

This PDF is a selection from an out-of-print volume from the National Bureau of Economic Research

Volume Title: The Economics of Information and Uncertainty

Volume Author/Editor: John McCall, ed.

Volume Publisher: UMI

Volume ISBN: 0-226-55559-3

Volume URL: <http://www.nber.org/books/mcca82-1>

Publication Date: 1982

Chapter Title: Statistical Decision Theory Requiring Incentives for Information Transfer

Chapter Author: Jerry R. Green

Chapter URL: <http://www.nber.org/chapters/c4433>

Chapter pages in book: (p. 77 - 106)

3 Statistical Decision Theory Requiring Incentives for Information Transfer

Jerry R. Green

A model of decision making under uncertainty is presented in which one agent receives information and transmits it to another who makes a decision that affects them both. Because their utilities differ, the former will not necessarily transmit the observation accurately to the latter.

We study the problem faced by the decision maker. He must balance the potentially conflicting goals of efficient actions versus accurate information transmission. Some characteristics of optimal action plans are derived.

We also study the effect of improving the information structure on the value of the problem. Better information may be harmful, in general; but a sufficient condition for it to be beneficial is that there are only two possible observations.

3.1 Introduction

In organizations, large and small, the locus of decision making is often separate from that of information gathering. When all agents within such a system have common values, there is no incentive for the information gatherers to distort their observations. But this is often not the case. The design of the decision making procedures for the organization must balance the feasibility of eliciting accurate information transmission against the efficiency of the decisions taken.

In this paper a simple model of this type is presented.¹ Optimal decision rules are analyzed, and properties of these rules are derived in some

Jerry R. Green is professor of economics at Harvard University.

This research was conducted while the author was an Overseas Fellow of Churchill College, Cambridge. It was supported by NSF Grants SOC 71-03803 and APR 77-06999 to Harvard University.

special cases. We also study the impact of improving the information structure used by the information gatherer on the decision maker's welfare.

The following results are demonstrated:

1. If there are only two possible actions and two possible observations, either a first-best can be attained or else there is no value to the information at all.
2. If there are more than two possible actions, it may be the case that a randomized decision rule is required to attain the optimum. Moreover, such a rule may place a positive probability on selecting actions that are dominated.
3. If there are more than two observations, the optimum may involve a situation where the decision maker knows that the information gatherer is transmitting the observations imperfectly. Thus, information may be of some value, but not as much as it would have been had it been received by the decision maker directly.
4. The value of improved information may be negative.
5. However, if we are comparing information structures both of which involve only two possible observations (but an arbitrary number of states and actions), then the criterion of "more informativeness" used in statistical decision theory always agrees with the ordering relevant to the decentralized problem we have posed.

There are three strands of literature with a good deal of similarity to the model studied herein: team theory, principal-agent problems, and incentive compatibility (with preferences unknown). Before proceeding to the main part of the paper, a brief comparison of our model with each of these is probably worthwhile.

In team theory the actions of each individual are taken separately, perhaps after some communication,² whereas in this model there is a central decision maker who chooses a single action affecting both his welfare and that of the information gatherer. Those papers in team theory that do address issues of partially conflicting goals have retained this feature of decentralized actions. They also have tended to allow monetary transfers that, in conjunction with the team's decision rules, affect the members' incentives. Nevertheless, at least in a formal sense, this model lies close to team theory in spirit.

Principal-agent problems share the feature of partially conflicting goals.³ However, the locus of information gathering and action is coincident rather than separate. Their interesting features arise, rather, from the imperfect way in which risks can be shared and in the conflict between such risk sharing and the appropriate motivation of the agent. In the present model there is no explicit risk sharing possible, although considerations of this sort do indirectly affect the dependence of optimal actions on the observations transmitted.

Incentive compatibility has primarily concentrated on the imperfect, or more explicitly, dispersed, knowledge about the preferences of the individuals composing the economy.⁴ In this paper such information is perfect. The element of gamesmanship that is present is that the decision maker must bind himself in advance to a course of action that depends on the information transmitted to him. The information gatherer will optimize his transmission against this background—and the decision maker must take this into account *ex ante*. By using this “Stakelberg” format, the problem becomes trivial from a game-theory point of view. This approach has been chosen for its simplicity and also because we believe it more closely characterizes the reality of such situations than do other possibilities.

There is a further difference between most of the papers in incentive compatibility and this one. These papers concentrate on the possibility or impossibility of attaining true Pareto optima, that is, allocations compatible with feasibility in the physical sense but unconstrained by informational imperfections. When Pareto optima are attainable, it is necessary to elicit all of the relevant information about the environment. Therefore, all attention in this literature has been centered on designing mechanisms where truth-telling, or its essential equivalents, has been required.⁵ In this paper the potential for full optimality is present only in some very special cases. The second-best nature of the optimum imposed by the necessity for incentives in the information transmission process is explicitly recognized. It is not the case, therefore, that we have imposed truth-telling as an absolute constraint. Rather, we let the problem itself dictate the (constrained) optimal accuracy of the information that is to be elicited.

The remainder of the paper is organized as follows: Section 3.2 describes the model and defines the constrained optimization of the decision maker. In section 3.3 several results concerning the nature of the constrained optimum are proved. Examples of various types of phenomena possible in such a second-best situation are given and discussed. Section 3.4 concerns the value of improving the information structure. A brief conclusion follows.

3.2 Description of the Model

We consider an individual who must choose from among a collection $A = \{a_1, \dots, a_K\}$ of *actions*. The result of these actions is uncertain at the date at which the choice must be made. It depends on which *state of nature* from among those in $\Theta = \{\theta_1, \dots, \theta_M\}$ will arise. We assume that this individual is a von Neumann–Morgenstern expected utility maximizer and that his *utility* is given by the entries in the $K \times M$ matrix

$$U = (u_{km}),$$

where u_{km} is the utility if a_k is chosen and θ_m occurs.

Before the action must be chosen, some information relevant to the prediction of the state of nature can be obtained. This is modeled by introducing a set of *observations* $Y = \{y_1, \dots, y_N\}$ and for each $\theta_m \in \Theta$ a probability vector $(\lambda_{m1}, \dots, \lambda_{mN})$, where λ_{mn} is to be interpreted as the conditional probability of observing y_n given that the true state is θ_m . One writes the $M \times N$ *likelihood matrix*

$$\Lambda = (\lambda_{mn}).$$

An *information structure* is described completely by its likelihood matrix (the set Y being relevant only through the number of points it contains, which is given implicitly in the dimensionality of Λ). One nevertheless speaks of the information structure (Y, Λ) , or simply Λ , without fear of confusion.

The individual has a *prior probability* distribution $\pi = (\pi_1, \dots, \pi_M)$ over Θ , where π_m is the probability of θ_m .

In the absence of information the act a_m for which $U\pi$ has its maximal component would be selected.

If the individual could observe $y \in Y$ before choosing an act, he would compute the optimal act by forming his posterior probability according to Bayes's rule:

$$p(\theta_m | y_n) = \frac{\lambda_{mn} \pi_m}{\sum_{m'=1}^M \lambda_{m'n} \pi_{m'}}$$

Writing

$$P = (p(\theta_m | y_n))_{\substack{m=1, \dots, M \\ n=1, \dots, N}}$$

as the $M \times N$ *matrix of posteriors*, the matrix product UP determines his optimal action plan according to the location of the maximal elements within each column.

This is the standard problem of statistical decision theory. This paper addresses a problem of essentially the same nature. It differs in that the observation $y \in Y$ is not received directly by the individual who chooses actions but rather by someone else. To distinguish them clearly, we will call this other person the *agent* and the individual we have been discussing above, the *planner*.

The agent may have a different utility function from the planner. The agent's utility depends on the planner's action and on the realized state. Call his utility matrix $U' = (u'_{km})$.

The model works roughly as follows: The planner announces the way in which his action will depend on the observation transmitted to him.

