining $n - s$ variables.

Proposition D Let $F : \mathbb{R}^4 \rightarrow \mathbb{R}$ be continuously differentiable and sensitive to all its variables⁴² Suppose that

- (I) three of F 's variables can be smoothly aggregated into two
- (II) no two of F 's variables may be smoothly aggregated into one
- (II) no three of F 's variables may be smoothly aggregated into one.

Then:

- (IV) there is a two-person F -realizing mechanism which has the graph shown in Figure 3, is smooth (its sending functions are continuously differentiable), and has burden 3 for both persons
- (V) there is no smooth two-person F -realizing mechanism in which each person has burden at most 3 and one of them has burden less than 3.

⁴²I.e., there is a neighborhood on which all partial derivatives are nonzero.

2.7.8 Goal functions for which Perfect Aggregation is possible. It is now clear that one faces formidable difficulties in characterizing the efficient speak-once-only mechanisms that realize a given goal function. Is there some general class of goal functions for which the task becomes much easier? One such class consists of the functions for which Perfect Aggregation (PA) is possible. A goal function $F : E_1 \times E_k \rightarrow \mathbb{R}$ has the PA property if it behaves like a sum: for *any* partitioning of the external variables, the variables in each set of the partitioning can be replaced by a *single* scalar-valued aggregator. If F is a sum, then the aggregator for a given set is the sum of the set's variables. As before, we confine attention to the case where each environmental variable is a scalar, i.e., each E_j is a subset of the real line.

Formally:

The function $F : E_1 \times \dots \times E_k \rightarrow \mathbb{R}$ has the PA property if for ANY partitioning of $\{1, \dots, k\}$, say into sets $T_1, \dots, T_s, s \leq k$, there exist differentiable real-valued functions G, A_1, \dots, A_s such that for any e in $E = E_1 \times \dots \times E_k$ we have

$$F(e) = G(A_1(e_{T_1}), \dots, A_s(e_{T_s})).$$

Note that each of the functions A_1, \dots, A_s must itself have the PA property. Examples of a function with the PA property are:

- $e_1 + \dots + e_k$
- $e_1 \cdot e_2 \cdot \dots \cdot e_{k-1} \cdot e_k$
- $b^{e_1 + \dots + e_k}$.

If our goal function F has the PA property, we have a strong result about the structure of efficient F -realizing mechanisms. It says that we can confine attention to trees if we seek mechanisms that are efficient with respect to number of persons and individual burdens.

Proposition E

Suppose F is real-valued and has the PA property. Suppose that $\Lambda = \langle G, (\vec{\rho}_{k+1}, \dots, \vec{\rho}_n) \rangle$ is an F -realizing mechanism and that the graph G , with received-from sets R_{k+1}, \dots, R_n , is not a tree. Then there exists an F -realizing mechanism $\Lambda' = \langle G', (\vec{\rho}'_{k+1}, \dots, \vec{\rho}'_{n'}) \rangle$, with received-from sets $R'_{k+1}, \dots, R'_{n'}$ such that

- (a) G' is a tree
- (b) $n' \leq n$
- (c) for every $i \in \{k+1, \dots, n'\}$, we have

$$\#R'_i \leq \#R_i,$$

where the symbol $\#$ means “number of elements in”.⁴³

The Leontieff Theorem It is remarkable that many years ago Leontieff (1952), inspired by the search for useful production functions in economics, and unaware of any implications for the theory of organization, proved a theorem that anticipated the Abelson theorem that we considered above. (Abelson’s theorem can be viewed as a generalization of Leontieff’s). The theorem says that any function with the PA property is essentially a sum.⁴⁴

Proposition F (Leontieff)

Let $F : E_1 \times \cdots \times E_k \rightarrow \mathbb{R}$ be a differentiable function with the PA property. Then there exist differentiable functions $\phi, \phi_1, \dots, \phi_k$ such that for all (e_1, \dots, e_k) in $E_1 \times \cdots \times E_k$,

$$F(e) = \phi(\phi_1(e_1) + \cdots + \phi_k(e_k)).$$

So the PA property is not only a very strong one (yielding Proposition E) but it is also more restrictive than one might have supposed. One can develop further propositions about trees that realize PA goal functions, by imposing further conditions on the costs associated with the individual burdens. For example one might attach different costs to the “observation” burden (the number of variables that a node $i > k$ receives from nodes in $\{1, \dots, k\}$), than the other burdens. Then one simple proposition says that if observation cost is linear in the number of observed variables, while the cost associated with the receipt of non-environment variables is increasing in the number of those variables, then the only efficient goal-realizing tree — and hence (by Proposition E) the only efficient goal-realizing design — has just one person, who observes all k environment variables.

Replacing the Perfect Aggregation property by a more general additive-structure property

Once we drop the strong PA requirement, it is natural to begin by studying goal functions that lack the PA property but have an additive structure somewhat more general than the “generalized sum” $\phi(\phi_1(e_1) + \cdots + \phi_k(e_k))$ of the Leontieff Theorem. Let us relabel the environment components e_1, \dots, e_k so that they become the m real variables x_1, \dots, x_m . For $1 \leq r < m$, we shall say that $F : \mathbb{R}^m \rightarrow \mathbb{R}$ is r – additive on the open set $S = S_1 \times \cdots \times S_m$ in \mathbb{R}^m , if there exist C^∞ functions $H, H_1, \dots, H_r, V_1, \dots, V_r$ such that for all $(x_1, \dots, x_m) \in S$, we have

$$(1) \quad F(x_1, \dots, x_m) = H\left(H_1(V_1(x_1) + \cdots + V_1(x_m)) + \cdots + H_r(V_r(x_1) + \cdots + V_r(x_m))\right).$$

⁴³The proof has two steps. Step I Since F has the PA property, it is clear that given any n^* -node directed tree G^* which meets the requirements of the graph of a speak-once-only mechanism on the k environment variables, one can construct an F -realizing mechanism such that (i) its graph is the tree G^* , and (ii) exactly one real number is sent over each arc. (In that mechanism some functions ρ_{ij} may be the identity function).

Step 2 Without loss of generality we may suppose that on any arc of G a single real number is sent. (If more than one number is sent, then they can be replaced by an aggregator). Now convert the given graph G into a tree as follows: Whenever there is more than one directed path between a node $i \neq n$ and the node n , delete all but one of those paths. The result will be a tree, say G' , with fewer than n nodes. The tree retains the original action-taking node. By Step 1, we can construct a new F -realizing mechanism with the tree G' as its graph and with exactly one real number sent over each arc. So the new mechanism satisfies conditions (a), (b) and (c). ■

⁴⁴Leontieff did not use our terminology, but his theorem can be restated in the form that follows.

We shall say that F is *weakly r -additive on S* if

$$(2) \quad F = H\left(H_1(V_1(x_1) + \cdots + V_1(x_m)), \dots, H_r(V_r(x_1) + \cdots + V_r(x_m))\right).$$

We shall say that F is *minimally r -additive (minimally weakly r -additive) on S* if for any positive integer $r' < r$, F is r -additive on S but not r' -additive (weakly r -additive on S but not weakly r' -additive). An economic interpretation might be as follows. The function F gives the organizational action appropriate to the environment (x_1, \dots, x_m) . That action depends on the *aspects* of the environment given by the V_i . There are $r < m$ aspects. The i th aspect influences the appropriate action through the function H_i . We may suppose that the goal function F has been parsimoniously written, so that minimality is assured.⁴⁵

AN EXAMPLE: Suppose there are four real environment variables: w, x, y, z and consider once again our function

$$(3) \quad F^{**}(w, x, y, z) = (w + x + y + z) + (wxyz) + \left(\frac{1}{w} + \frac{1}{x} + \frac{1}{y} + \frac{1}{z}\right)^2.$$

This function is 3-additive on any open set in which all four variables are nonzero, since it can be written

$$F^{**} = (w + x + y + z) + \exp(\ln w + \ln x + \ln y + \ln z) + \left(\frac{1}{w} + \frac{1}{x} + \frac{1}{y} + \frac{1}{z}\right)^2.$$

The results we obtained above imply that F^{**} is *minimally 3-additive*.

By contrast consider our function F^* which we obtain by deleting the exponent 2 in the third term of (3). The function F^* is 3-additive but it is not minimally 3-additive, since it can also be written in the 2-additive form

$$\left(w + \frac{1}{w} + x + \frac{1}{x} + y + \frac{1}{y} + z + \frac{1}{z}\right) + (wxyz).$$

2.7.8 Conjectures about the aggregation properties of goal functions with an additive structure

What can be said about efficient speak-once-only mechanisms which realize a goal function that has an additive structure? If certain key conjectures were established, we could develop a substantial theory about such mechanisms. First let us define $F : \mathbb{R}^m \rightarrow \mathbb{R}$ to be *t -aggregative on*

⁴⁵A possible scenario: The organization operates a production facility in each of m locations. Each location always produces a fixed proportion of the total quantity produced by the organization. That total quantity is the action to be chosen. The appropriate total quantity depends on total production cost, which depends, in turn, on the amounts of each of r centrally procured inputs required for each unit of product in each location. For every location, say location i , the r input requirements are determined by the environmental parameter x_i , which describes the current technology in location i . The r input requirements are $V_1(x_i), \dots, V_r(x_i)$.

Once the case of a scalar-valued additive goal function F is understood, one could turn to the case of a vector-valued additive goal function F . That would, of course, permit more flexible scenarios, including scenarios wherein the organization's action is a vector, specifying (for example) an output level for each of the m locations.

the open set $S \subseteq \mathbb{R}^m$, if for any proper subset W of $\{1, \dots, m\}$, there exist \mathcal{C}^2 real-valued functions G, A_1, \dots, A_t , with $t < \#W$, such that for all $x = (x_W, x_{-W})$ in S we have⁴⁶

$$F(x_1, \dots, x_m) = G(A_1(x_W), \dots, A_t(x_W), x_{-W}).$$

Thus, on the set S , it is enough to know x_{-W} and the values of the t aggregators A_1, \dots, A_t , in order to compute F . However we partition the m variables into the sets W and $-W$, we need no more than t aggregators of the variables in W in order to compute F .

The following key conjectures remain open.

Conjecture 1

If F is symmetric and r -additive on S and t -aggregative on S , with $t < r < m$, then F is also t -additive on S .

Conjecture 2 (stronger than Conjecture 1)

If F is symmetric and t -aggregative on S , with $t < m - 1$, then F is t -additive on S .

