in

Arrow and the Ascent of Modern Economic Theory

Edited by
George R. Feiwel
Alumni Distinguished Service Professor and Professor of Economics, University of Tennessee

NEW YORK UNIVERSITY PRESS
Washington Square, New York
1987

13 Incentive-based Decentralization: Expected-Externality Payments Induce Efficient Behaviour in Groups

John W. Pratt and Richard Zeckhauser*

1 INTRODUCTION

1.1 Endowment Contributed by Kenneth Arrow

Thirty-five years ago Kenneth Arrow asked a profound question: Is it 'formally possible to pass from a set of known individual tastes to a pattern of social decision-making, the procedure in question being required to satisfy certain natural conditions'? (Arrow, 1951, p. 2). He laid out an appealing set of conditions and demonstrated that the answer was 'No'. The vast literature that followed, frequently played musical chairs with his requirements while it tiptoed along the border of infeasibility.

A second major strand of Arrow's work on collective choice processes examines the circumstances under which the market can serve as an effective decentralised decision mechanism. Arrow and Debreu (1954) assume that individuals' preferences are unknown, hence the

* Pratt's research was supported by the Associates of the Harvard Business School. Zeckhauser's research was supported by NSF grant SOC77-16606 and by the Energy and Environmental Policy Research Center, Harvard University. Valuable comments were provided by Roy Radner, John Riley, and participants in the NSF-NBER Seminar on Decentralization, San Diego, in the 8th Management Science Colloquium of the Kansai Economic Research Institute, Osaka, where the research was originally presented, and subsequently, in various seminars at the University of Chicago, Harvard, MIT, and Stanford.

process must elicit the requisite information. However, they allow for a medium of exchange, for transferable utility. (A monetary measuring rod – quite appropriate in this context – violates the Independence of Irrelevant Alternatives Condition of Arrow's aforementioned impossibility result.) Arrow and Debreu and their successors demonstrated that if all relevant markets can be established, and if preferences and production sets are not too ill behaved, then efficient competitive outcomes will be generated by market processes (see also Arrow and Hahn, 1971).

In quite separate lines of important and still continuing inquiry, Arrow has addressed many of the major issues of welfare economics related to market performance in the presence of such imperfections as externalities and asymmetric information. (His work is surveyed in the introduction to this volume.)

1.2 Decentralization

These three bodies of work together illuminate a central question of microeconomics: 'Under what conditions and with what mechanisms can we allocate resources efficiently in our society?' A vast literature on topics ranging from business management to socialism has addressed this question. Within economics it has taken the form: 'When and how can local decision units be induced to provide information and to take actions that co-ordinate appropriately and are in the interest of a company, an organization, or society as a whole? In other words, how can we effectively decentralize decision-making responsibility?' This paper addresses that question allowing for externalities and for asymmetric information, quite possibly about preferences, but permitting monetary transfers among the participants.

Effective decentralization is a major concern for both public and private organizations. The government, for instance, may seek to induce the industrial firms in a particular geographic region to make appropriate trade-offs between environmental quality and production cost. A business firm with many subordinate decision-making units will try to develop mechanisms to ensure that the units undertake actions that support the profits of the firm as a whole. To maintain a firm-wide reputation for product quality, for example, some prodding from headquarters is needed to give each self-interested profit centre sufficient encouragement to maintain or raise its own quality level. Another

example: architects' fee schedules should be designed so that they have an appropriate interest in holding down construction costs.

If all agents had a common objective, it could be maximized by a team, whose members observe information and follow strategies formulated by a central authority (Marschak and Radner, 1972). Because interests are fully shared, team members have an automatic incentive to co-operate in achieving the team optimum, subject to whatever information-transmission constraints and costs the world imposes. The centre assesses the joint distribution of information the agents will observe and then specifies what messages each agent should send and what actions he should take depending upon the information he has, so as to maximize the joint payoff for the team. (Hurwicz, 1959, is a classic paper on the more general subject of informational efficiency in resource allocation.)