The agent sees y_n and forms his posterior. Since we will assume that the agent and planner share the same prior distribution, this is just the n th column of the matrix P . The agent then transmits some element $y' \in Y$ which may or may not be the truth y_n . Thus, there are three factors influencing the choice of y' : the true value of the observation y_n , the agent's utility U' , and the courses of action to which the planner has committed himself.

We will be considering this system from the planner's point of view, under the assumption that both the information structure (Y, Λ) and the agent's utility (U') are known to him. It is clear that if $U = U'$, then there is no barrier to the policy of announcing the optimal action plan, which will induce the agent to transmit the true observation in every instance. But if $U \neq U'$, then announcing the optimal action plan may induce the agent to lie, and therefore may produce suboptimal results.

Because of these considerations, it may be the case that an announced action plan that is stochastic is necessary to achieve the best results.⁶ We define the planner's *random action plan* denoted by an $N \times K$ Markov matrix

$$Z = (z_{nk}),$$

where z_{nk} is the probability that the planner will choose a_k given that the agent responds with y_n . It is in the spirit of the literature on incentive compatibility and optimal taxation that the planner actually will make his choices according to Z rather than choose either the unconstrained best action at that stage or the best action from among those k' for which $z_{nk'} > 0$ (and which he therefore could claim to have produced as the result of the promised randomization).

The responses of the agent are thus determined by the maximal elements from within each column of the $N \times N$ matrix:

$$ZU'P.$$

We will define the *response pattern* R induced by Z as follows: Let R be an $N \times N$ matrix in which every column is a probability vector concentrated on the maximal elements of $ZU'P$ within that column. Since U' and P are fixed and are known to the planner, we are interested in the correspondence that sends Z into the set of all matrices R generated in this way. Letting \mathcal{M} and \mathcal{M}^T be the set of all Markov matrices and their transpositions, then

$$\phi: \mathcal{M} \rightarrow \mathcal{M}^T,$$

$$\phi(Z) = \{R \mid R \text{ is a response pattern induced by } Z\},$$

is this correspondence.

Typically, the maximal elements in all of the columns of $ZU'P$ will be unique, in which case of course $\phi(Z)$ consists of a single matrix of zeros and ones.

The planner tries to choose Z to his own best advantage. If Z induces the response pattern $R \in \phi(Z)$, the actual actions that the planner will be taking as they depend on the true observation of the agent are given by $R^T Z$, where R^T denotes the transpose of R .

We can now calculate the planner's expected utility attained when he announces the action plan Z . Let the probability vector $q = (q_1, \dots, q_N)^T$ be the probabilities of observing the various points in Y when the information structure is (Y, Λ) and the prior is π . We have $q = \Lambda^T \pi$. Let the symbol $\hat{\cdot}$ over a vector denote the diagonal matrix with that vector on the diagonal. The planner's expected utility when R is the agent's response pattern is given by

$$\text{trace } R^T Z U P \hat{q},$$

or, admitting the possibility that the agent will choose the best response pattern from among those to which he is indifferent, we have

$$(1) \quad \max_{R \in \phi(Z)} \text{trace } R^T Z U P \hat{q}.$$

The *planner's problem* is to maximize (1) over $Z \in \mathcal{M}$, knowing the correspondence ϕ .

The second-best nature of this problem is clear from the fact that a single decision maker would maximize $\text{trace } Z U P \hat{q}$, with Z unconstrained, but here, the range of the correspondence that sends Z into $\{Z' \in \mathcal{M} \mid Z' = R^T Z, \text{ for } R \in \phi(Z)\}$ may be much smaller than \mathcal{M} .

The fact that ϕ is an upper semicontinuous correspondence, but not a continuous one, means that the objective function (1) will not necessarily be continuous in Z . It will be upper semicontinuous (as a function), so a maximum will surely exist. In this section we show that it can be solved by converting it into a linear programming problem.

Nonstochastic response patterns give rise to mappings of the set of observations into itself. We will say that ρ is a *response rule* associated with a response pattern $R_\rho \in \mathcal{M}^T$ whenever

$$\begin{aligned} r_{\rho(y_n)y_n} &= 1, \\ r_{y_n'y_n} &= 0 \quad \text{for all } y_n' \neq \rho(y_n). \end{aligned}$$

Our search for the (second-best) optimum for the planner's problem involves a comparison of the optimal action plans inducing all the N^2 possible response rules. For each fixed ρ we are therefore limited to the action plans $Z \in \phi^{-1}(R_\rho)$. These are defined by a system of linear inequalities as follows: let $ZU'P = V$ be an $N \times N$ matrix, whose typical element is $v_{y_n'y_n}$. We need to have⁷

$$(2) \quad v_{\rho(y_n)y_n} \cong v_{y_n'y_n} \quad \text{for every } n, n'.$$

One can then express the planner's problem for each fixed ρ as

$$(3) \quad \max_{Z \in \mathcal{M}} \text{trace } R_\rho Z U P \hat{q}$$

subject to (2). This is a linear programming problem. The overall optimum is found by comparing the values of these problems over all possible response rules.

Although these remarks seem to imply that we will have to look at N^2 separate problems, we will now argue that *it suffices to solve the one for which ρ is fixed at the identity.*

Theorem 1

The optimum in the planner's problem can be found by solving

$$(4) \quad \max_{Z \in \mathcal{M}} \text{trace } Z U P \hat{q}$$

subject to

$$(5) \quad v_{y_n'y_n} \cong v_{y_n'y_n} \quad \text{for every } n, n',$$

where

$$Z U' P = V = (v_{y_n'y_n}).$$

Proof

Suppose the optimum of (3) subject to (2) exceeded the optimum of (4) subject to (5) for some response rule ρ other than the identity. Let the optimal action plan be Z . We will show that a trivial modification of Z leads to another action plan \hat{Z} for which (5) is satisfied and such that the value of (4) is the same as the value of (3).

This construction is done in a very simple way: Let $\hat{Z} = R_\rho^T Z$. The n th row of \hat{Z} is the $\rho(n$ th) row of Z . Since there are no rows of \hat{Z} which were not rows in Z , the effective choices open to the agent are, if anything, more limited. But since he chose the random action corresponding to the $\rho(n$ th) row of Z in response to y_n , he will now choose (or at least be indifferent to the choice of) the n th row of \hat{Z} . This proves that (5) is satisfied.

Since the random actions chosen in response to the observations are $R^T Z$ under Z and R as well as under \hat{Z} and the truthful response rule, (4) and (3) have the same value. QED.

It will often be the case that the solution to (4) subject to (5) will involve a Z matrix with several rows identical. That means that "truthful" responses are elicited in a situation with fixed alternatives. The essential response pattern is distorted by the necessity to use the same actions in several cases—which means that these observations are really not being

distinguished in any genuine sense. Nevertheless, theorem 1 is a useful computational tool since it reduces an apparently complex situation to a single linear program.

3.3 Optima in the Planner's Problem

In order to develop some further intuition about the nature of this problem we will now examine several special cases.

The first develops the general proposition that with two actions and two possible observations the problem of the planner exhibits a type of degeneracy.

Theorem 2

If $K = N = 2$, the planner's problem has an optimum of one of the following two forms: (i) A first best is attainable. (ii) The agent's information is valueless in that the action plan is independent of the observation transmitted.