Weak versions of these conjectures:

Put “weakly” in front of “ r -additive” and “ t -additive”

Suppose F is symmetric, and *minimally* r -additive with $r < n - 1$. Suppose Conjecture 1 holds. Then we can claim that if we want to aggregate $x = (x_W, x_{-W})$ in S , with $r < \#W$, into as few variables as possible, while retaining our ability to compute F , we cannot do better than to use

$$(4) \quad G(A_1, \dots, A_r, x_{-W}) = H\left(H_1\left(A_1 + \sum_{\ell \in -W} V_1(x_\ell)\right) + \dots + H_r\left(A_r + \sum_{\ell \in -W} V_r(x_\ell)\right)\right),$$

where

$$(5) \quad A_i \equiv \sum_{\ell \in W} V_i(x_\ell), i = 1, \dots, r.$$

If it were possible to aggregate the variables x_W into fewer than r numbers — i.e., if there existed G, A_1, \dots, A_q , with $q < r$, such that $F = G(A_1(x_W), \dots, A_q(x_W), x_{-W})$ — then Conjecture 1 says that F is q -additive. But that contradicts the assumed minimal r -additivity of F . Note that the *converse* of Conjecture 2 holds. If F is t -additive, then we can write F in a t -aggregative form, using (4) and (5) (with $r = t$). Note also that *Conjecture 2 is correct for the case $m = 3$* . To see this, call the three environment variables x, y, z . Since $m - 1 = 2$, Conjecture 2 takes the following form:

⁴⁶Here $-W$ denotes the complement of W .

If F is symmetric and there exist \mathcal{C}^2 functions G, A such that

$$F(x, y, z) = G(A(x, y), z)$$

then F is 1-additive.

But that is just the simplest (three-variable) version of the Leontieff Theorem.

A modest proposition about mechanisms with a tree structure . Consider again our function

$$F^* = \left(w + \frac{1}{w} + x + \frac{1}{x} + y + \frac{1}{y} + z + \frac{1}{z} \right) + (wxyz).$$

As we saw in Figure 3, F^* can be realized by a two-person mechanism whose graph is a six-node tree. Each person's burden is 3. There is, however, no F^* - realizing mechanism whose graph is a tree with all the individual burdens being two or less. Our discussion of the application of Proposition (*) to F^* showed that while F^* is 2-aggregative it is not 1-aggregative. (If we choose the set W to contain three variables, then it is not the case that a single aggregator provides the information about those variables that is needed to compute F^*). Generalizing from the example of F^* and the Figure-3 tree, we can obtain a rather modest proposition:

Proposition G

Suppose the function

$$F : \mathbb{R}^k \rightarrow \mathbb{R},$$

where $k > 2$, is 2-aggregative but not 1-aggregative. Then in any two-person tree which realizes F , each person's burden is at least 2.

RESEARCH CHALLENGE #6: Find interesting classes of goal functions F for which Conjectures 1 and 2 hold.

2.7.10 Computing a speak-once-only mechanism's delay A mechanism's delay is the total elapsed time until the action taker has completed his computation of the organization's final action. We have considered delay informally, and we have illustrated delay, but we have not yet discussed how one might compute it. To do so we first define delay more carefully. We start by making the following simplifying assumption.

- (α) The environment variables change their values every G time units. For all the speak-once-only mechanisms we consider, G is sufficiently large that the mechanism's delay is less than G .

We repeat, in addition, three assumptions already made in our introductory discussion of delay in Section 2.7.1.

- (β) It takes one time unit for a person to receive or observe one (scalar-valued) variable.

- (γ) A person receives all the message variables that the mechanism specifies, and does all the observing required of him, before he sends messages or computes the final action.
- (δ) It takes zero time for a person to compute a message variable to be sent onward, and for the action taker to compute and put into force the final action.

Thus if a person observes X (scalar-valued) environment variables, receives Y (scalar-valued) message variables from other persons, and computes Z scalar-valued functions of the observed environment variables and the received message variables (to be sent forward to other persons), then $X + Y$ time units are required for him to complete these tasks. He is silent until the mechanism has operated for at least $X + Y$ time units.

Next, note that for a given speak-once-only mechanism, the sequence of tasks may be subject to choice. Suppose for example that in a given three-person mechanism Person 1 observes three environment variables, Person 2 observes three other environment variables, and each sends forward to Person 3 (the action taker) values of two scalar-valued functions. In one possible sequence, Person 1 first observes all his environment variables, then Person 2 observes all of his, and immediately after that (with no further time required), both persons compute and send forward the values of their scalar-valued functions. That would appear to be an inefficient sequence, since time would be saved if 1 and 2 did their observing simultaneously. Nevertheless it is a possible sequence. In general we may define a *protocol* π for a given mechanism as a specification of exactly what tasks (observing, sending, receiving, computing) each person performs at each time in a sequence of times (one time unit apart), terminating with the action taker's final-action computation. *A protocol is required to respect the above assumption γ .* (A formal definition of a protocol can be provided, but it requires an array of new symbols).

It will be useful to use the term *entity* for a node of the mechanism's directed graph, whether that node identifies an environment variable or a person. For the protocol π we may define a *time-to-completion function* f_π on the set of entities. For Person p , the number $f_\pi(p)$ is the shortest time after which Person p "does nothing", i.e., he does no further observing, receiving, sending, or computing. If the entity j is an environment variable (i.e., it is a node belonging to the set $\{1, \dots, k\}$), then we define $f_\pi(j)$ to be zero. If the protocol π is used, then the total elapsed time until the mechanism terminates is $f_\pi(n)$ where (as before) n is the action taker. Our definition of delay is then as follows:

Definition: *A speak-once-only mechanism's delay is the minimum value, over all protocols π , of $f_\pi(n)$. A protocol π for which $f_\pi(n)$ is minimal will be called a minimal protocol for the mechanism.*

A mechanism's delay may be computed recursively. To simplify notation, let T denote the function f_π , where π is minimal, and write T_p instead of $T(p)$. (Thus T_p is the smallest time until p does nothing further). We shall use the term "listens to" to cover both observation of an environment variable and receipt of a message variable. Suppose p listens to just one entity, say 1, from whom he receives N_1 variables. Then $T_p = T_1 + N_1$, since p has to wait T_1 units for 1 to finish and then takes N_1 units to listen to the N_1 variables sent to him by 1. At that point, p is done. Now suppose that p listens to more than two entities. We make the following claim:

Claim *There is a minimal protocol in which p listens to each of the persons from whom he receives in the order in which those persons are done with their sending to him.*

To argue this claim, start by supposing that p listens to two persons, say 1 and 2; 1 sends N_1 variables to P and 2 sends N_2 variables. Assume that $T_1 \leq T_2$. Then there is a minimal protocol in which

$$T_p = \max(T_1 + N_1, T_2) + N_2.$$

To see this, note first that p must wait T_1 time units for 1 to finish. He spends the next N_1 time units listening to 1. If T_2 has elapsed by then, (i.e., $T_2 \leq T_1 + N_1$), then p spends the next N_2 time units listening to 2. If T_2 has not yet elapsed ($T_2 > T_1 + N_1$), then p must wait for T_2 units to elapse and then spends the next N_2 units listening to 2. If p listens to three persons, say 1,2,3, with $T_1 \leq T_2 \leq T_3$, then we obtain, for some minimal protocol:

$$T_p = \max\left(\max(T_1 + N_1, T_2) + N_2, T_3\right) + N_3.$$

Recalling that $T_p = 0$ if the node p corresponds to an environment variable, we now see that the following general procedure computes T_p for a person p in some minimal protocol.

- (1) Label the entities that p listens to as $1, 2, \dots, m$, where $T_1 \leq T_2 \leq \dots \leq T_m$.
- (2) Initialize values $v = T_1 + N_1$ and $k = 1$. If $k = m$, then set $T_p = v$ and stop.
- (3) Replace v with $v^* = \max(v, T_{k+1}) + N_{k+1}$ and replace k with $k + 1$.
- (4) If $k + 1 < m$, repeat step (3). If $k + 1 = m$, then set $T_p = v^*$ and stop.

Note that the recursive scheme just described is indeed realized by some protocol for a given mechanism, since the four steps can clearly be carried out for a person p who observes environment variables but does not receive messages. Since a speak-once-only mechanism has no directed cycles, it follows that the four steps can also be carried out for a person p who receives messages. Since we can, in particular, carry out the four steps for $p = n$ (the action taker), the four steps provide a recursive way of computing the delay; the delay is $T_n = f_\pi(n)$ where π is a minimal protocol. Roughly speaking, the steps insure that each person is “as busy as possible”. Each person begins the work of listening as soon as signals (environment variables or message variables) begin to arrive, and no person stays idle while there is work (listening) that he could be doing. (Assumption (γ) is crucial for this claim).

Some preliminary computer experiments with the recursive procedure suggest that typically the time required to compute T_n rises rapidly as n and k grow. It seems likely that one can construct improved versions which exploit the properties of narrow classes of goal functions.

RESEARCH CHALLENGE # 7: For interesting classes of goal functions, refine the delay-computation algorithm so that it performs rapidly for mechanisms that realize a function in the class.

Examples of the computation of delay

(a) In a one-person mechanism, where the action taker observes all environment variables, delay is just the number of those variables.

(b) Suppose there are $k = 2^M$ environment variables, where M is a positive integer not less than 2. Suppose the mechanism's graph is a binary tree with a person at each non-leaf node. Then delay is $2 \log_2 k = 2M$. Similarly if $k = t^M$, where t and M are integers greater than 2, and the mechanism's graph is a t -ary tree with a person at each non-leaf node, then delay is $t \log_t k = tM$. (Since each person listens to t variables, it takes t time units for all the persons at a given tier of the tree to finish, and all persons in the tier finish at the same time; there are M tiers, so delay is tM). For such a mechanism no protocol can achieve a delay less than tM . In contrast, a one-person mechanism has a delay of $k = 2^M$ as well as a burden of 2^M

(c) Consider once again the goal function

$$F^{**}(w, x, y, z) = (w + x + y + z) + (wxyz) + \left(\frac{1}{w} + \frac{1}{x} + \frac{1}{y} + \frac{1}{z} \right)^2.$$

Consider the following two non-tree mechanisms that realize F^{**} . For each of them delay can be checked visually or it can be found using our recursive four-step procedure.

FIGURES 4 AND 5 HERE

The six-person Figure-4 mechanism has a property of “U-forms” in the economic literature, briefly mentioned in the Introduction. Each environment variable is observed by just one person (a specialist). Persons 1 and 2 are the observers. They report to a “sum” specialist (Person 3), a “product” specialist (Person 4) and a “reciprocal” specialist (Person 5). Those three specialists report, in turn, to the action taker (Person 6), who then has all the information needed to compute F .