If objectives differ, there are still circumstances in which efficiency might be achieved, even without transfers. Under *laissez-faire*, self-interested maximizing decisions by each agent maximize total welfare unless there are interdependencies (externalities) whereby one agent's choice directly affects the welfare of another. (This holds even when agents have information useful to one another, providing they signal honestly when it costs them nothing to do so.) Alas, the market cannot always work its miracles: externalities do occur.

Despite interdependencies and differing objectives, a ruthless dictator could maximize total welfare if there were no uncertainty or if all information were eventually public or verifiable at negligible cost. Death to him who does not act to maximize total welfare. Equivalently, a central authority providing sufficient financial incentives could 'pin' the optimum.

But information may not be monitorable. The important literature of demand-revealing processes, starting with Vickrey's (1961) contribution, shows that interdependencies can be handled even in some situations with asymmetric information. The primary question it addresses is how can we simultaneously: elicit information from individuals about their valuations of goods, and use that information to allocate resources efficiently?

The procedures rely on financial incentives. The now-recognised guiding principle (see Tideman and Tullock, 1976), is to charge each individual the cost to all others of his choice, or pay him the benefit. Thus, in a Vickrey auction, the high bidder pays the valuation of the person who is denied. In a binary public choice, any individual who changes the decision, say from 'build' to 'not build', must pay the net valuation of all other individuals in favour of the opposite course. Groves (1973) employs and extends this approach, showing its applicability to a class of situations in which individuals signal to a centre, the centre replies, and then all choose their actions. These mechanisms make truthfulness a dominant strategy for each participant by paying or charging him for his revealed externality, that is, the actual net benefits or costs his action conveys to the other participants as implied by their actions and/or statements.

Besides the limited range of problems handled, a significant disadvantage of demand-revealing processes is that the incentive payments ordinarily create a deficit or surplus that must be absorbed by some external party. Is there a way to balance the incentive-payment budget? In many situations only the procedure just described makes truthfulness a dominant strategy (see Green and Laffont, 1979, and references therein). A weaker type of equilibrium must be accepted. And it can be implemented in some circumstances, given that if knowledge is probabilistic then incentives need only reward truthful revelation on an expected-value basis.

1.3 Arrow's Related Contribution

Once again it is appropriate to start with the work of our honoree. The work most closely related to ours is Arrow (1977, 1979) and d'Aspremont and Gérard-Varet (1975, 1979).[1]

Arrow includes a discussion placing his problem in the context of recent developments in economic theory:

> I wish to suggest a different approach to demand revelation, which achieves efficiency and avoids the waste of resources in the incentive payments. As might be expected, it makes stronger assumptions, in this case, assumptions about the expectations that each agent has about each other's valuations. (Arrow 1979, p. 27).

Arrow formalizes the collective decision problem as a game of incomplete information with one stage of revelation, a collective decision rule, and incentives induced by transfer payments. He assumes no income effects, specifically, utility functions linear in income whose

expectations the players seek to maximize. He uses an equilibrium concept in extensive form, which he observes precludes reduction to normal (strategy) form because 'the second player is optimizing given all previous history', making the equilibrium perfect. He describes perfection as 'the game-theoretic counterpart of the principle of optimality in dynamic programming' (Arrow, 1979, p. 29; see note 2 for further discussion). Except for this sequential refinement, the solution concept is essentially Nash equilibrium, but it is often called Bayesian as well, since it relies on expected values defined by a common original prior distribution.

For a collective decision, assuming agents' information is one-dimensional and continuous, Arrow derives the first-order conditions that the transfers must satisfy locally to make truth-telling an equilibrium, and second-order conditions for a local maximum. He shows that the first-order conditions can be integrated – that is, there exist transfers satisfying them – if the agents have probabilistically independent information. He also shows that the second-order conditions are satisfied if the collective decision rule has a property he calls 'responsive' and that maximizing the sum of the agents' utilities, assuming truth-telling, has this property, namely, an agent's expectation of the product of his marginal utility for the public good and the marginal change in decision induced by a change in his information is always positive under truth-telling.