Proof

Let $\bar{U}' = U'P$ be any 2×2 matrix. Consider $ZU'P$ given by

$$\begin{pmatrix} z_1 & 1-z_1 \\ z_2 & 1-z_2 \end{pmatrix} \begin{pmatrix} \bar{u}'_{11} & \bar{u}'_{12} \\ \bar{u}'_{21} & \bar{u}'_{22} \end{pmatrix} = \begin{pmatrix} z_1(\bar{u}'_{11} - \bar{u}'_{21}) + \bar{u}'_{21} & z_1(\bar{u}'_{12} - \bar{u}'_{22}) + \bar{u}'_{22} \\ z_2(\bar{u}'_{11} - \bar{u}'_{21}) + \bar{u}'_{21} & z_2(\bar{u}'_{12} - \bar{u}'_{22}) + \bar{u}'_{22} \end{pmatrix}.$$

If $\phi(Z)$ contains the truthful response pattern $R_p = I$, then we must have

$$(6) \quad \begin{aligned} z_1(\bar{u}'_{11} - \bar{u}'_{21}) &\geq z_2(\bar{u}'_{11} - \bar{u}'_{21}), \\ z_2(\bar{u}'_{12} - \bar{u}'_{22}) &\geq z_1(\bar{u}'_{12} - \bar{u}'_{22}), \quad z_1, z_2 \in [0, 1]. \end{aligned}$$

There are obviously two cases according to the sign of $(\bar{u}'_{11} - \bar{u}'_{21})(\bar{u}'_{12} - \bar{u}'_{22})$. If this is positive, the agent prefers one of the acts to the other independent of his observation; in that case, clearly, he will always respond with whichever element of Y associates a higher probability to the act he prefers. Only the action plan inducing his indifference, $z_1 = z_2$, remains within the truth-telling domain. The planner has thus chosen to act independently of the information reported, so that the best action is simply

$$\arg \max_{a \in A} U\pi.$$

In the more interesting case of $(\bar{u}'_{11} - \bar{u}'_{21})(\bar{u}'_{12} - \bar{u}'_{22}) < 0$ we have from (6) that either $z_1 \geq z_2$ (when $\bar{u}'_{11} - \bar{u}'_{21} > 0$) or $z_2 \geq z_1$ (when $\bar{u}'_{12} - \bar{u}'_{22} > 0$).

Consider the problem that the planner would face if he had perfect information:

$$\max_{Z^* \in \mathcal{M}} \text{trace } Z^*UP\hat{q}.$$

If the information is at all relevant to the planner, the nature of this optimum involves either $z_1^* = 1, z_2^* = 0$, or $z_2^* = 1, z_1^* = 0$. Therefore, when the truth-telling constraint is $z_1 \geq z_2$, either the first-best can be implemented (when $z_1^* = 1, z_2^* = 0$) or else the best policy is to set $z_1 = z_2 = 1$ if

$$\arg \max_{a \in A} U\pi = a_1$$

or $z_1 = z_2 = 0$ if

$$\arg \max_{a \in A} U\pi = a_2.$$

Other cases are completely symmetric. QED..

In the case of two states and two acts ($K = M = 2$, but N arbitrary), similar restrictions on the form of the optimum can be established. Here, however, it is the agent's optimum which might be implemented.

Theorem 3

If $K = M = 2$, the solution to the planner's problem must be one of the following two types: (i) a choice of action independent of the information transmitted; (ii) an action plan that implements the first-best solution for the agent.

Note that the second-best nature of the problem from the planner's point of view is in full force here. The agent's optimal action plan may be quite different from the planner's.

Proof

Let P be the posterior matrix in which the observations have been reordered (if necessary), so that

$$(7) \quad p_{11} > p_{12} \dots > p_{1N},$$

and thus

$$p_{21} < p_{22} \dots < p_{2N}, \text{ as } p_{1n} + p_{2n} = 1 \quad \text{for all } n.$$

Consider $u_{11} - u_{21}$ and $u_{12} - u_{22}$. If these are of the same sign, the optimal action plan is obviously of type (i). Therefore, we assume without loss of generality that $(u_{11} - u_{21}) > 0$ and $(u_{12} - u_{22}) < 0$.

If $u'_{11} - u'_{21}$ and $u'_{12} - u'_{22}$ have the same sign, then the agent will always respond with the same $y_n \in Y$, namely, that one for which z_{n1} is maximal (minimal) if they are both positive (negative). This would also lead to a solution of type (i).

Let us therefore take the case of $u'_{11} - u'_{21} > 0, u'_{12} - u'_{22} < 0$.

$$\text{Let } \mathcal{N}_P = \{n \mid u_{11}p_{1n} + u_{12}p_{2n} > u_{21}p_{1n} + u_{22}p_{2n}\},$$

$$\mathcal{N}_A = \{n \mid u'_{11}p_{1n} + u'_{12}p_{2n} > u'_{21}p_{1n} + u'_{22}p_{2n}\}.$$

These are the sets of observations that would lead the planner and the agent, respectively, to choose a_1 if they were to have perfect information about that observation.

Because of (7) we see that \mathcal{N}_p and \mathcal{N}_A can be described by two "cutoff" indices n^*_p and n^*_A :

$$\mathcal{N}_p = \{n \mid 1 \leq n \leq n^*_p\},$$

$$\mathcal{N}_A = \{n \mid 1 \leq n \leq n^*_A\}.$$

Again, we can assume without loss of generality that $n^*_p < n^*_A$, for if $n^*_p = n^*_A$, then the first-best for the planner and the agent coincide and this action plan can be implemented with truthful responses.

We know that the first-best action plan for the agent, where $z_{n1} = 1$ if $n \in \mathcal{N}_A$, and $z_{n1} = 0$ otherwise, induces honest responses and can thus be implemented. Suppose that another action plan Z^* is the true optimum in the planner's problem. Let n_{\max} and n_{\min} be indices of the observations such that

$$z^*_{n_{\max}1} \geq z^*_{n1} \geq z^*_{n_{\min}1} \quad \text{for all } n.$$

The agent's response pattern will then be of the form

$$\rho(y_n) = y_{n_{\max}} \quad \text{if } n \in \mathcal{N}_A,$$

$$\rho(y_n) = y_{n_{\min}} \quad \text{if } n \notin \mathcal{N}_A.$$

Therefore, the planner's expected utility level is

$$\begin{aligned} & [u_{11}z^*_{n_{\max}1} + u_{21}(1 - z^*_{n_{\max}1})] \left(\sum_{n \in \mathcal{N}_A} p_{1n}q_n \right) \\ & + [u_{12}z^*_{n_{\max}1} + u_{22}(1 - z^*_{n_{\max}1})] \left(\sum_{n \in \mathcal{N}_A} p_{2n}q_n \right) \\ & + [u_{11}z^*_{n_{\min}1} + u_{21}(1 - z^*_{n_{\min}1})] \left(\sum_{n \notin \mathcal{N}_A} p_{1n}q_n \right) \\ & + [u_{12}z^*_{n_{\min}1} + u_{22}(1 - z^*_{n_{\min}1})] \left(\sum_{n \notin \mathcal{N}_A} p_{2n}q_n \right). \end{aligned}$$

Because $n \notin \mathcal{N}_A \rightarrow n \notin \mathcal{N}_p$, the planner can increase his expected utility by decreasing $z^*_{n_{\min}1}$ to zero. Thus, $z^*_{n_{\min}1} = 0$.

Suppose that $z^*_{n_{\max}1} < 1$. Let us take the difference between the optimal expected utility of the planner and that obtained at the agent's first-best, which is implemented by $z^*_{n_{\max}1} = 1$:

$$\begin{aligned} (8) \quad & u_{11}(z^*_{n_{\max}1} - 1) + u_{21}(-z^*_{n_{\max}1}) \left(\sum_{n \in \mathcal{N}_A} p_{1n}q_n \right) \\ & + u_{12}(z^*_{n_{\max}1} - 1) + u_{22}(-z^*_{n_{\max}1}) \left(\sum_{n \in \mathcal{N}_A} p_{2n}q_n \right). \end{aligned}$$

Note that the derivative of (8) in $z^*_{n_{\max 1}}$ is

$$(9) \quad (u_{11} - u_{21}) \left(\sum_{N_A} p_{1n} q_n \right) + (u_{12} - u_{22}) \left(\sum_{N_A} (1 - p_{1n}) q_n \right).$$

Thus, if (9) is positive, then $z^*_{n_{\max 1}}$ should be increased to one, implementing the agent's first-best. If (9) is negative, then $z^*_{n_{\max 1}}$ should be zero. It follows that $z^*_{n_1} = 0$ for all n and that the solution is of type (i).