The five-person Figure-5 mechanism has a property of “M-forms” in the economic literature: each environment variable is separately observed by several persons. Persons 1, 2, and 3 all observe w, x, y . For those variables, Person 1 is a “sum” specialist, Person 2 is a “product” specialist, and Person 3 is a “reciprocal” specialist. Person 1 reports directly to Person 5 (the action taker), but Persons 2 and 3 report to an intermediary (Person 4), who also observes z directly, as does Person 5. We find that

- the Figure-4 mechanism uses 6 persons. Five of them have a burden of 2 and one of them (the action taker) has a burden of 3. Delay is 7.
- the Figure-5 mechanism uses 5 persons. Each has a burden of 3. Delay⁴⁷ is 6.

The comparison illustrates a general conjecture:

⁴⁷At the end of the first three periods, 1, 2, and 3 have each learned w, x, y , while 4 and 5 have learned z . No messages have yet been sent. In the fourth period, 5 receives $w + x + y$ from 1, while 4 receives wxy from 2. In the fifth period, 4 receives $\frac{1}{w} + \frac{1}{x} + \frac{1}{y}$ from 3. In the sixth period, 5 receives $wxyz + \left(\frac{1}{w} + \frac{1}{x} + \frac{1}{y} + \frac{1}{z} \right)^2$ from 4 and is now able to compute F .

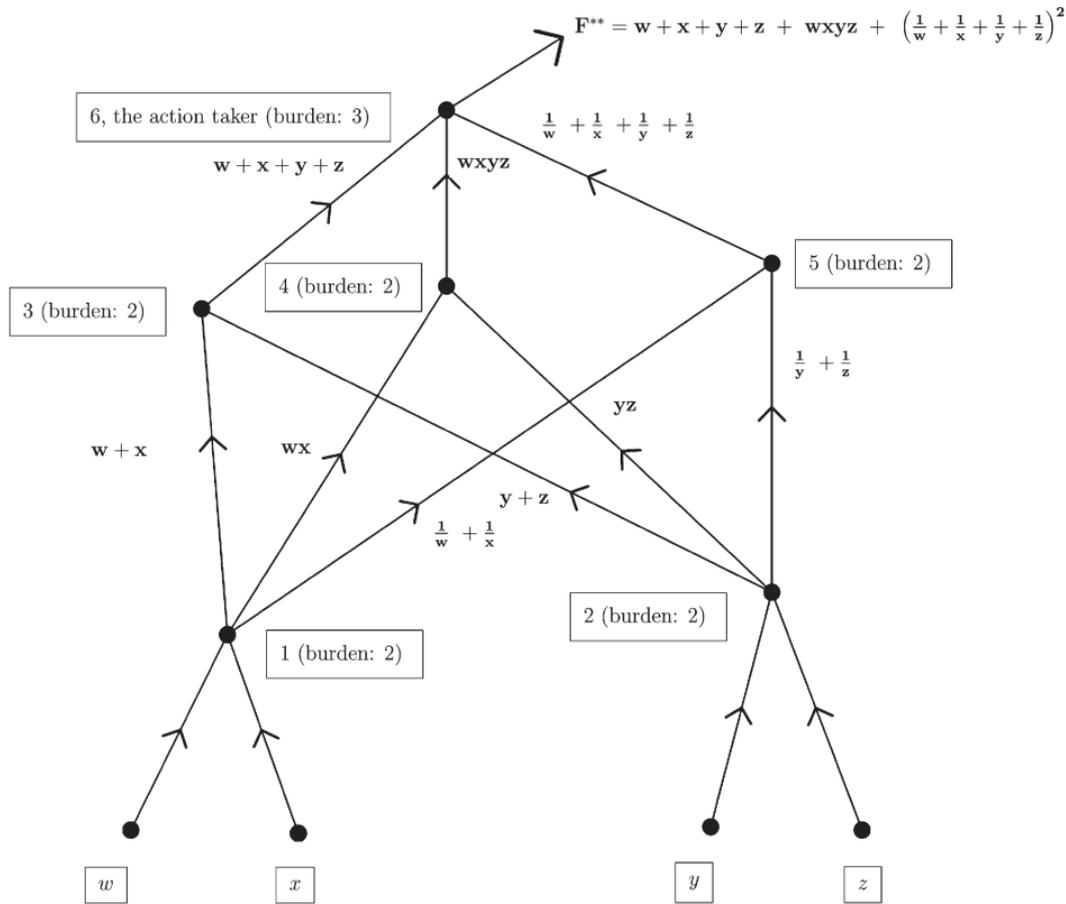


Fig. 4. This is a six-person nontree F^{**} realizing mechanism with delay 7. It has the “U-form” property: each environment variable is observed by a single person.

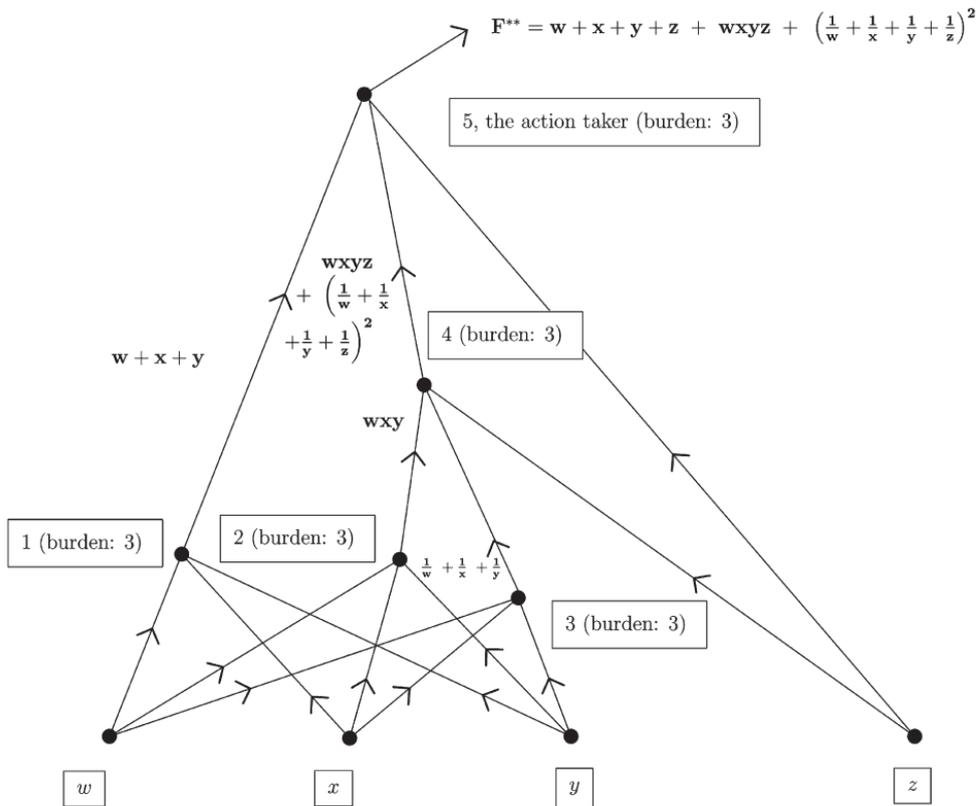


Fig. 5. This is a five-person nontree F^{**} realizing mechanism with delay 6. It has the “M-form” property: each environment variable is observed by more than one person.

Conjecture § Consider a goal function F which is a sum of three scalar-valued functions of four environment variables and can not be expressed as a function of fewer than three scalar-valued functions. There does not exist a non-tree F -realizing mechanism which dominates all the other non-tree F -realizing mechanisms with regard to number of persons, individual burdens, and delay.

RESEARCH CHALLENGE # 8: Apply the recursive delay-computation algorithm to establish Conjecture § and similar conjectures.

3. MODELS IN WHICH THE DESIGNER IS CONCERNED WITH INCENTIVES AS WELL AS INFORMATIONAL COSTS

A gigantic literature considers the designing of schemes that induce the self-interested members of an organization to make choices that meet a given organizational goal. One branch of this literature studies contracts. Another studies “implementation”; it considers schemes in which each self-interested member chooses an individual message, knowing that an organizational action will be chosen once all the messages are sent, knowing the outcome function that will be used to assign an action to the individual messages, and knowing the impact of each action on his own personal welfare. Such a scheme defines a message-choosing game. If the scheme is properly designed then it “implements” a goal correspondence that has been given to the designer: at an equilibrium of the game, the messages chosen are those for which the assigned action (prescribed by the outcome function) is goal-fulfilling.⁴⁸ Contracting schemes and goal-implementing mechanisms have informational costs, and it would be of great interest to be able to find the cheap ones. Informational costs are often discussed informally. Indeed some of the classic goal-implementing mechanisms (e.g., those that lead the organization’s self-interested members to make the correct choice as to the quantity of a public good) are striking in their informational simplicity.⁴⁹ But it is very rare for a paper to present a goal-implementing mechanism and then to argue its informational minimality in a precise way. The managerial accounting literature studies mechanisms or contracting schemes to be used by a profit-maximizing principal and the self-seeking “responsibility centers” (cost or profit centers) of a firm. The information transmitted by the centers (perhaps by using transfer prices) is far more modest than, for example, that transmitted in a Direct Revelation mechanism which achieves the same result. But again, informational minimality is not formally shown.⁵⁰ Occasionally, however, papers in this literature explicitly measure an informational cost, e.g., the cost of reducing the variance in the noise that accompanies a signal sent from the agents to the principal.⁵¹

⁴⁸Two excellent introductory surveys of the implementation literature are: R. Serrano, 2004 and M.O. Jackson, 2001.

⁴⁹Groves and Ledyard, 1977 is the classic paper on public-good provision in an organization whose members differ with regard to their private valuation of the good, using a mechanism that appears to be informationally cheap, although no formal claim about its informational minimality is made. By contrast, informational cost is formally studied in a “non-incentive” paper by Sato (1981). That paper finds a lower bound to the message space size for mechanisms used by an economy which seeks an efficient allocation of resources to the production of public goods, if we assume that the agents voluntarily follow the mechanism’s rules.

⁵⁰See, for example, three papers by N. Melumad, D. Mookherjee, and S. Reichelstein: (1992), (1995), and (1997).

⁵¹See, e.g., A. Ziv (2000).

We now turn to some research in which initial steps toward incorporating informational costs as well as incentives are taken.