1.4 The Central Question

We are concerned more broadly with problems involving externalities, uncertainty, private information, and differing objectives among players, making team theory inapplicable, dictatorship and other 'pinning' mechanisms inadequately informed, and *laissez-faire* also inefficient. Furthermore we consider multi- as well as single-stage contexts.

Efficiency may still be achievable if appropriate financial incentives can be created to induce agents to take actions that are optimal for the group. The incentives we consider are transfer payments that may take the form of penalties, subsidies, compensations, taxes, and so on. We refer to our co-ordination mechanism as *incentive-based decentralization*. Our central research question is:

Under what circumstances can a group using incentive-based decentralization achieve as high an expected value as a team?

1.5 Description of a Group

We use the term *group* to signify that the decision-making units of concern are neither as cohesive as a team nor totally unco-ordinated. A group is composed of self-interested agents together with a central authority which implements a scheme of transfers among them intended to induce collectively desirable behaviour.

The scheme may be designed by the agents acting collectively and drawing up contracts and incentive agreements. The central authority is then merely a referee and clearing house. It monitors publicly available information and collects and dispenses payments on the basis of that information. Alternatively, the central authority might also design and impose the incentive scheme, if it has the power to do so. It could also be, at the same time, one of the agents. It is not, however, an external source of funds. The funds for making the transfer payments are extracted from and paid to the agents in the group. Thus the transfer budget must always balance. (A surplus could not be efficient.)

Our primary interest is in situations where the central authority can monitor agents' observations incompletely if at all. For this reason, among others, we seek here a scheme of incentives and penalties that is just powerful enough to bring each agent's objective into alignment with the group objective. The use of such gentle incentives may also be advantageous from the standpoint of fairness, individual initiative, public acceptability, and robustness against errors in assumptions. (For a variety of policy-relevant contexts, Schultze (1977) has outlined the attractive features of using incentives, in contrast to 'command-and-control' approaches, to promote public purposes.)

1.6 The Model

The model we shall consider is general with respect to uncertainty, private information, externalities, and agents' objectives. It also allows an arbitrary ordering of observations, signals, and actions. At each of any number of stages, each agent has three possible activities: observing information, sending signals, and taking actions. He may undertake none, some, or all of these activities at any particular stage. For convenience, we assume that each agent's signals and actions at each stage become known to all other agents at the next stage, though delays are easily accommodated in the model. Information private to an agent may or may not become known to other agents or the central authority later.

The central authority determines the incentive transfers among the agents on the basis of the information it obtains and the agents' actions and signals. Transfers can be made at each stage or only at the conclusion of the process.

It is assumed throughout that all agents are risk neutral and had common prior beliefs before any information was observed. The objective is therefore simply to maximize the expected total payoff to all agents. Once this is done, all points on the Pareto-efficient frontier can be achieved through lump-sum transfers.

1.7 Payment of the Expected Externality – Incentive-based Decentralization

All of our positive results revolve around the extension of the simple principle – most widely discussed in the context of pollution – of internalizing an externality: Pay each agent the total benefit to all other parties of his action (or change in action). This gives him an incentive to maximize total benefit (see Pigou, 1960). In extending this approach, we allow for uncertainty by using appropriate expected values, for the multistage incentive effect of budget balancing by recursion, and for private information by finding conditions under which the resulting transfers, that lead to a Pareto-efficient equilibrium, can be determined from information known to the central authority. Our formulation, like Arrow's, is a game of incomplete information, with a sequential Bayesian Nash equilibrium concept and incentives induced by transfer payments under the assumption of utility linear in income.