Finally, we must take the case where the planner and the agent have diametrically opposed goals: $u'_{11} - u'_{21} < 0$ and $u'_{12} - u'_{22} > 0$, in addition to our maintained hypotheses $u_{11} - u_{21} > 0$, $u_{12} - u_{22} < 0$.

The set N_A can then be characterized as

$$N_A = \{n | N \geq n \geq n^{**}_A\}.$$

There are two possibilities according to whether or not $n^{**}_A \leq n^*_p$. If so, then, the decision to choose a_1 independent of the observation transmitted dominates anything that can be sustained by a nontrivial response function (i.e., anything obtained by $z^*_{n_{\max 1}} > z^*_{n_{\min 1}}$). If not, then a_2 dominates. QED.

The simple types of optima observed when either $K = M = 2$ or $K = N = 2$ do not persist in more complex cases. There are essentially two phenomena at issue. They highlight the second-best nature of the problem we have posed and contrast it with those that have been studied in the earlier literature on incentive compatibility.

The planner's expected utility would improve if he could obtain more accurate answers. But to do so he may have to commit himself to an action plan that deviates from his own (unconstrained) choice, given the information that he will eventually possess. A potential tension is thus established between the nonoptimality of actions and the inaccuracy of responses.

When there are more than two actions, some positive probability may even have to be placed on an action that is *dominated* from the planner's point of view, in order to induce truthful responses by the agent. The cost of this nonoptimality may be outweighed by the value of the information so obtained. This is the content of the next example.

When there are more than two possible observations, the information may be of some value but may be imperfectly elicited in that truth-telling is obtained in an action plan Z with some identical rows. Observations corresponding to identical rows of Z are being lumped together in the information transmission process for all practical purposes. Note that, by theorem 3, such a situation must involve either more than two acts or more than two states; otherwise, either the information would be valueless or a genuine truth-telling situation would exist.

Example 1

Let $N = 2$, $M = 2$, $K = 3$, and

$$P = \begin{pmatrix} 3/4 & 1/4 \\ 1/4 & 3/4 \end{pmatrix} \quad q = \begin{pmatrix} 1/2 \\ 1/2 \end{pmatrix},$$

$$U' = \begin{pmatrix} 0 & -1 \\ 2 & 0 \\ 0 & -4 \end{pmatrix} \quad U = \begin{pmatrix} -2 & 3 \\ 0 & 0 \\ -1 & -1 \end{pmatrix}.$$

Let

$$Z = \begin{pmatrix} z_{11} & z_{12} & 1 - z_{11} - z_{12} \\ z_{21} & z_{22} & 1 - z_{21} - z_{22} \end{pmatrix},$$

incorporating the Markovian nature of Z in the notation.

To induce truth-telling we must choose Z so that the matrix ZUP has the larger element of each column on the diagonal. Writing

$$ZUP = \begin{bmatrix} (3/4)z_{11} + (10/4)z_{12} - 1 & (9/4)z_{11} + (14/4)z_{12} - 3 \\ (3/4)z_{21} + (10/4)z_{22} - 1 & (9/4)z_{21} + (14/4)z_{22} - 3 \end{bmatrix},$$

we see that the necessary and sufficient conditions for truth-telling are

$$(10) \quad (10/3)(z_{22} - z_{12}) < (z_{11} - z_{21}) < (14/9)(z_{22} - z_{12}).$$

In addition, we require

$$(11) \quad z_{ij} \geq 0 \quad \text{for } i, j = 1, 2, \quad z_{11} + z_{12} \leq 1, \quad z_{21} + z_{22} \leq 1.$$

Note that from the planner's point of view, act a_3 is dominated by random acts assigning weight to a_1 and a_2 with at least half the probability on the latter. In an unconstrained statistical decision problem, a_3 would never be selected. But we will now observe that a_3 is part of the optimal action plan because of its role in inducing the agent to tell the truth.

Following the solution method outlined in section 3.2, we compute the optimum assuming the planner is constrained to choose $Z \in \phi^{-1}(I)$; that is,

$$\max 1/2 \text{ trace } ZUP$$

subject to (10) and (11). This optimum is given by

$$Z = \begin{pmatrix} 0 & .643 & .357 \\ 1 & 0 & 0 \end{pmatrix},$$

and the value of the objective function⁹ is 0.696.

Note that this value is in between what can be achieved in the first-best ($z_{21} = z_{12} = 1$ and all other z 's zero), which gives an expected utility of 0.875, and what can be achieved independent of information ($z_{11} = z_{21} = 1$ and all other z 's zero), which yields 0.5.

Example 2

In this example we show how a situation of truth-telling can actually transmit information imperfectly in that the optimal Z matrix has some, but not all, rows identical. By theorem 1 we will obtain "truth-telling" as a response rule, but the planner will not be able to utilize the distinctions between certain observations.

Let

$$U'P = \begin{pmatrix} 1 & 3/2 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix},$$

$$UP = \begin{pmatrix} 5 & -1 & -1 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix},$$

$$q = (1/3, 1/3, 1/3)^T.$$

Solving the problem by writing the single linear program (4) subject to (5) yields the solution

$$Z = \begin{pmatrix} 1 & 0 & 0 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \end{pmatrix}$$

in which the value to the planner is $4/3$. Note that observations y_1 and y_2 are not being distinguished here. The first-best is $5/3$, and the best unconditional action (a_1) would yield 1.

Although truthful responses are elicited, the agent is indifferent between this and the response pattern

$$R = \begin{pmatrix} 1 & 1 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 1 \end{pmatrix}.$$

3.4 The Value of Improving the Agent's Information

We now turn to the second theme of this paper: is it true that planners would necessarily want to improve the information structures of their agents? To investigate this issue it is necessary to have a criterion through which the superiority of one information structure over another can be asserted.

In statistical decision theory, the following definition of *more informativeness*, or superiority, is standard: (Y, Λ) is said to be more informative than (Y', Λ') if, for any U and any π , the maximal expected utility attainable with (Y, Λ) is at least as large as that attainable under (Y', Λ') .

Blackwell (1951) has given a criterion equivalent to the fact that (Y, Λ) is more informative than (Y', Λ') : there exists a Markov matrix B such that $\Lambda' = \Lambda B$.

The elements of B can be interpreted¹⁰ as if they were the conditional probabilities of $y' \in Y'$ given $y \in Y$. But actually the *joint* distribution of y and y' may not be like this, and in any event it is not specified. The theorem simply says that if one information structure is better than another, one can imagine that the poorer one is a garbled version of the better one; for this reason it is called quasi garbling.¹¹

When (Y, Λ) is more informative than (Y', Λ') and when a prior π is fixed so that $q = \Lambda^T \pi$ and $q' = \Lambda'^T \pi$ can be computed, there exists another matrix C , which relates the two posterior matrices according to¹²

$$(12) \quad \begin{aligned} \text{i) } & P' = PC, \\ \text{ii) } & q = Cq', \\ \text{iii) } & C \text{ is the transpose of a Markov matrix, } C \in \mathcal{M}^T. \end{aligned}$$

Moreover, the existence of such a matrix C implies that the matrix B whose elements are given by

$$(13) \quad b_{nn'} = \frac{c_{nn'} q'_{n'}}{q_n}$$

will satisfy $B \in \mathcal{M}$, $\Lambda' = \Lambda B$; hence, the existence of such a C is equivalent to the more informativeness criterion.

When the agent's information structure changes, there are two effects on the planner's problem. First, there is a new function ϕ describing the response rule. The observation sets Y and Y' may be entirely different. But even if Y and Y' have equal numbers of elements, the new information will generally alter the accuracy of the responses made for any action plan. Second, there is a direct beneficial impact on the utility attainable for the planner under any fixed response function. The trouble is, of course, that taking advantage of this superior information may entail using an action plan at which the requisite pattern of information transfer will not be forthcoming from the agent.