3.1 The message-space size required for implementation of a goal Suppose, as before, that each person i has a privately observed local environment e_i whose possible values comprise the set E_i . In most implementation discussions, e_i determines i 's preferences over the possible organizational actions (and may contain other information about the organization's environment as well). Then an n -person implementation scheme, often called a *game form*, has two elements: an n -tuple $S = (S_1, \dots, S_n)$ of individual strategy spaces, and an outcome function $h : S \rightarrow A$, where A is the set of possible organizational actions. In particular, a strategy $s_i \in S_i$ may be a rule that tells i how to behave at each stage of a T -step message-announcement process, where the organization's action is a function of the step- T announcements. It tells him what message m_i , in a set M_i of possible messages, to announce at step t , given what he knows about the announcements thus far and given his current e_i . The organizational action is therefore a function of the current $e = (e_1, \dots, e_n)$. Given his current e_i , and given any strategy choices by the others, person i (who knows the function h) is able to rank any two of his own strategies with regard to his own welfare. For each e , we therefore have a game, denoted Γ_e , and we may choose a solution concept — say Nash equilibrium — to identify the game's equilibria. The scheme $\langle S, h \rangle$ implements a goal correspondence G from $E = E_1 \times \dots \times E_n$ to A if $h(s) \in G(e)$ whenever $s = (s_1, \dots, s_n)$ is a Nash equilibrium of the game Γ_e .

Now consider the equilibrium strategy profiles of the game Γ_e . As before, let m denote (m_1, \dots, m_n) and let M denote $M_1 \times \dots \times M_n$. Define the correspondence

$$\mu_i(e_i) = \{m \in M : \text{for some } e^* \in E, m \text{ is the step-}T \text{ announcement for } (s, e^*), \\ s \text{ is an equilibrium of } \Gamma_{e^*}, \text{ and } e_i = e_i^*\}.$$

Let \tilde{h} be an outcome function from M to A , with the following property: $\tilde{h}(m) = a$ if (1) $a = h(s)$ and (2) for some $e \in E$, s is an equilibrium strategy profile of the game Γ_e and m is the step- T message for (e, s) . The triple $\langle M, (\mu_1, \dots, \mu_n), \tilde{h} \rangle$ is a (*privacy-preserving*) *mechanism on E in our previous sense*. Moreover it realizes the goal correspondence G .

One can ask: among the G -implementing mechanisms (obtained from a G -implementing game form in the manner just described), which ones have a minimal message-space size? Or, less ambitiously, is there a useful lower bound to such a mechanism's message-space size? The message-space requirements for implementation of a goal correspondence are certainly going to be harsher, in general, than the requirements for realization alone. "How much harsher?" is a difficult question, and very few papers have addressed it.⁵²

3.2 Models in which the organization's mechanism is partly designed by its self-interested members, who bear some of the informational costs.

3.2.1 A model in which the organization's decentralized self-interested members choose their search efforts, a "Decentralization Penalty" results, and the effect of improved

⁵²Three of them are: S. Reichelstein (1984); S. Reichelstein and S. Reiter (1980), G. Tian (1990).

search technology on that penalty can be traced Consider a three-person organization⁵³ consisting of two Managers, called 1 and 2, and a Headquarters (HQ). There is a changing external environment $e = (e_1, e_2)$. Manager $i, i = 1, 2$, observes e_i , whose set of possible values is E_i . HQ does no observing. Whenever there is a new value of $e = (e_1, e_2)$, HQ learns something about it from the managers, though their knowledge of the new e is imperfect. Having received reports about the new e from the managers, HQ chooses an *action* a . The organization then collects a payoff $W(a, e)$. [Example: The organization makes a product which it markets in two locations. The variable e_i is the product's price in location i next week. The action a has two components a_1 and a_2 , where a_i is the amount to be shipped to location i for sale next week. The payoff $W(a, e)$ is the profit earned next week.]

Manager i always learns something about a new e_i in the same way. He has a finite partitioning of E_i , which he has chosen once and for all out of some set of available partitionings. Whenever e_i changes, he conducts a *search* to find that set of his partitioning which contains the new e_i . Let σ_1, σ_2 denote the two chosen partitionings. Manager i reports the set he finds, say $S_i \in \sigma_i$, to HQ. Then HQ chooses an action $\hat{a}(S_1, S_2)$ which maximizes the conditional expectation $\mathcal{E}\left(W(a, e) \mid e_1 \in S_1, e_2 \in S_2\right)$, where the symbol \mathcal{E} again denotes expectation. Let $\hat{\Omega}(S_1, S_2)$ denote the value of that conditional expectation when the maximizing action \hat{a} is used. The *highest attainable expected gross performance* of the two chosen partitionings, which we simply call *gross performance* (for brevity), is the expected value of $\hat{\Omega}$, where the expectation is taken over all the possible pairs (S_1, S_2) . Our symbol for gross performance is just Ω . Thus

$$\Omega(\sigma_1, \sigma_2) = \mathcal{E}_{(S_1, S_2) \in \sigma_1 \times \sigma_2} \hat{\Omega}(S_1, S_2).$$

Even though the technology of search may be advanced, the managers' search has a cost, since it requires human time and human expertise. We suppose that the cost of searching the partitioning σ_i is measured by $\theta \cdot C(\sigma_i)$, where θ and the values taken by the function C are positive real numbers. **When search technology improves, θ drops.** Consider two definitions of $C(\sigma_i)$. The first is simply the number of sets in σ_i . The second is the "Shannon content" of σ_i , i.e.,

$$- \sum_{S_i \in \sigma_i} (\text{probability of } S_i) \cdot (\log_2(\text{probability of } S_i)).$$

The number-of-sets measure ignores the fact that some sets occur more frequently than others. But it is an appropriate measure if the searchers who assist the manager have to maintain their ability to distinguish between the sets. That may require substantial training, and the number-of-sets cost measure may be viewed as the opportunity cost of the investment made in such training. As the technology of search improves, the training required to distinguish among a given number of sets becomes less costly. On the other hand, the Shannon content *is* sensitive to the set probabilities. Using the most elementary of the theorems in the Information Theory which Shannon founded (the noiseless coding theorem)⁵⁴, one can show that if θ is small then the partitioning's Shannon content approximately equals the average number of steps required when a searcher follows a well-chosen binary

⁵³Studied in Marschak (2004).

⁵⁴See, for example, Abramson (1963), pp. 72-74.

tree to find the set in which his current e_i lies. Then θ times the Shannon content is (approximately) the average amount of dollars paid for searcher time when search is efficient. That drops when search technology improves (i.e., when θ drops).⁵⁵

We now compare a *centralized* organization with a *decentralized* one. Here we use the term “decentralized” in a new way, appropriate to the discussion of incentives. In the decentralized organization the managers are self-interested and are free to make choices that were not available to them in the centralized case. In the centralized case, monitoring (or perhaps socialization) insures that each manager strikes the correct balance between his search cost and the gross performance that his search permits. So the chosen partitionings, say σ_1^*, σ_2^* are *first best*. They maximize

$$\Omega(\sigma_1, \sigma_2) - \theta \cdot C(\sigma_1) - \theta \cdot C(\sigma_2).$$

In the decentralized case, there is no monitoring and each self-interested manager i is free to choose his preferred σ_i . He bears the associated cost, namely $\theta \cdot C(\sigma_i)$. But he is rewarded with a share of the expected gross performance Ω . So the two decentralized managers play a *sharing game* in which Manager 1 chooses σ_1 , Manager 2 chooses σ_2 , and each manager i collects an (average) payoff equal to

$$\rho \cdot \Omega(\sigma_1, \sigma_2) - \theta \cdot C(\sigma_i),$$

where ρ , each manager’s share, is greater than zero and not more than $\frac{1}{2}$. In the decentralized case, the chosen partitionings comprise a Nash equilibrium of the sharing game.

The sharing-game interpretation of decentralization seems very natural if one seeks a workable and plausible model. “Profit-sharing” is, after all, the oldest (and simplest) of the schemes that one sees when one looks at decentralized real-world attempts to reconcile individual incentives with organizational goals. Such schemes will not perform as well as first-best choices would, but they may be the best practical alternative if the only way to ensure first-best choices is to engage in intrusive policing, or to adopt (in a real-world setting) some version of the sophisticated but informationally costly monitoring schemes that theorists have proposed.

Now let $(\sigma_1^\dagger, \sigma_2^\dagger)$ denote a pair of decentralized partitionings, chosen at an equilibrium of the sharing game. Our central interest is the *Decentralization Penalty when θ is the level of search technology*. The Penalty associated with $\sigma_1^*, \sigma_2^*, \sigma_1^\dagger, \sigma_2^\dagger$ is

$$P(\theta) = \left[\Omega(\sigma_1^*, \sigma_2^*) - \theta C(\sigma_1^*) - \theta C(\sigma_2^*) \right] - \left[\Omega(\sigma_1^\dagger, \sigma_2^\dagger) - \theta C(\sigma_1^\dagger) - \theta C(\sigma_2^\dagger) \right].$$

In a *shirking equilibrium* of the decentralized sharing game, the managers’ total search expenditures are less than the first-best expenditures. In a *squandering equilibrium* the reverse is true. It turns out that under plausible assumptions a search technology improvement (a drop in θ) *decreases* the Decentralization Penalty associated with a squandering equilibrium, but its effect on the penalty associated with a shirking equilibrium is not so clear-cut.

⁵⁵Once the two partitionings σ_1, σ_2 have been specified, they define — using our previous terminology — a speak-once-only mechanism with HQ as the action-taker. For every e , the action-taker computes the function $F(e)$ which takes the value $\hat{a}(S_1, S_2)$ when $e_1 \in S_1, e_2 \in S_2$.

First consider a “short-run” setting, where a finite collection of possible partitionings is available to each manager, and the technology parameter θ changes but its values lie in some interval $[R, S]$, where $0 < R < S$. The interval is sufficiently small that when θ drops within that interval there is no change in the decentralized managers’ partitioning choices or in the first-best partitionings. For all θ in the interval $[R, S]$, the former partitionings are $\sigma_1^\dagger, \sigma_2^\dagger$ and the latter are σ_1^*, σ_2^* . We can rearrange the expression for the associated Decentralization Penalty to obtain:

$$P(\theta) = \Omega(\sigma_1^*, \sigma_2^*) - \Omega(\sigma_1^\dagger, \sigma_2^\dagger) + \theta \cdot \left[\left(C(\sigma_1^\dagger) + C(\sigma_2^\dagger) \right) - \left(C(\sigma_1^*) + C(\sigma_2^*) \right) \right].$$

Examining the term in square brackets, we immediately see that if squandering occurs at the decentralized partitionings $\sigma_1^\dagger, \sigma_2^\dagger$, then a drop in θ *decreases* the Decentralization Penalty, but if shirking occurs at the decentralized partitionings, then a drop in θ *increases* the Decentralization Penalty.