Because we employ the concept of expected externalities, we obtain a global equilibrium directly, without differentiation or restriction to one-dimensional or even continuous variables. We allow agents not only to send messages for use in a known collective decision rule but also to take acts affecting other agents and to communicate selectively with other agents in arbitrarily many stages. Our results apply to utilitarian objectives quite generally, not only to utilitarian collective decisions. They apply in other cases as well, but it is not immediately obvious how these compare with other 'responsive' cases. Besides knowing their own tastes, our agents may have all kinds of information that is not necessarily independent. In some situations our conditions might be hard to interpret or verify. Fortunately, under our predecessors' assumptions of independent information, utilitarian objective, and one stage, our mechanism is implementable if each agent's welfare is unaffected by other agents' private information. This is already more general than the usual models of collective decision, and it is by no means our most general model.

1.8 Examples of Efficient Incentives

The expected externality can be employed to engender efficiency in the sale of an item, with the buyer's and seller's reservation prices unknown. When presenting an early version of this paper at Stanford University, Zeckhauser offered to sell Kenneth Arrow a necktie featuring a pattern of extraordinary creativity and colours of great warmth, altogether befitting the intended purchaser. The bargaining mechanism reflected our recommended procedure. The audience helped to define the probability distribution of Arrow's reservation price for the tie. The expectation of the difference, if positive, between his reservation price and the seller's offering price was the expected externality of the seller's offer. With this amount paid by Arrow to the seller, the offering price will be the seller's reservation price, and a sale will always occur when there is any overlap in the bargaining range (see Chatterjee, Pratt and Zeckhauser (1978) for analytic details).

Several further essential elements of our approach will be illustrated in a pollution control situation. Upstream is a firm that discharges wastes into a river. Downstream is a firm that suffers damages r if the wastes are untreated. Upstream can treat the wastes at a cost c and reduce the damages to $r/2$. Thus the payoffs are:

Upstream Action:	*Treat*	*Don't Treat*
Upstream payoff	$-c$	0
Downstream payoff	$-r/2$	$-r$

Only Upstream knows its cost c, and only Downstream knows its sensitivity r. Assume that c is uniformly distributed on (0, 1) and r is uniformly distributed on (1, 2). Both are exogeneous. (For instance, r may vary according to Downstream's production process, this being dictated by external considerations such as conditions in the product markets.)

Example 1 No communication

Suppose first that Upstream must act with no further knowledge of Downstream's sensitivity. The benefit that Upstream conveys to Downstream by treating, relative to not treating as the benchmark, is $r/2$. At the time of its decision, Upstream's expectation of this externality, which we call the expected externality, is 0.75, since r is uniformly distributed on (1, 2). Paying this amount to Upstream if it treats induces it to treat whenever $c<0.75$. This is the desired group-optimal strategy subject to the no-communication conditions. To balance the budget, given that there is no other player or external banker, Downstream must make this payment. This has no incentive effect on Downstream since Downstream neither acts nor signals.

Example 2 Communication

Now suppose that, before Upstream acts, Downstream sends a signal s of its sensitivity. The group optimum is for Downstream to signal honestly $s=r$ and for Upstream to Treat if $c<s/2$. To make this a Nash equilibrium strategy for Upstream, merely pay him the expected (here also actual) externality he conveys, given honest signalling by Downstream. That externality is $s/2$ for Treat and 0 for Don't, the benchmark. Budget balance requires that Downstream make this payment.

Downstream's signal of sensitivity, whenever it influences Upstream to Treat rather than Don't, conveys two different externalities. There is a direct negative externality, Upstream's treatment cost c, and an indirect positive externality, the later incentive payment $s/2$ paid by Downstream to Upstream. Honest signalling can be made a Nash equilibrium strategy for Downstream by paying him the expected value of the sum of these two externalities given appropriate action by Upstream.