By presenting a counterexample, we now proceed to establish that an improvement in information may be harmful. The analysis of the counterexample will rely heavily on the properties of optimal action plans derived in theorem 3 for the case $K = M = 2$.

Example 3

Let $K = M = 2$ and

$$\pi = (1/2, 1/2),$$

$$P = \begin{pmatrix} 3/4 & 1/2 & 0 \\ 1/4 & 1/2 & 1 \end{pmatrix},$$

$$q = (1/2, 1/4, 1/4)^T,$$

$$P' = \begin{pmatrix} 3/4 & 1/4 \\ 1/4 & 3/4 \end{pmatrix},$$

$$q' = (1/2, 1/2)^T.$$

One can compute directly that the matrix C given by

$$C = \begin{pmatrix} 1 & 0 \\ 0 & 1/2 \\ 0 & 1/2 \end{pmatrix}$$

satisfies (12) and hence establishes that the information structure for which P is the posterior is superior to that for which P' is the posterior.

Take the utility matrices for the planner and the agent given by

$$U = \begin{pmatrix} 2 & 0 \\ 0 & 4 \end{pmatrix}$$

and

$$U' = \begin{pmatrix} 4 & 0 \\ 0 & 2 \end{pmatrix}.$$

By theorem 3, we know that the optimum for the planner's problem in either case is either a situation in which the information is not being used at all or one in which the agent is attaining his own first-best. The value of the planner's problem without information is the maximal element of $U\pi$, which is 2, and occurs when a_2 is chosen with certainty.

The agent's problem with the poorer information structure is optimized with the action plan

$$Z' = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix},$$

which is also the action plan for the planner's first-best and generates an expected utility of $9/4$ ($= \text{trace } Z'UP'q'$) for the planner. As this is above 2, the action plan Z' is the overall optimum for the information structure P' .

With the better information structure the action plan that implements the first-best for the agent is

$$Z = \begin{pmatrix} 1 & 0 \\ 1 & 0 \\ 0 & 1 \end{pmatrix}.$$

But this lowers the planner's expected utility to 2 ($= \text{trace } ZUPq$). Thus, the planner can get at most an expected utility of 2 under the better information structure, by theorem 3, and he will prefer to give the agent the worse information structure.

Despite the rather negative conclusion of the previous example there is one case in which it can be proved that superior information is necessarily beneficial. This is in the comparison of information structures each of which has only two possible observations. Such systems are called *binomial channels* in the information theory literature.

Theorem 4

The value of the planner's problem under (Y, Λ) is necessarily at least as high as that under (Y', Λ') if (Y, Λ) is more informative than (Y', Λ') and both are binomial channels.

This theorem will be proved in two steps. First, we will establish that if Z is an action plan compatible with truth-telling under (Y', Λ') , then Z will also elicit truth-telling under (Y, Λ) . Second, we prove that if an action plan Z that induces truth-telling is optimal under (Y', Λ') and continues to induce truth-telling under (Y, Λ) , then the planner's expected utility improves. Since the optimum is always attained by truth-telling, by theorem 1, this would suffice.

By rearranging the labeling of $y'_1, y'_2 \in Y'$ we can, without loss of generality, establish that the Markov matrix B (defined by (13)) has its maximal element in each row on the diagonal.

Lemma

If C and B are 2×2 matrices related by (13) and B has its maximal element in each row on the diagonal, then C satisfies the same property.

Proof

Straightforward.

Lemma

Let C have its largest element in each row on the diagonal. Let $ZU'PC$ have its maximal element in each column on the diagonal; then $ZU'P$ has the same property.

Proof

Denote $ZU'P = \bar{Z}$ and

$$C = \begin{pmatrix} c_{11} & c_{12} \\ 1 - c_{11} & 1 - c_{12} \end{pmatrix}.$$

We need to establish that

$$(14) \quad c_{11}(\bar{z}_{11} - \bar{z}_{12}) + \bar{z}_{12} > c_{11}(\bar{z}_{21} - \bar{z}_{22}) + \bar{z}_{22},$$

$$(15) \quad c_{12}(\bar{z}_{21} - \bar{z}_{22}) + \bar{z}_{22} > c_{12}(\bar{z}_{11} - \bar{z}_{12}) + \bar{z}_{12},$$

$$(16) \quad c_{11} > c_{12}$$

together imply

$$(17) \quad \bar{z}_{11} > \bar{z}_{21} \text{ and } \bar{z}_{22} > \bar{z}_{12}.$$

Another way of writing (14) and (15) is

$$(18) \quad c_{11}(\bar{z}_{11} - \bar{z}_{21}) > (1 - c_{11})(\bar{z}_{22} - \bar{z}_{12}),$$

$$(19) \quad c_{12}(\bar{z}_{11} - \bar{z}_{21}) < (1 - c_{12})(\bar{z}_{22} - \bar{z}_{12}).$$

So we want to show that (16), (18), and (19) imply (17). From (18) and (19) and because $c_{11}, c_{12} \in [0, 1]$ we have the implications

$$\bar{z}_{22} - \bar{z}_{12} > 0 \text{ implies } \bar{z}_{11} - \bar{z}_{21} > 0,$$

$$\bar{z}_{11} - \bar{z}_{21} > 0 \text{ implies } \bar{z}_{22} - \bar{z}_{12} > 0.$$

Therefore, $\bar{z}_{11} - \bar{z}_{21}$ and $\bar{z}_{22} - \bar{z}_{12}$ must have the same sign. If they are both negative, then we have

$$\frac{c_{11}}{1 - c_{11}} < \frac{\bar{z}_{22} - \bar{z}_{12}}{\bar{z}_{11} - \bar{z}_{21}} < \frac{c_{12}}{1 - c_{12}},$$

contradicting (16). Therefore, (14), (15), and (16) imply $\bar{z}_{11} - \bar{z}_{21}$ and $\bar{z}_{22} - \bar{z}_{12}$ are both positive, which is (17). QED.

Recall that $q = Cq'$ is the distribution of the $y \in Y$ when q' is the distribution of $y' \in Y$. Since Z elicits truth-telling behavior under either information structure, we want to prove

Lemma

$$\text{Trace } \bar{Z}C\hat{q}' < \text{trace } \bar{Z}(C\hat{q}').$$

Proof

Straightforward computations lead to

$$\text{trace } \bar{Z}C\hat{q}' = (\bar{z}_{11}c_{11} + \bar{z}_{12}(1 - c_{11}))q'_1 + (\bar{z}_{21}c_{12} + \bar{z}_{22}(1 - c_{12}))q'_2$$

and

$$\text{trace } \bar{Z}(C\hat{q}') = \bar{z}_{11}(c_{11}q'_1 + c_{12}q'_2) + \bar{z}_{22}((1 - c_{11})q'_1 + (1 - c_{12})q'_2).$$

Therefore, we want to show that

$$(20) \quad (\bar{z}_{22} - \bar{z}_{12})(1 - c_{11})q'_1 + (\bar{z}_{11} - \bar{z}_{21})c_{12}q'_2 \geq 0.$$

By the previous lemma, $\bar{z}_{22} - \bar{z}_{12} > 0$ and $(\bar{z}_{11} - \bar{z}_{21}) > 0$. Thus, (20) follows from the nonnegativity of C . QED.

Combining these three lemmas, we see that the action plan Z satisfies the truth-telling constraints for the better information structure and induces a higher expected utility, which proves theorem 4.

3.5 Conclusions

A simple model of decision making in a two-person organization has been presented that attempts to capture the separation between information gathering and action. The problem has been shown to be equivalent to a linear program. Special situations in which a first-best could be attained have been analyzed.

In general, the optimal action plan may be random and may place some weight on dominated strategies. It may elicit truthful responses, but not in a way that allows effective use of this information.