Now let us turn to a “long-run”, where an infinity of possible partitionings is available and, in addition, the first-best and decentralized (sharing-game equilibrium) partitionings change whenever θ changes. Let $D_i^*(\theta) > 0$ identify Manager i ’s first-best partitioning and let $D_i^\dagger(\theta) > 0$ identify the decentralized Manager i ’s partitioning at an equilibrium of the sharing game. Let $\theta \cdot C(D)$ be the cost of the partitioning D , where C is increasing. Then for a given θ , the Decentralization Penalty is

$$P(\theta) = \left(\Omega(D_1^*(\theta), D_2^*(\theta)) - \theta C(D_1^*(\theta)) - \theta C(D_2^*(\theta)) \right) - \left(\Omega(D_1^\dagger(\theta), D_2^\dagger(\theta)) - \theta C(D_1^\dagger(\theta)) - \theta C(D_2^\dagger(\theta)) \right).$$

Suppose the functions D_i^*, D_i^\dagger are differentiable.⁵⁶ For every θ , the first best partitionings D_i^* have to satisfy the first-order condition for the maximization of $\Omega - \theta C_1 - \theta C_2$, and the decentralized partitionings D_i^\dagger have to satisfy the first-order condition for D_1 to be a best reply (in the sharing game) to D_2 , and vice versa. When we compute the derivative P' , while inserting the first order conditions, we obtain:

$$\begin{aligned} P'(\theta) = & \left[C(D_1^\dagger(\theta)) + C(D_2^\dagger(\theta)) - \left(C(D_1^*(\theta)) + C(D_2^*(\theta)) \right) \right] \\ & - \left[\theta \left(\frac{1}{\rho} - 1 \right) \left(D_1^{\dagger'}(\theta) C'(D_1^\dagger(\theta)) + D_2^{\dagger'}(\theta) C'(D_2^\dagger(\theta)) \right) \right]. \end{aligned}$$

Here primes denote derivatives and ρ , with $0 < \rho \leq \frac{1}{2}$, is each manager’s share of the gross performance Ω in the sharing game which the managers play when they are decentralized. Now suppose we know the following:

(#) **when search technology improves (θ drops), the decentralized managers spend more on search, i.e., costlier partitions are chosen at the equilibria of the sharing game**

⁵⁶The differentiability assumption is made for analytic convenience. Typically one would want each of the available partitionings to be identified by a positive integer T . (Thus Manager i ’s external variable e_i might have the interval $[A, B]$ as its support, with $0 < A < B$, and T might identify the partitioning in which that interval is divided into T sub-intervals of equal length). Under very plausible assumptions about Ω , the finding that the derivative P' is negative (when we treat D_i^*, D_i^\dagger as continuous and differentiable rather than integer-valued) implies that for sufficiently small θ , P remains a decreasing function when we restrict the functions D_i^*, D_i^\dagger to integer values.

Then $D_i^\dagger(\theta) < 0, i = 1, 2$ for all $\theta > 0$. Since C is increasing and $\frac{1}{\rho} - 1 > 0$, the second of the two terms in square brackets is negative or zero. *If there is squandering at the decentralized partitionings*, then the first of the two terms in square brackets is positive and hence the entire expression for P' is positive, i.e., *the Decentralization Penalty shrinks when IT improves*. If there is shirking, then the sign of P' is indeterminate unless we make further assumptions.

Some of the scarce empirical literature on the effect of IT (Information Technology) improvement on organizational structure suggests that as IT improves, the case for “decentralizing” an organization becomes stronger.⁵⁷ The model of managers who engage in search is a first attempt to see why this might be so. If the IT improvement in question is the explosive growth of efficiently searched databases, and if “decentralized” managers can reasonably be modelled as players of a sharing game, then it is of considerable interest to find the conditions under which improved search technology indeed implies a drop in the Decentralization Penalty. An explicit formula for the Penalty has been found for certain cases.⁵⁸ But there also examples of a W and a probability distribution on the e_i in which the decentralized sharing game has a squandering equilibrium.⁵⁹ The effect of a drop in θ on the Decentralized Penalty associated with squandering is strikingly unambiguous, both in the “short-run” setting, and (provided the statement (#) holds) in the long-run setting as well. So it would be valuable to understand the squandering phenomenon much better. For a very wide class of sharing games a sufficient condition for *ruling out* squandering is *complementarity*: an increment in Manager 1’s search expenditures improves (or does not damage) Manager 2’s “marginal productivity”, i.e., the increment in Ω due to an increment in Manager 2’s search expenditures.⁶⁰ On the other hand a class of functions Ω has been found for which the following is true: if Ω exhibits sufficiently strong “*anticomplementarity*”, then there are squandering equilibria in the sharing game defined by Ω and the shares ρ_i . (Anticomplementarity means that an increment in Manager 1’s search expenditures damages

⁵⁷See the papers mentioned in the Introduction: Bresnahan, Brynjolffson, and Hitt (2000), (2002).

⁵⁸Suppose, in particular, that the payoff function W has the linear/quadratic structure discussed in section 2.6 above, where the techniques of the Theory of Teams were considered. For that W , person-by-person-satisfactoriness is both necessary and sufficient for a team action rule to be best for a given information structure. That allows us to find the best action rules used by HQ once each manager tells HQ the set in which his current local environment lies. Suppose further that each e_i is uniformly distributed on a closed interval, that each manager i divides his interval into T_i subintervals of equal length, and that each choice of T_i defines one of the manager’s possible partitionings. It then turns out that (1) in the unique equilibrium of the sharing game, each manager shirks (chooses a T_i that is lower than the first-best T_i), (2) in the “long-run” setting (where a small change in the technology parameter θ leads to new equilibrium values of the T_i as well as new first-best values) the Decentralization Penalty indeed drops when θ drops. That is true for the number-of-sets cost measure as well as the Shannon cost measure.

⁵⁹In the examples found so far, there is a finite set of possible values for each e_i .

⁶⁰Consider n -person sharing games in which each person i chooses a strategy x_i , bears the nondecreasing cost $c_i(x_i)$, and receives a reward $R_i(z)$ when the organization earns the revenue $z = A(x_1, \dots, x_n)$, where A is nondecreasing in each of its arguments. Suppose the functions R_i obey a “nondecreasing residual” property: when z increases by any positive amount Δ , the sum of the rewards does not rise by more than Δ . (That is satisfied, for example in the “balanced-budget” case where $z = \sum R_i(z)$ as well as the “constant-share” case $R_i(z) = \rho_i \cdot z$ which we have been considering.) Suppose A obeys complementarity: if we look at the increment in A when any player spends more, we find that the increment does not drop when some other player spends more. Then in any equilibrium no player squanders relative to an efficient (first-best) (x_1, \dots, x_n) , where A minus the cost sum is maximal. Every player spends the efficient amount or less. The argument holds whether the strategy sets are finite or are continua. The argument is provided in Courtney and Marschak, 2004.

Manager 2’s marginal productivity).

In an informal way, one can imagine that the much discussed but seldom modeled phenomenon of “information overload” can indeed lead to anticomplementarity and to squandering. As Manager 1 increases his search effort and the level of detail in his report to HQ, he greatly diminishes the value of Manager 2’s report, because Manager 2’s report then becomes somewhat *redundant*, i.e., its marginal contribution is small. The managers might then be trapped in an equilibrium where Manager 1’s partitioning is very costly, Manager 2’s partitioning is cheap, and the sum of the costs is higher than the cost sum for a first-best partitioning pair. A rough conjecture would be as follows:

Conjecture (*): For sets $G \subseteq E_1, H \subseteq E_2$, let $\hat{a}(G, H)$ denote a value of a which maximizes the conditional expected value of the payoff $W(a, e_1, e_2)$, given that $e_1 \in G, e_2 \in H$. Given the probability distribution of $e = (e_1, e_2)$, the payoff function W , and partitionings σ_1, σ_2 , let us measure the marginal contribution of σ_2 by

$$\mathcal{E}_{Q_1 \in \sigma_1, Q_2 \in \sigma_2} \left[\mathcal{E} \left(W(\hat{a}(Q_1, Q_2), e_1, e_2) \mid e_1 \in Q_1, e_2 \in Q_2 \right) \right] - \mathcal{E}_{Q_1 \in \sigma_1} \left[\mathcal{E} \left(W(\hat{a}(Q_1, E_2), e_1, e_2) \mid e_1 \in Q_1 \right) \right].$$

Then if the marginal contribution of σ_2 is “sufficiently small but not too small”, (σ_1, σ_2) is a squandering equilibrium of the sharing game.

Here “but not too small” seems appropriate, since if σ_2 makes an extremely small marginal contribution, then Manager 2 finds that when compared to the one-set partitioning (the cheapest partitioning), his share of σ_2 ’s extra revenue fails to justify its extra cost.

RESEARCH CHALLENGE #9: To better understand the squandering phenomenon, find probability distributions on the local environments, and payoff functions W , such that

- Condition (#) holds: when search technology improves, the decentralized managers spend more on search at the sharing-game equilibrium. (Then, in the “long-run” setting, a search-technology improvement diminishes the Decentralization Penalty associated with squandering)
- Conjecture (*) holds.

3.2.2 The effect of search-technology improvement on the Coordination Benefit. In addition to studying the Decentralization Penalty, one can conduct a parallel investigation of the effect of improved search technology on the Coordination Benefit. That might be motivated by informal suggestions, in some empirical studies, that improved IT leads to greater “lateral communication”⁶¹ To define coordination in our two-manager model, let the organization’s action have two components: $a = (a_1, a_2)$, where a_i is associated with manager i . Then coordination for the centralized organization (where managers are loyal team members and are not self-interested) means that HQ chooses a_i as a function of both Managers’ search results, while no coordination means that each a_i is chosen as a function of i ’s results only. The Coordination Benefit is the improvement in expected net payoff — i.e.,

⁶¹Again, see Bresnahan, Brynjolffson, and Hitt, 2002.

Ω minus the sum of the search costs — when we move from no coordination to coordination. To define the concept in the decentralized (sharing-game) case, suppose that HQ no longer exists. Instead, the self-interested Manager i chooses a_i himself. Coordination now means that i bases his choice on both manager’s search results (after “lateral communication” occurs), and no coordination means that he bases it on his own results only. Two different sharing games are thus defined. The Coordination benefit is the improvement in net payoff when we move from an equilibrium of the first game to an equilibrium of the second. Classes of probability distributions and payoff functions have been found for which it is indeed the case that when search technology improves (θ drops), the Coordination Benefit rises, in both the centralized and the decentralized cases.⁶²

3.2.3 Extending the model so that managers may be decentralized with regard to “production” as well as search; studying the effect of a search-technology improvement on the Decentralization Penalty when the “degree of decentralization” is varied. Suppose that in our two-manager model the organizational payoff (or gross profit) function W takes the following form:

$$W = V(a, e) - K_1(a_1) - K_2(a_2).$$

For the action $a = (a_1, a_2)$, the function K_i is Manager i ’s *local production cost* while $V(a, e)$, which may reflect externalities between the two managers, is the gross revenue earned by the organization when the environment is e . We modify our model of the decentralized organization. Each manager i now chooses (1) a partitioning of E_i (as before) and (2) a rule α_i which assigns a value of a_i to every possible search result. He receives a share $\rho \cdot V$ of V , but has to pay a share γ of his production cost, with $0 < \gamma < 1$. Thus his net payoff in the game is the expected value of $\rho \cdot V$ minus the expected value of $\gamma \cdot K_i$ (when he uses the rule α_i) minus the search cost for his chosen partitioning. In the first-best (centralized) situation, by contrast, the partitionings and the rules α_i are chosen so as to maximize the expected value of $V - K_1 - K_2$ minus the total search costs. We can call γ the *degree of decentralization*. As γ rises towards one, each manager bears more and more responsibility for his own production costs. We can again study the Decentralization Penalty, i.e., the amount by which the expected value of $[V - K_1 - K_2] -$ (the total search costs) in the decentralized case falls short of its first-best value. The Penalty depends on γ and on θ . We now have a new question.