We now compute this expected value for each possible s. The sum is $s/2-c$ when Upstream Treats and 0 otherwise. Given s, Upstream should Treat whenever $c<s/2$. Since c is uniformly distributed on (0, 1), the probability of treatment is then $s/2$, the conditional expected value of c given that Upstream does Treat is $s/4$, and the conditional expected value of the sum $s/2-c$ given treatment is $s/2-s/4=s/4$. The unconditional expected value of the sum is the probability of Treat, times the conditional expected value given treatment, $(s/2)(s/4)=s^2/8$. This is the required incentive payment to Downstream. Making Upstream pay it

to balance the budget has no incentive effect on Upstream, since Upstream, having no earlier choices, can have no influence on Downstream's signal.

The net transfer from Upstream to Downstream, which makes group-optimal behaviour a (Bayesian) Nash equilibrium strategy for both, is the incentive payment to Downstream, namely $s^2/8$, less the incentive payment to Upstream, which is $s/2$ if he Treats and 0 otherwise. The net transfer is:

Payment from Upstream to Downstream as function of Upstream action and Downstream signal s

Upstream action:	Treat	Don't treat
	$s^2/8 - s/2$	$s^2/8$

A constant could be added to both entries, perhaps to make participation attractive to both parties, or in pursuit of some distributional goal. Changing the benchmark, for instance to Treat as the base case, also adds a constant or, in more complex problems, a quantity independent of the agents' actions. Group-optimal behaviour remains a Nash equilibrium under all such changes.

Can self-interested agents with private information always be induced to behave in a group-optimal manner through incentive arrangements of this type? Unfortunately not. Two kinds of difficulty are possible. One occurs if one agent's private information directly affects another's payoff. For example, if a polluter knows the expected health effects of its emissions on a particular day, while the recipient who suffers the impact does not know, and no one can monitor it, then efficiency cannot be induced through incentive payments. If the polluter can be paid an incentive only on the basis of what he reports and the action he takes, then with given treatment costs he will always make the same reports and take the same action, regardless of the health effects he actually predicts. The optimum is not achieved. The other kind of difficulty occurs when the unsignalled information of one agent is probabilistically dependent on the unsignalled information of another. Agents may then know more about expected externalities than the central authority does; the authority may therefore be unable to induce group-optimal behaviour. In the absence of these two specific difficulties, expected-externality methods obtain positive results.

1.9 Outline of the Analysis

We move from the simple to the complex. Section 2 considers single-stage situations. Here extracting the funds for one agent's incentive payment from other agents has no incentive effect on the latter. The key analytic question is when the information available to the central authority suffices to determine the expected externalities.

Section 3 shows that incentive-based decentralization can be employed with an arbitrary number of stages of observations, acts, and signals. In multistage environments, requiring one agent to contribute towards another agent's expected-externality payment has an incentive effect on his earlier choices. Fortunately, a conceptually interpretable mechanism is available to compensate and restore incentives for appropriate signals and actions. In our multistage discussion, we do away with the intuitive but unnecessary distinctions between signals and acts in favour of a simpler and more general framework. We show that,

> 'the team optimum can be achieved by a group using transfer mechanisms that rely solely on the observed actions and signals of agents and information monitored by the central authority, if any information observed by an agent that is not monitored by the central authority is irrelevant to all other agents' valuation functions and either signalled or conditionally independent of all unsignalled, unmonitored information possessed by all other agents.'

This strong positive result has surprising implications. For example, even if it is impossible to monitor the information of agents, the team optimum may be achievable if each agent is required to signal fully (honestly or dishonestly) the information he observes.

Many generalizations and applications of our methods are possible. Section 4 presents some of the most salient and some further remarks.

2 INCENTIVES LEADING A GROUP TO A TEAM OPTIMUM – ONE STAGE

2.1 Problem Formulation

We consider one-stage problems first. Since signals affect only later stages, they are superfluous here. Assumptions 1′ and 2′ apply solely to the one-stage situation. They will be replaced later.