Then we studied the behavior of the value of the problem when there is an improvement in the agent's information structure. This value may respond perversely in general. However, when the information structures being compared involve only two possible observations, the two criteria must agree.

Notes

1. We do not present any explicit models in which the structure developed here could be applied. Some possible examples are stockbroker-client or physician-patient relations; relations between superiors and subordinates within a firm, or between central management of a firm and that of a subdivision. More general models in which there are several players but emphasis on a separation between decision making and information gathering will be studied in later work. This would extend the range of applicability to situations in which there is some gamesmanship required among the information gatherers, as, for example, among several divisions of the same firm which are to be allocated resources for investment.

2. The basic reference is Marschak and Radner (1972). The problem of incentives in team theory has been treated by Groves (1973), retaining the feature of decentralized actions in the hands of the direct receivers of information.

3. See Ross (1973), Wilson (1969), Shavell (1979), and Holmström (1979).

4. A partial summary of the work in this area is given in Green and Laffont (1979). Dasgupta, Hammond, and Maskin (1979) have been particularly concerned with the Nash and Stakelberg equilibrium approaches. Gibbard (1977) and Zeckhauser (1969) use random decision devices, but in a different spirit than we will do here.

Most closely related to the present paper is Myerson (1979). Although formally concerned with unknown preferences, his analysis is similar in spirit to ours. He allows for random decision. He shows that using the space of possible utilities as a strategy space is adequate for the implementation of constrained optima as Bayesian equilibria (Nash equilibria in the game of incomplete information). However, the example he presents has only two possible utilities, and therefore cannot display some of the phenomena we look at in section 3.3 of this paper.

Rosenthal (1978) proves that full information Pareto optimality is generally unattainable in such games, and he provides a sufficient condition under which a Bayesian equilibrium can implement such optima.

5. Groves and Ledyard (1977) have a strategy space much smaller than the space of unknown parameters, but the domain over which its equilibria exist is hard to delineate. Myerson (1979) and Rosenthal (1978) impose truth-telling as a constraint in their definition of equilibrium. The actual constrained Pareto optima we obtain in this paper do involve

truth-telling. However, it is not imposed in the solution concept, but rather emerges as a characteristic of the second-best. See theorem 1.

6. See example 1.

7. We will always follow the practice of writing the constraints on responses of the agent as weak inequalities, expressing the fact that if he is indifferent between response patterns, we can assume that he will make the indicated choice.

8. See example 2.

9. The true optimum differs slightly because of rounding error. This calculation was performed on an electronic computer.

10. See Marschak and Miyasawa (1968) or McGuire (1972).

11. Another advantage of having this theorem is that it provides a constructive method to check this relationship between information structures. The existence of such a matrix can be determined by solving a system of linear inequalities, whereas the definition requires, in principle, that a comparison be made over all U, π .

12. See Marschak and Miyasawa (1968, p. 154). Our equation (13) is their (8.23).

References

- Blackwell, D. 1951. Comparison of experiments. In J. Neyman, ed., *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability*, pp. 93–102. Berkeley and Los Angeles: University of California Press.
- Dasgupta, P.; Hammond, P.; and Maskin, E. 1979. The implementation of social choice rules. *Review of Economic Studies* 66 (no. 2): 185–216.
- Gibbard, A. 1977. Manipulation of schemes that mix voting with chance. *Econometrica* 45 (no. 3): 665–82.
- Green, J., and Laffont, J.-J. 1979. *Incentives in public decision-making*. Amsterdam: North-Holland.
- Groves, T. 1973. Incentives in teams. *Econometrica* 41 (no. 4): 617–32.
- Groves, T., and Ledyard, J. 1977. Optimal allocation of public goods: A solution to the 'free-rider' problem. *Econometrica* 45 (no. 4): 783–809.
- Holmström, B. 1979. Moral hazard and observability. *Bell Journal of Economics* 10 (no. 1): 74–91.
- McGuire, C. B. 1972. Comparison of information structures. In Radner, R., and McGuire, C. B., eds., *Decision and organization*, chap. 5. Amsterdam: North-Holland.
- Marschak, J., and Miyasawa, K. 1968. Economic comparability of information systems. *International Economic Review* 9 (no. 2): 137–74.
- Marschak, J., and Radner, R. 1972. *Economic theory of teams*. Cowles Foundation Monograph 22. New Haven: Yale University Press.
- Myerson, R. 1979. Incentive compatibility and the bargaining problem. *Econometrica* 47 (no. 1): 61–74.
- Rosenthal, R. 1978. Arbitration of two-party disputes under uncertainty. *Review of Economic Studies* 45 (no. 3): 595–604.

- Ross, S. 1973. The economic theory of agency: The principal's problem. *American Economic Review* 63 (no. 2): 134-39.
- Shavell, S. 1979. Risk sharing and incentives in the principal and agent relationship. *Bell Journal of Economics* 10 (no. 1): 55-73.
- Wilson, R. 1969. The structure of incentives for decentralization under uncertainty. In *La décision*, pp. 287-308. Paris: CNRS.
- Zeckhauser, R. 1969. Majority rule with lotteries on alternatives. *Quarterly Journal of Economics* 83: 696-703.

Comment Andrew Postlewaite

Green presents a model of decision making under uncertainty in which one agent receives information and transmits it to a principal. The principal then makes a decision (chooses an act) on the basis of the information received. Both the principal and the agent then receive a payoff which depends on both the state of the world and the action chosen by the principal. It is assumed that the principal must commit himself to an action plan which determines which action (possibly a probability distribution over several actions) he will take given any information signal transmitted by the agent. Thus, the principal is faced with a possible incentive problem. If the principal commits himself to taking the action which is best for him given any information signal, this may be harmful to the agent. In this case it could be in the agent's interest to transmit incorrect signals. Hence, the principal, recognizing the incentives he creates for the agent, assumes that the agent will behave in a purely self-interested manner and will transmit whatever (possibly false) information signal which according to the principal's action plan gives the agent his highest payoff.

It is assumed that both the principal and the agent know the probabilities of the states, the probabilities of the signals the agent could receive in each state, and the payoffs each will receive given the state and action taken by the principal. Thus, for any action plan the principal contemplates, he can precisely calculate what signal the agent will transmit given the information signal that the agent has observed (received). This, with the probabilities of the states and information signals, is sufficient for the principal to determine his own expected payoff for any action plan. He chooses the action plan which maximizes his expected payoff.

Through a series of examples we will examine the possible effects of improved information on both the principal and the agent. In each of the examples, the set of actions the principal can take are $A = \{a_1, a_2, a_3, a_4\}$

Andrew Postlewaite is associated with the Wharton School of the University of Pennsylvania, Philadelphia.

	Y
A	
1	
2	
3	
4	

and the states of the world are $Y = \{y_1, y_2, y_3, y_4\}$; the payoffs to the principal and the agent for each action and state of the world are given in the matrices in figure C3.1.

For all of the examples we assume probability $(y_i) = 1/4, i = 1, 2, 3, 4$. For each of the examples we will consider two information structures. In the first structure, W , the agent will observe s_1 if either y_1 or y_2 is the true state and s_2 if either y_3 or y_4 is the true state. In each example we will compare the information structure W with an improved structure I where the agent observes the state precisely; i.e., he observes s_i if y_i is the true state, $i = 1, 2, 3, 4$.

Before beginning the examples it will be helpful to write down the matrices of expected payoffs for the action and information signals for the worse information structure W . These are given in figure C3.2.

If action a_1 is taken when the signal s_1 is observed by the agent, the principal's payoff is either $2b$ or 0 depending upon whether y_1 or y_2 is the true state. Since each has probability $1/2$ given s_1 , the principal's expected payoff is b in this case. The other numbers for the principal and the agent are calculated in this manner.

The first case we will examine is that where the improved information structure I yields higher expected payoffs for both the principal and the agent.