RESEARCH CHALLENGE #10: When is it the case that a search technology improvement (a drop in θ) “justifies” a higher degree of decentralization? When is it the case, in other words, that the value of γ which minimizes the Decentralization Penalty for $\bar{\theta}$ exceeds the value of γ which minimizes the Penalty for $\theta > \bar{\theta}$?

This is a subtle question. In answering it, a key role is played by another question: for a given level of search technology (a given θ), is the value (to a manager) of another dollar spent on search higher for low γ or for high γ ? What is intriguing about the latter question is that one’s off-the-cuff intuition can go either way. One could very roughly argue that “when γ goes up, each manager ends up with

⁶²Suppose, once again, that the payoff function W has the linear/quadratic structure and each e_i is uniformly distributed on a closed interval. Then the methods of the Theory of Teams again allow us to find explicit formulae for the action rules used by HQ in the centralized/ coordinated and centralized/uncoordinated cases, and the rules used by each manager in the decentralized/coordinated and decentralized /uncoordinated cases. In all cases one finds (for the number-of-sets cost measure and for the Shannon cost measure) that the Coordination Benefit rises when θ drops.

less for every action, so learning a little more about e_i becomes less valuable to the manager”. But one could just as well argue, in an equally primitive way, that “when γ goes up, it becomes more important for each manager to get his action just right, so learning more about e_i becomes more valuable.”

One could start learning about the matter by studying a *one-manager version*. To give it a concrete setting, let the organization consist of a manufacturer and a franchised retailer. The manufacturer can produce and ship the product quantity q to the retailer at a cost $C(q)$. He charges the franchised retailer a fee of γ for each unit shipped. We can interpret γ as the “degree of decentralization”; the higher γ is set, the higher the retailer’s contribution to the manufacturing cost. Once the manufacturer chooses γ , the retailer is entirely on her own. She sells the received product in a local market. The demand curve she faces is

$$P = e - Q,$$

where P is price, Q is quantity, $e > 0$ and $Q \leq e$. But e is a random variable. It changes each “week” and its possible values comprise an interval $[A, B]$ with $B > A$. The retailer has to commit to the quantity she orders for next week’s sales before she knows next week’s demand curve. By spending money on search (market research) she can learn about next week’s e . She chooses a T -interval partitioning of $[A, B]$ from a set of available partitionings and searches to find the interval in which next week’s e lies. Denote the T intervals $I_1, \dots, I_t, \dots, I_T$. The search costs the retailer the amount $\theta \cdot T$. Once the current interval, say I_t , has been found, the retailer chooses a quantity $\hat{Q}^{T\gamma}(t)$ so as to maximize the following expectation:

$$\pi(t, \gamma) = \mathcal{E} \left[(e - Q) \cdot Q - \gamma \cdot Q \mid e \in I_t \right].$$

The retailer chooses T so as to maximize

$$\sum_{t \in \{1, \dots, T\}} (\text{prob. of } I_t) \cdot \pi(t, \gamma) - \theta T.$$

Let $\hat{T}(\gamma, \theta)$ denote the retailer’s chosen T . Then the manufacturer’s expected profit depends on both γ and θ . It is

$$\sum_{t \in \{1, \dots, \hat{T}(\gamma, \theta)\}} (\text{prob. of } I_t) \cdot \left(\gamma \cdot Q^{\hat{T}(\gamma, \theta), \gamma}(t) - C(Q^{\hat{T}(\gamma, \theta), \gamma}(t)) \right).$$

He chooses the franchise fee γ so as to maximize expected profit. We ask: *When is it the case that a drop in θ leads the manufacturer to raise γ ? When is it the case that the drop leads the manufacturer to lower γ ?*⁶³

3.3 Networks of self-interested decision-makers, who bear the network’s informational costs

There is an abundant literature on the formation of networks, where every node is occupied

⁶³Preliminary exercises suggest that both can happen. Going back to the general one-manager problem, consider two examples: (1) $V(a, e) = ea$ and $K(a) = a^2$; (2) $V = J - \frac{1}{ae}$, $K = b \ln a$, where $J > 0, b > 0, a > 0$. In both cases, the Manager collects (for given a, e) the amount $V - \gamma K$, where $0 \leq \gamma \leq 1$. In both cases, let e be uniformly distributed on $[A, B]$ and let the partitioning defined by T consist of T equal-length intervals. The partitioning costs the Manager θT . It turns out that if we keep θ fixed, then in the first example, raising γ leads the manager to choose a *lower* value of T , while in the second example it leads him to choose a *higher* value.

by a self-interested person, who obtains information from those persons to whom he is linked.⁶⁴ That literature usually *takes as given* a “value function” defined on the possible networks. The function expresses the network’s collective performance. Much of the literature then considers alternative “allocation rules”. Such a rule assigns a share of the network’s value to each of its participants. Suppose we specify a rule, and we also specify the cost of forming a link, borne by the two persons who form the link. Then we have a game, in which each person chooses the links to be formed with others, and each person collects his share of value (for the resulting network) minus his link costs. The equilibria of these games are studied, with particular attention paid to the structure of the equilibrium networks. (When is the equilibrium network a tree? When is it a ring?) The literature has developed a great many interesting results about equilibrium networks.

In this framework, we could, in principle, study the impact of improved Information Technology. *Given* a value function and an allocation rule, how does a drop in link costs (due to IT improvement) change the structure of the equilibrium networks?

Suppose, however, that we venture in a new direction and we no longer take the value function and the allocation rule as exogenously given. Instead we let each of n persons, say person j , observe some external random variable e_j . In addition, person j learns the current value of e_k , for all persons k who are *neighbors* of j , i.e., there is a link between i and k . Finally, we let j be a decision-maker. He responds to the current value of his own and his neighbors’ external environments by taking an *action*. He then collects a gross revenue U_j , which depends on his action, the others’ actions, and the current value of $e = (e_1, \dots, e_n)$. Let V_j denote the expected value of U_j . Person j ’s net payoff is V_j minus his share of the cost of his links. (In one version we assume that a link’s cost is shared equally between its two users).

In our network formation game, each person chooses his *intended* neighbors and he also chooses a rule that tells him, what action to choose for each value of his intended neighbors’ and his own current external variables. When all n persons have made these choices, a network emerges. That network has a link between j and some k if and only if both j and k have chosen the link.

RESEARCH CHALLENGE #11: A collection of related questions awaits exploration:

- Which networks and action-rule n -tuples are *stable*, i.e., they have the property that if a player changes his intended neighbors or his action rule (while no one else makes any changes) then his expected payoff drops or stays the same?
- Which networks and action-rule n -tuples are *efficient*, i.e., $\sum_{i=1}^n V_i$ minus the total link costs is as high as possible?
- When is a stable network efficient and vice versa?
- When is a stable network connected (there is a path between any two persons)? When is an efficient network connected?
- When is a stable network a tree and when is it a ring? When is an efficient network a tree and when is it a ring?

⁶⁴Two surveys are: M.O. Jackson (2003), (2004).

- For the case of a tree, does its number of *tiers* shrink as link costs go down — i.e., can we substantiate the classic claim that an improvement in IT indeed leads to the “flattening of hierarchies”?
- When does an inefficient stable network exhibit shirking (total link costs are less than in an efficient network) and when does it exhibit squandering (total link costs exceed those of an efficient network)?

Clearly this is an ambitious agenda. The existing literature has made progress on some of these questions, because it suppresses the decision-making role of each person and takes network value and its allocation as given. The agenda is even more challenging in the proposed new research path, where each person is explicitly a decision-maker. Nevertheless, progress has been made in the following case.⁶⁵ Suppose the highest attainable value of the expected-revenue function V_j has the property that it depends *only on the number of j 's neighbors*.

Here is a simple example where this is indeed so. Let each person j have one unit of a product to sell. Each of the n persons is located at a market where the product can be sold. The random variable e_i is the product's current price at location i . Given what he knows about the prices at his neighbors' locations, person j has to choose the location where he will sell his product. That can be his own location, it can be the location of a neighbor, or it can be the location of a non-neighbor. Now suppose each e_i takes just two values: zero and one, each with probability $\frac{1}{2}$, whatever all the other e_k ($k \neq i$) may be. Suppose Person i is given a neighbor set, say the set N_i . He wants to choose a selling location so as to maximize the expected value of his revenue. It is clear that he may as well choose his own location if he observes $e_i = 1$ or if no location k in the set N_i has $e_k = 1$. But if he sees $e_i = 0$ and some neighbor k has $e_k = 1$, then he chooses the location k . His highest attainable expected revenue depends only on $|N_i|$, the *number* of his neighbors. It is

$$h(|N_i|) = 1 - \left(\frac{1}{2}\right)^{|N_i|}.$$

The function h has a fortunate “diminishing marginal product property”: it increases more and more slowly as $|N_i|$ increases.

When the number of neighbors is all that matters, and each person's highest attainable expected revenue depends on the number of his neighbors through an increasing function h which (as in the above example) has the diminishing marginal product property, then a fairly complete analysis is possible. Let J be the link cost, and let $g(J)$ denote the largest value of $|N_i|$ for which $h(|N_i| + 1) - h(|N_i|) \geq J$. Let half of the cost of each link be paid by each of its users. The results one can establish include the following:

- (i) Let us measure a network's net performance by $\sum_{i=1}^n h(|N_i|) - (\text{total link costs})$. If a network is stable and its net performance is maximal among all stable networks, then it is efficient.
- (ii) If $n - 1 \leq g(J)$, then both stable and efficient networks are connected.
- (iii) If $g(J) < \frac{n-1}{2} < g(\frac{J}{2})$, then an efficient network is connected but a stable network need not be.

⁶⁵The results given here are due to Xuanming Su.

- (iv) The net performance of an efficient network rises as J drops (i.e. IT improves), but more and more slowly.
- (v) In an inefficient stable network there is shirking, not squandering.