Case I. Suppose

$$e > d > b > c > 0$$

and

$$j > h > f > g > 0.$$

Consider an action plan for the principal given as follows: $a_4 = f(s_1), a_3 = f(s_2)$; i.e., the principal takes a_4 when s_1 is transmitted and a_3 when s_2 is transmitted. Notice that this guarantees the agent (expected) payoff j

	Principal	Agent

A \ Y	1	2	3	4
1	2b	0	2c	0
2	0	2c	0	2b
3	2d	0	2e	0
4	0	2e	0	2d

A \ Y	1	2	3	4
1	0	2f	0	2g
2	2g	0	2f	0
3	2h	0	2j	0
4	0	2j	0	2h

Fig. C3.1

		Principal		Agent			
		S		S			
A	S	S ₁	S ₂	A	S	S ₁	S ₂
1		b	c	1		f	g
2		c	b	2		g	f
3		d	e	3		h	j
4		e	d	4		j	h

Fig. C3.2

regardless of the signal if he transmits the information correctly. Since we assumed j was larger than any payoff in his matrix for information structure W , it is in his best interests to transmit the signal correctly. This yields the principal a payoff of e , which is the best he can get for the structure; thus, this is the "solution" for the structure W .

Now consider the improved structure I and the action plan

$$a_3 = f(s_1) = f(s_3),$$

$$a_4 = f(s_2) = f(s_4).$$

This gives the agent h in states 1 and 4 and j in states 2 and 3 if he transmits the signals correctly. These by assumption are the largest payoffs he can get in the respective states. They similarly yield the principal either d or e , which are his highest possible payoffs given the states. Hence, this action plan is the solution for information structure I .

Note that here the expected payoffs for the principal and the agent are $d + e$ and $h + j$, respectively, as compared with e and j for information structure W .

The fact that the improved information structure I is better for both the principal and the agent in this example is not surprising. Given the assumed inequalities, the interests of the two agents coincide for the relevant actions a_3 and a_4 . The common interests yield "truth-telling" on the agent's part for that action plan which gives the principal his highest payoff for any observed signal. In this case we should expect improved information to yield higher payoffs for both principal and agent than the worse information structure.

The next example shows how the improved information structure can hurt both principal and agent.

Case II. Suppose

$$b > d + e > \frac{1}{2}(b + c) > c > 0 \quad d, e > 0$$

and

$$f > h + j > \frac{1}{2}(f + g) > g > 0 \quad h, j > 0.$$

For information structure W consider the action rule

$$a_1 = f(s_1),$$

$$a_2 = f(s_2),$$

The agent receives f if he reports the signal truthfully. Since, by assumption, f is the largest payoff possible, he will do so. Since this yields the principal b , his highest possible payoff, this action plan is optimal.

While in information structure W the agent's and principal's interests essentially coincide, this dramatically changes when we go to the improved information structure I . Here there is absolute conflict between the agent and principal for actions a_1 and a_2 : if one gets a positive payoff, the other receives nothing. Clearly the information the agent transmits is worthless to the principal if he takes only actions a_1 and a_2 . Any use he makes of the signals the agent sends can be exploited by the agent. The best the principal can do if he is to use only actions a_1 and a_2 is to ignore the signals received, e.g., use an action plan such as

$$a_i = f(s_i) \quad i = 1, 2, 3, 4$$

which would yield the principal $\frac{1}{2}(b + c)$. If, however, he chooses action plan

$$a_3 = \hat{f}(s_1) = \hat{f}(s_3),$$

$$a_4 = \hat{f}(s_2) = \hat{f}(s_4),$$

the agent will find it in his best interest to transmit the correct signals. This gives $d + e$ to the principal, which by assumption is better than he can do using only a_1 and a_2 . But this yields both the agent and the principal less than their respective payoffs under the worse information structure W .

So far, we have only shown that the action plan \hat{f} for information structure I leaves the players worse off than in the worse information structure W . But \hat{f} is not in fact the optimal information structure for I . Since the agent has a strict incentive to transmit signals truthfully (i.e., this yields him strictly higher payoff than any other response pattern), the principal could use a stochastic action plan which would put a small positive probability on a_1 and a_2 for the appropriate signals without destroying the agent's incentives. This would leave the principal slightly better off and the agent slightly worse off than under \hat{f} . The probability that can be placed on a_1 and a_2 without destroying the incentives (and thus making the information worthless here) becomes arbitrarily small as h gets arbitrarily close to j . Thus, we see that if the assumed inequalities hold and also if h is sufficiently close to j , the improved information

structure I leaves both agent and principal worse off than they are in structure W .

A few words about this example. Green makes the point that the principal's optimal action plan may involve positive probabilities on dominated actions. If $c > e$ and $g > j$, that is precisely what is happening here. Action a_1 dominates a_3 and a_2 dominates a_4 , yet the optimal action plan involves placing most of the probability on a_3 and a_4 . This example shows clearly how the problem of the agent's incentives makes this necessary.

A second aspect to point out is how the better information structure makes both players worse off. Actions a_1 and a_2 are the best actions to take from a joint point of view so long as states $\{y_1, y_2\}$ can be distinguished from $\{y_3, y_4\}$. The worse information structure allows an action plan to distinguish them because the interests of the two coincide. The better information structure gives enough information to put the players in conflict with each other as to which of the actions a_1 or a_2 is best. Given the framework of this model, both agents would theoretically agree to use information structure W rather than the improved structure I .

This raises a question which is essentially outside of this model. How is the information structure determined, and what is the "information about the information"? Although both principal and agent do better under W than I , this is because the optimal action plan under W picks only a_1 and a_2 . Given this fact, the agent has an incentive to secretly "get the better information" in I . Thus, even the availability of the better information structure I would prevent the principal from using only actions a_1 and a_2 unless he could ensure that the agent would not observe the more precise information signals in I . We will say a bit more below about the questions of who chooses the information structure and the degree of common knowledge about the information structure being used.

So far we have provided examples of an improved information structure either benefiting both principal and agent or hurting both. It is also possible that the improved information structure I can help the agent and hurt the principal.

Case III. Suppose

$$b > c > d > e > 0 \quad b > d + e$$

and

$$f > g > j > h > 0 \quad f < j + h.$$

This example is similar to Case II. For structure W , the optimal action rule is $a_1 = f(s_1)$, $a_2 = f(s_2)$. This induces truthful responses from the agent and yields payoffs of b and f to the principal and agent, respectively.

If the information structure is I and h is very close to j , the optimal action rule puts nearly all probability on a_3 and a_4 , yielding approximately $d + e$ to the principal and $j + h$ to the agent; from the assumed inequalities we see that the agent is better off and the principal worse off.

Each of the three cases above was explicitly or implicitly mentioned in Green's paper. The last case—that in which the agent is worse off and the principal better off under the improved information structure—was not. At first glance one might think this case particularly paradoxical. Cannot the agent simply "ignore" the better information? Yet as Case II illustrated, the structure of the model essentially prevents an agent from ignoring information. The principal commits himself to an action rule, and if he knows that information is available he can design his action rule so that an agent would be even worse off if he did ignore the information. It is this Stackelberg-like feature which has the principal but not the agent commit himself which gives rise to these phenomena.

In fact a simple transformation of Case III gives rise to an example where the agent is worse off and the principal better off under the improved information structure.

Case IV. Suppose

$$b > c > d > e > 0 \quad b < d + e$$

and

$$f > g > j > h > 0 \quad f > j + h.$$

An argument as in Case III shows that if h is close to j , the payoffs to the principal and agent are b and f , respectively, under W and approximately $d + e$ and $h + j$, respectively, under I . That the principal is better off and the agent worse off follows from the assumed inequalities.

As we mentioned above, the question of how the information structure is determined becomes important. Does the agent or the principal decide on the structure? Is the structure itself a matter of negotiation? In the examples in both Case II and Case IV, if the agent had the ability to commit himself to the worse information structure, he would *if* he could also convince the principal that he had done so.

It also should be pointed out that the results of the examples would be drastically changed even if we left aside the question of determination of the information structure, but only reversed the roles of leader and follower in the model. Consider a situation where the agent commits himself to a response rule and the principal then chooses an action. It is straightforward to construct examples where both principal and agent are better off with the roles of leader-follower reversed. In such cases we might expect the reversed roles "institutionalized." The model examined in the paper is but one of a class of similar models. Which model is

operative for a particular application would perhaps be endogenously determined in an expanded setting.

Comment John G. Riley

In the basic model of information acquisition as developed by Blackwell (1951) and Marschak and Miyasawa (1968), an individual may be viewed as paying a predetermined fee to an information service. The latter then delivers one of a set of messages which have the effect of updating the individual's beliefs as to the likelihood of the various states of nature. Swept into the background is the issue of how the purchaser of information (hereafter referred to as the "principal") verifies the accuracy of the information.

Even if the "agent" operating the information service simply collects and processes data, verification of computations, etc., is a costly process. However, more commonly the principal relies at least in part on the expertise of the agent. Verification is then only indirect. The principal either develops a long-term relationship with the agent or, alternatively, relies on the reputation the latter has developed.

Green's starting point is to consider what sort of contracts between principal and agent might be negotiated when such indirect verification is not available. For the problem to be interesting it is necessary that both the purchaser and gatherer of information have an interest in the final outcome. In general these interests will not coincide and, as a result, the agent may have an incentive to reveal false information. When such is the case, Green examines the potential benefits to the principal if he precommits himself, for every feasible message, to a particular action (or, more generally, to probability weights on different actions).

The first result is that the principal can do no better than induce truthful information transfer. Appropriately interpreted this is essentially definitional. Let there be n possible messages from agent to principal. The latter announces that if he receives message m_i , he will take probabilistic action $z_i = (z_{i1}, \dots, z_{in})$, where z_{ij} is the probability assigned to action j . Effectively the agent then freely chooses from the set of probabilistic actions $Z = \{z_1, z_2, \dots, z_n\}$. Thus, without loss of generality we may relabel the elements of this set so that z_i^* is the action chosen by the agent when the *truth* is message i .

Using this result, it is then shown that if there are only two states of nature and either two actions or two possible messages, the principal's optimal response is (i) to not make use of the information, or (ii) to accept

the message without any precommitment strategy. In more complex environments it is shown that probabilistic precommitments can increase the expected utility of the principal. An example is provided in which the principal finds it desirable to introduce an action that is strictly dominated by another action. To understand this result, I find figure C3.3 helpful. On the horizontal axis is the possibility p of state 1, and on the vertical axis expected utility of the principal, which we may write as

$$U(p; a_i) = pu_1(a_i) + (1-p)u_2(a_i),$$

where $u_s(a_i)$ is the utility of the principal in state s if he takes action a_i . Without the services of the information gathering agent the principal, with prior probability $p_0 = .5$, chooses action a_1 . Suppose that the agent receives one of the following two messages with equal probability:

$$m_1: p = .8, \quad m_2: p = .2.$$

If the agent always tells the truth, the principal has an expected utility of AB when $m = m_1$ and CD when $m = m_2$. Since each message is received with equal probability the ex ante expected utility is EG and the net value of the information is FG .




Fig. C3.3 The value to the principal of receiving information.

Suppose next that action a_2 is the dominant action from the viewpoint of the agent. Then rather than reveal the truth he will always pass on the message $m = m_1$. Obviously this is valueless to the principal. To combat this problem, the principal announces that if the message received is $m = m_1$ he will adopt action a_2 with probability z_{12} and action a_3 with probability z_{13} . The expected utility of the principal is then the weighted average

$$z_{12}U(p;a_2) + z_{13}U(p;a_3).$$

As long as action a_3 is sufficiently worse for the agent than a_1 when the truth is $m = m_2$, and better than a_1 when the truth is $m = m_1$, his optimal response is to reveal the truth. As a result the expected utility of the principal is AJ if $m = m_1$ and CD if $m = m_2$, and the information service has an expected value of FH .

In the final section of his paper Green asks whether a more informative agent is necessarily more valuable to the principal. Except in a very special case he finds that the opposite may be true: even after an optimal revision of his precommitment strategy the principal may be worse off. Again I find a diagram helpful in understanding this conclusion. The top part of figure C3.4 once again depicts the choices open to the principal for the case of two states and two actions. The bottom part illustrates the expected utility of the agent. With only the two messages $m = m_1$, $m = m_2$ the interests of agent and principal always coincide. However, suppose that message m_2 is really a "garbled" version of the following two messages;

$$m_{2\alpha}: p = 0, \quad m_{2\beta}: p = 0.4,$$

that is, the agent receives a message but is not sure whether it is $m_{2\alpha}$ or $m_{2\beta}$. If he assigns equal probabilities to the two alternatives, his information is summarized by the message m_2 . Clearly an agent who is able to unscramble such a message is sensibly described as being better informed.

As depicted in figure C3.4 revelation of the unscrambled truth is of no additional benefit to the principal since his optimal action is a_1 for $m = m_{2\alpha}$ and $m = m_{2\beta}$. However, the interests of principal and agent no longer always coincide. If $m = m_{2\beta}$, the optimal action from the viewpoint of the agent is a_2 . Therefore, the agent has an incentive to conceal the fact that he has unscrambled message m_2 and to reveal the false message m_1 . But suppose the principal were to learn that the agent had become better informed. With only two actions Green's theorem 3 tells us that the best the principal can do is either to ignore the agent entirely or to accept his message as if it were the truth. If he does the latter, he knows that when the agent announces m_1 there is one chance in three that $p = .4$ and two chances in three that $p = .8$. The principal then assigns a probability of $\frac{1}{3}(.4) + \frac{2}{3}(.8) = .66$ to state 1, and his expected utility falls




Fig. C3.4 More informative information leaving both principal and agent worse off.

from AB to AC . A similar argument for the bottom part of figure C3.4 establishes that the agent's expected utility rises from $A'B'$ to $A'C'$. Of course this is also the outcome when the principal is unaware that the agent has become better informed.

The other alternative open to the principal is to ignore the agent and use only his prior probability of state 1 (.5). Clearly this is preferred since it yields an expected utility of AD while lowering the expected utility of the agent to $A'D'$. Therefore, if the information contained in message m_2

becomes unscrambled, the expected utility of the principal drops from AB to AD and the expected utility of the agent drops from $A'B'$ to $A'D'$. This is another simple example of the phenomenon first noted by Andrew Postlewaite: additional information can hurt both parties. Indeed, in this example the coarser information benefits both principal and agent while the finer information is valueless!

I shall conclude with some more general remarks. The main inference that I draw from Green's paper is that we cannot expect to obtain more than a multitude of special cases unless the model is refined in some way. One line of attack would be to focus more closely on the way in which the agent is compensated by the principal. Implicitly compensation takes the form of a fixed fee paid in advance. However, by contracting to pay a fee contingent upon the eventual state of nature the principal might be able to eliminate, or at least reduce, the extent to which his interests conflict with those of the agent. A second approach would be to take explicit account of the repetitive nature of informational exchanges. In each of the applications suggested by the author (stockbroker-client, physician-patient, manager-worker) it is hardly surprising that the information purchaser is usually observed in a long-term relationship with the information gathering agent. If precommitment strategies of the sort described above have practical relevance, surely they are only likely in situations which are sufficiently repetitive to allow the principal to learn the interests of the agent and vice versa.

References

- Blackwell, D. 1951. Comparison of experiments. In Neyman, J., ed., *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability*, pp. 93-102. Berkeley and Los Angeles: University of California Press.
- Marschak, J., and Miyasawa, K. 1968. Economic comparability of information systems. *International Economic Review* 9, no. 2: 137-74.