One might view the requirement that “only the number of neighbors matters” as very strong. One might guess, in particular that it rules out externalities, wherein one person’s action affects the payoffs collected by others. (There are no externalities in our selling-location example). But this need not be so, as the following example shows: Each person j again observes a random variable e_j and again learns the current e_k for every neighbor k . In response, he chooses an action a_j (a real number). When the action vector $a = (a_1, \dots, a_n)$ has been chosen, then for a given $e = (e_1, \dots, e_n)$, the entire organization obtains a payoff

$$W(a, e) = \sum_{i=1}^n e_i a_i - \sum_{i=1}^n \sum_{j=1}^n q_{ij} a_i a_j,$$

where the matrix $((q_{ij}))$ is positive definite (which insures that there is a unique W -maximizing a). Person i ’s own revenue is a share of W , namely $r \cdot W$, where $0 < r \leq \frac{1}{n}$. Thus there are externalities between the actions. Let the e_i be independently distributed and let each be normally distributed with mean zero. For any given network, consider the possible action rules for each person j . The rule assigns a value of a_j to each value of (e_j, e_{N_j}) , where e_{N_j} denotes the external variables observed by j ’s neighbors and hence learned by j as well. Because of the linear/quadratic payoff structure we may use, once again, the methods of the Theory of Teams. It turns out, in spite of the externalities, that in the rule n -tuple which achieves the highest expected value of W attainable by the network, each person’s action depends only on the number of his neighbors. If we now turn to the network-formation game, we find that for a given network, each person’s highest attainable expected payoff (the highest attainable expected value of $r \cdot W$), again depends only on the number of his neighbors, whatever action rules the $n - 1$ persons may have chosen. Moreover the function h , which gives that highest attainable expected payoff, has the required diminishing marginal product property.

4. ORGANIZATIONAL MODELS IN WHICH THE PRIMITIVE IS A “TASK”, “PROBLEM”, “PROJECT”, OR “ITEM”.

There are formal models of organization which do not explicitly consider the actions the organization takes in response to a changing environment. Instead the model supposes that “tasks” “problems”, “projects”, or “information items” flow in from the outside world, and the organization has to process them. They become the primitives of the model and are not further explained. One seeks to characterize organizations which process the flow efficiently, in some appropriate sense. Just as in the approaches we have discussed, processing costs (or capacities) are part of the model. To illustrate, we briefly discuss three studies that share this approach but are otherwise rather different.⁶⁶

In Sah and Stiglitz (1986), the organization has to judge “projects”. A project has a net benefit x , where x is a random variable with known distribution. The organization consists of evaluators. Once

⁶⁶Other studies, which use similar primitives and pay some attention to informational costs, include the following: M. Keren, M and D. Levhari (1983); M. Beckman (1983); A. Arenas, A. Cabrales, L. Danon, A Diaz-Guilera, R. Guimera, and F. Vega-Redondo (2003); B. Visser (2000).

he has received a project, an evaluator either chooses to reject it or to pass it on to another evaluator for further evaluation. The total expected benefit of the project portfolio selected by a "hierarchy" is compared to that achieved by a "polyarchy". In a hierarchy, the evaluators are arranged in two "bureaus", called 1 and 2. All projects first flow to a member of Bureau 1, who either rejects it or passes it onto Bureau 2, where a final judgment is made. The only projects reviewed by Bureau 2 are those it receives from Bureau 1. In a polyarchy, the organization is divided into two "firms". Each receives half of the projects. If it accepts the project no further evaluation occurs; if it rejects the project then the other firm performs a final evaluation. The quality of the evaluators can be varied, by changing the conditional distribution of x given the "reject" judgment. Costs could be attached to quality improvement but the 1986 paper only sketches how that might be done.⁶⁷

In Bolton and Dewatripont (1994), the main primitive is a "cohort of M information items" received by the organization from an external source. Each item is a "type" of information about the outside world. All cohorts yield the same "value" R to the organization, once they are "processed". ("Processing" is another primitive, not further explained). In order for the organization to realize the cohort's value, at least one agent must receive the entire cohort. (Thus far the problem can be restated as the choice of a speak-once-only mechanism, with an action-taker who needs to know the entire cohort). While one can study delay (time until at least one agent knows the entire cohort), the paper emphasizes another question. It supposes that economies of scale are achieved when a given agent processes a given type more and more frequently. A network of agents is given and each can be made a specialist in one or more items; he processes only those. For each given network, one seeks an assignment of types to agents so as to minimize "the total labor time spent per processed cohort". In efficient networks the labor time for a best assignment is minimal. Some suggestive properties of efficient networks are worked out.

In Garicano and Rossi-Hansberg (forthcoming), the agents who form organizations belong to a "knowledge economy". In each time period each agent receives a problem whose level of difficulty can vary, and solves it if its difficulty is below her level of knowledge. ("Problem", "level of difficulty", "solve", and "level of knowledge" are primitives, not further explained). Each problem is identified by a value of Z , a nonnegative number; a higher Z means the problem is harder. The population of possible problems has a known probability distribution over the possible values of Z . Each agent is endowed with a "cognitive ability" α , a random variable with known distribution. By incurring a cost, an agent can learn to handle all problems of difficulty up to a given level z . The cost is increasing in z and decreasing in α . An agent receives an income, which is increasing in the proportion of problems the agent is able to solve. An agent seeks to maximize income minus learning costs. But high-ability agents can help low-ability agents, and that leads to the formation of organizations. The structure of those organizations in an equilibrium of the economy is studied.⁶⁸

5. CONCLUDING REMARKS

⁶⁷In a further paper (Sah and Stiglitz (1988)), the organization becomes an n -person "committee". All n persons judge each project and it is accepted if $k \leq n$ members judge it favorably. This time the tradeoff between portfolio quality and cost is central. The cost measure is simply n itself.

⁶⁸Another model in which "task" is a primitive is developed in Malone (1987) and Malone and Smith (1988). These papers study a number of cost measures for a given organizational structure.

We have followed a modeling path with four elements: environment, action, goal, and informational cost. The models describe, in detail, the process by which the organization reaches new actions when the environment changes, i.e., they describe the mechanism that the organization chooses. This is not an easy research path. A great deal more work needs to be done before one can judge whether the effort is worthwhile. In the present state of knowledge there is a formidable gulf between the propositions that can be proved and the complexities of organization and information as they appear in the real world. Yet without theory, it is hard to make sense of that reality and to see why casual claims, like those about the impact of IT on organizations, might or might not be true. Perhaps models that omit one or more of our four elements, or forego the detailed description that the concept of mechanism permits, or use entirely different primitives, may prove more useful guides — for the time being — to persons who are able to observe real organizations and assemble new datasets. But surely all four of our elements arise in real organizations, and real organizations follow *some* procedure in choosing new actions. It seems inevitable that future empirical work will eventually try to examine that procedure in detail and will deal, in one way or another, with all four of our elements.

Comparing mechanisms with regard to informational cost is particularly tricky. Minimal message-space size provides one fundamental way of judging the complexity of an organizational goal, and it tells us, in a preliminary and abstract way, how expensive one goal is relative to another if the goal is to be met by an organization which is decentralized in the sense that each member privately observes some aspect of the environment. But delay, the number of persons, and the individual communication and computation burden that each person faces are all important as well. Modeling of those costs is still in an early stage.

We have looked primarily at work conducted by economic theorists. But there are parallel efforts in computer science and in artificial intelligence. Theoretical research that bridges disciplines is finally emerging.⁶⁹ There are also many parallel efforts by social scientists who are not economists (e.g., persons in the Organizational Behavior field). A review of that literature would doubtless paint a very different picture as to what has been learned or could be learned about the effect of IT advances on organizational structure. Economic theorists are endowed (or perhaps burdened!) with a certain point of view when they approach such a challenging question. That point of view has thoroughly permeated this Chapter.

REFERENCES

- H. Abelson, 1980, “Lower bounds on information transfer in distributed computations,” *Journal of the ACM*, 27, 384.

⁶⁹As noted in Section 2.2.8, the computer science specialty called Communication Complexity deals with minimal privacy-preserving mechanisms (protocols) for an n -person organization whose goal is the computation of an action that depends on n variables, each privately observed by one of the n persons. Both communication-complexity results and economic-theory results are central to the bridge-building paper by Nisan and Segal (2005) on the allocation of many objects among many persons, mentioned in Section 2.2.8. At present it seems fair to say that there is very much more work by theoretical computer scientists influenced by economic ideas, than work by economic theorists influenced by computer science ideas. In particular, computer scientists have explored the merits of market-based approaches to computational problems. Many papers which do so could be cited. A few of them are: X. Deng, C. Papadimitriou, and M. Safra (2002); W. Walsh, M. Yokoo, K. Hirayama, and M. Wellman (2003); W. Walsh and M. Wellman (2003).

- N. Abramson, 1963, *Information Theory and Coding*, McGraw Hill, New York.
- Aizpura, J. and A. Manresa, 1995, A decentralized and informationally efficient mechanism realizing fair outcomes in economies with public goods, *Economic Design*, 1, 141.
- Alon, N. and D. Kleitman, 1986, Covering a square by small perimeter rectangles, *Discrete and Computational Geometry*, 1, 1.
- Arenas, A., L. Cabrales, et al, 2003, Optimal information transmission in organizations: search and congestion”, working paper, Universitat Pompeu Fabra.
- Apostol, T., 1957, *Mathematical Analysis*, Addison-Wesley, Reading, Mass.
- Beckman, M., 1983, *Tinbergen Lectures on Organization Theory*, Springer, Heidelberg.
- Bolton, P. and M. Dewatripont, 1994, The firm as a communication network”, *Quarterly Journal of Economics*, CIX, 809.
- Bresnahan, T. E. Brynjolfsson, and L.Hitt, 2000, Technology, organization, and the demand for skilled labor, *The New Relationship: Human Capital in the American Corporation*, Eds. M. Blair and T. Kochan, Brookings Institution, Washington.
- Bresnahan, T. E. Brynjolfsson, and L.Hitt, 2002, Information technology, workplace organization, and the demand for skilled labor, *Quarterly Journal of Economics*, CXVII , 339.
- Calsamiglia, X., 1977, Decentralized resource allocation and increasing returns”, *Journal of Economic Theory*, 14, 263.
- Calsamiglia, X. and Kirman, A., 1998, A unique informationally efficient decentralized mechanism with fair outcomes, *Econometrica*, 61, 1147.
- Chander, P., 1983, On the informational size of message spaces for efficient resource allocation processes”, *Econometrica*, 51, 919.
- Courtney, D. and T. Marschak, 2004, Sharing games, working paper.
- Creti, A., 2001, Firms’ organization and efficient communication networks, *The Manchester School*, 69, 77.
- Deng, X, C. Papadimitriou, and M. Safra, 2002, On the complexity of equilibria”, *Symposium on the Theory of Computing*.
- Garicano, L. and E. Rossi-Hansberg, forthcoming, Organization and inequality in a knowledge economy”, *Journal of Political Economy*.
- Groves, T. and J. Ledyard, 1977, Optimal allocation of public goods, *Econometrica*, 45, 783.
- Harris, M., and A. Raviv, 2002, Organization design, *Management Science*, 48, 852.
- Hurwicz, L., 1960, Optimality and efficiency in resource allocation, *Mathematical Methods in the Social Sciences*, Eds. K. Arrow, S. Karlin, P. Suppes, Stanford University Press, Stanford Ca.

- Hurwicz, L., 1977, On the dimensional requirements of informationally decentralized Pareto-satisfactory processes", *Studies in Resource Allocation Processes*, Eds. K. Arrow and L. Hurwicz, Cambridge University Press, New York.
- Hurwicz, L., 1986, On informational decentralization and efficiency in resource allocation mechanisms, *Studies in Mathematical Economics*, Ed. S. Reiter, Mathematical Association of America, Washington.
- Hurwicz, L. and T. Marschak, 1985, Discrete allocation mechanisms: dimensional requirements for resource allocation mechanisms when desired outcomes are unbounded, *Journal of Complexity*, 1, 264.
- Hurwicz, L. and H. Weinberger, 1990, A necessary condition for decentralization and an application to intertemporal allocation", *Journal of Economic Theory*, 51, 313.
- Hurwicz, L. and T. Marschak, 2003a, Comparing finite mechanisms, *Economic Theory*, 21, 783.
- Hurwicz, L. and T. Marschak, 2003b, Finite allocation mechanisms: approximate Walrasian versus approximate direct revelation, *Economic Theory*, 21, 545.
- Hurwicz, L. and T. Marschak, 2004, The informational efficiency of finite price mechanisms, *Assets, Beliefs, and Equilibria in Economic Dynamics*, Eds. Aliprantis, C. et al, Springer, Heidelberg.
- Hurwicz, L. and S. Reiter, 2006, *Designing Economic Mechanisms*, Cambridge University Press, New York.
- Ishikida, T., and T. Marschak, 1996, Mechanisms that efficiently verify the optimality of a proposed action, *Economic Design*, 2, 33-68.
- Jackson. M.O., 2001, A crash course in implementation theory, *Social Choice and Welfare*, 18, 655.
- Jackson, M.O., 2003a, The stability and efficiency of economic and social networks, *Networks and Groups: Models of Strategic Formation*, Eds. B. Dutta and M. Jackson, Springer, Heidelberg.
- Jackson, M.O., 2003b, A survey of models of network formation: stability and efficiency, *Group Formation in Economics: Networks, Clubs, and Coalitions*, Eds. G. Demange and M. Wooders, Cambridge University Press, New York.
- Jordan, J., 1982, The competitive allocation process is informationally efficient uniquely, *Journal of Economic Theory*, 28, 1.
- Jordan, J., 1987, The informational requirements of local stability in decentralized allocation mechanisms, *Information, Incentives, and Economic Mechanisms: Essays in Honor of Leonid Hurwicz*, Eds. T. Groves, R. Radner, and S. Reiter, University of Minnesota Press. Minneapolis.
- Jordan, J., 1995, Information flows intrinsic to the stability of economic equilibrium, *The Economics of Informational Decentralization: Complexity, Efficiency, and Stability*, Ed. J. Ledyard, ed., Kluwer Publishing Company, Boston.
- Jordan, J. and D. Xu, 1999, On the Communication Complexity of Expected Profit Maximization, *Journal of Economic Theory*, 86, 185.
- Kaganovitch, M., 2000, Decentralization of intertemporal economies with discounted maximin criterion, *International Economic Review*, 41, 1021.

- M. Karchmer, 1989, *Communication Complexity: A New Approach to Circuit Depth*, MIT Press, Cambridge Mass.
- Keren, M., and D. Levhari, 1983, The internal organization of the firm and the shape of average costs, *Bell Journal of Economics*, XL, 474.
- Kushilevitz, E. and N. Nisan, 1997, *Communication Complexity*, Cambridge University Press, New York.
- Leavitt, H. and T. Whistler, 1958, Management in the 1980s", *Harvard Business Review*, 36, 6.
- Leontieff, W., 1947, A note on the interrelation of subsets of independent variables of a continuous function with continuous first derivatives, *Bulletin of the American Mathematical Society*, 53, 343.
- Lovász, L., 1990, *Communication complexity: a survey*, *Paths, Flows, and VSLI Layouts*, Ed. B.H. Korte, Springer-Verlag, Berlin.
- Malone, T., 1987, Modeling coordination in organizations and markets, *Management Science*, 53, 1317.
- Malone, T. and S. Smith, 1988, Modeling the performance of organizational structures", *Operations Research*, 36, 421.
- A. Manresa, 1993, An infinite dimensional extension of the theory of decentralized mechanisms, *Mathematical Social Science*, 26, 157.
- Marschak, J. and R. Radner, 1972, *The Economic Theory of Teams*, Yale University Press, New Haven.
- Marschak, T., 1959, Centralization and decentralization in economic organizations", *Econometrica*, 27, 399.
- Marschak, T., 1972, Computation in organizations: the comparison of price mechanisms and other adjustment processes", *Decision and Organization*, Eds. C.B. McGuire and R. Radner, North-Holland, Amsterdam.
- Marschak, T., 1987, Price versus direct revelation: informational judgments for finite mechanisms, *Information, Incentives, and Economic Mechanisms: Essays in Honor of Leonid Hurwicz*, Eds. T. Groves, R. Radner, and S. Reiter, University of Minnesota Press, Minneapolis.
- Marschak, T. and S. Reichelstein, 1995, Communication requirements for individual agents in networks and hierarchies, *The Economics of Informational Decentralization: Complexity, Efficiency, and Stability*, Ed. J. Ledyard, Kluwer Publishing Company, Boston.
- Marschak, T., and S. Reichelstein, 1998, Network mechanisms, informational efficiency, and hierarchies, *Journal of Economic Theory*, 44, 106.
- Marschak, T., 2004, Information technology and the organization of firms, *Journal of Economics and Management Strategy*, 13, 473.
- Melumad, N., D. Mookherjee, and S. Reichelstein, 1992, A theory of responsibility centers, *Journal of Accounting and Economics*, 15, 445.
- Melumad, N., D. Mookherjee, and S. Reichelstein, 1995, Hierarchical decentralization of incentive contracts, *RAND Journal of Economics*, 26, 654.

- Melumad, N., D. Mookherjee, and S. Reichelstein, 1997, Contract complexity, incentives, and the value of delegation, *Journal of Economics and Management Strategy*, 6, 257.
- Mount, K. and S. Reiter, 2002, *Communication and Complexity in Economic Behavior and Organizations*, Cambridge University Press, New York.
- Mount, K. and S. Reiter, 1996, A lower bound on computational complexity given by revelation mechanisms, *Economic Theory*, 7, 237.
- Mount, K. and S. Reiter, 1998, A modular network of bounded rationality, *Organizations with Incomplete Information*, Ed. M. Majumdar, Cambridge University Press, New York.
- Nisan, N. and I. Segal, 2005, The communication requirements of efficient allocations and supporting Lindahl prices”, *Journal of Economic Theory*.
- Osana, H., 1978, On the informational size of message spaces for resource allocation processes, *Journal of Economic Theory*, 17, 66.
- Osana, H., 2005, Externalities do not necessarily require larger message spaces for realizing Pareto-efficient allocations, *Review of Economic Design*, 19, 227.
- Radner, R., 1993, The organization of decentralized information processing, *Econometrica*, 62, 1109.
- Radner, R. and T. Van Zandt, 1995, Information processing in firms and returns to scale, *The Economics of Informational Decentralization: Complexity, Efficiency, and Stability*, Ed. J. Ledyard, Kluwer Publishing Company, Boston.
- Reichelstein, S., 1984, Incentive compatibility and informational requirements, *Journal of Economic Theory*, 32, 384.
- Reichelstein, S., and S. Reiter, 1988, Game forms with minimal message spaces”, *Econometrica*, 53, 661.
- Sah, K. and J. Stiglitz, 1986, The architecture of economic systems: hierarchies and polyarchies, *American Economic Review*, LXXVI, 716.
- Sah, K., and J. Stiglitz, 1988, Committees, hierarchies, and polyarchies, *Economic Journal*, 98, 451.
- Sato, F., 1981, On the informational size of message spaces for resource allocation processes in economies with public goods, *Journal of Economic Theory*, 24, 48.
- Saari, D., and C. Simon, 1978, Effective price mechanisms, *Econometrica*, 46, 1097.
- Segal, I, 2004, The communication requirements of social choice rules and supporting budget sets”, working paper.
- Serrano, R., 2004, The theory of implementation of social choice rules *SIAM Review*, 46, 377.
- Stoenescu, T. 2004, *Decentralized Resource Allocation Mechanisms in Networks*”, Ph. D. Dissertation, Department of EECS-Systems, University of Michigan.
- Tian, G., 1990, Completely feasible and continuous Nash implementation of the Lindahl correspondence with a message space of minimal dimension, *Journal of Economic Theory*, 51, 443.

- Tian, G., 2004, A unique informationally efficient allocation mechanism in economies with consumption externalities, *International Economic Review*, 45, 79.
- Van Zandt, T., 1999, Real-time decentralized information processing as a model of organizations with boundedly rational Agents, *Review of Economic Studies*, 66, 633..
- Van Zandt, T., 1998, Organizations with an endogenous number of information processing agents, *Organizations with Incomplete Information*, ED. M. Majumdar, Cambridge University Press, New York.
- Van Zandt, T., 2003a, Real-time hierarchical resource allocation, discussion paper.
- Van Zandt, T., 2003b, Real-time hierarchical resource allocation with quadratic payoffs, discussion paper.
- Visser, B., 2000, Organizational communication structure and performance, *Journal of Economic Behavior and Organization*, 42, 231.
- Walker, M., 1997, On the informational size of message spaces, *Journal of Economic Theory*, 15, 366.
- Walsh, W., M. Yokoo, K. Hirayama, and M. Wellman, 2003, On market-inspired approaches to propositional satisfiability”, *Artificial Intelligence*, 144, 125.
- Walsh, W. and M. Wellman, 2003, Decentralized supply chain formation: a market protocol and competitive equilibrium analysis, *Journal of Artificial Intelligence Research*, 19, 513.
- Williams, S., forthcoming, *Communication in Mechanism Design: a Differential Approach*, Cambridge University Press, New York.
- Ziv, A. 2000, Information technology and optimal firm structure, *Journal of Accounting Research*, 38, 297.