Assumption 1′: Agents' observations–
Each agent i ($i = 1, \ldots, n$) observes the value r_i of a random vector $\tilde{r}_i$, his *observation*. The joint distribution of the $\tilde{r}_i$ is common knowledge. (A tilde in our notation indicates a random variable or vector.)

Assumption 2′: Choice of act–
Each agent i must choose one act a_i depending only on his own observation r_i. The set of available acts depends at most on r_i.

Thus a strategy for agent i is a function A_i on the r_i-space with values $A_i(r_i)$ in the a_i-space. We may and do allow randomized strategies without additional notation by introducing if necessary a component of $\tilde{r}_i$ which is continuously distributed independently of all else. If the agents send signals to a central authority which then acts according to a known rule, as in the 'public goods' or 'collective decision' problem (Section 4), the signals may simply be regarded as actions or components of actions and the formulation of this section applies. If agent i has no information or no choice of action, the r_i-space or a_i-space is a single point. A passive agent is relevant if affected by others and of concern to the group or if contributing to budget balance.

Assumption 3: Receipt of direct benefits–
Each agent i receives direct benefits (possibly negative) determined by the observations and actions of all the agents. He values these direct benefits at v_i, which we shall call his *direct return.*

We assume that individuals are risk-neutral in v_i; they will be maximizing under conditions of uncertainty, and will attempt to maximize the expected value of $\tilde{v}_i$. This objective is embodied in Assumption 6 below. In the present situation, v_i is a function of $(a_1, \ldots, a_n, r_1, \ldots, r_n)$, but this and the following assumptions will apply in the multistage case as well.

Assumption 4: Receipt of transfer–
Each agent i also receives a transfer payment u_i.

The transfer u_i could be thought of in a variety of ways, including penalties, taxes on externalities, or compensation for damages. Our main purpose is to design transfers that will lead to team-optimal behaviour by the agents. Our derivation will separate u_i into an incentive payment s_i intended to influence agent i and a balancing payment t_i which is agent i's contribution to the other agents' incentive payments ($u_i = s_i - t_i$). For the moment, however, we leave the determinants of u_i in abeyance and require only that the transfers balance across all agents.

Assumption 5: Transfers balance–
The sum of all transfer payments must be 0 for every outcome:

$$\sum_i u_i = 0. \tag{1}$$

The foregoing assumptions describe the physical structure of the situation. We next indicate the objectives of the agents and the group. Later we will specify the sense in which these objectives can be met and we will discuss the determination of the probability distributions that arise in the course of the analysis.

Assumption 6: Self-interested agents–
Agent i's objective is to maximize the expectation of his direct return plus transfer, which we shall call his expected *net return,* $E\{\tilde{v}_i + \tilde{u}_i\}$.

Assumption 7: Group objective–
The group objective is to maximize $E\{\sum_i \tilde{v}_i + \tilde{u}_i\}$. Given Assumption 5, this reduces to maximizing $E(\sum_i \tilde{v}_i)$.

2.2 Derivation of Incentives

The incentive mechanism that we propose is based on paying each agent the expected return, conditional on the information available to him, that all other agents receive in consequence of his action. It may be described as 'paying the expected externality'.

Consider the situation from agent i's point of view as he decides upon an action, having observed r_i. Let agent i's expectation of agent j's direct return be

$$s_{ij}=\mathrm{E}\{\tilde{v}_j|r_i\}; \tag{2}$$

it is a function of agent i's action a_i and observation r_i and the other agents' strategies. It is determined as follows: Given a_i and A_i, v_j is a function of the other agents' actions and observations. Given their strategies, the other agents' actions are specified functions of their observations. Thus v_j becomes a function of the other agents' observations. The expectation of $\tilde{v}_j$ given r_i is determined by this functional relationship and by the distribution of the other agents' observations, given r_i.

$$\mathrm{E}\{\sum_j \tilde{v}_j|r_i\}=s_{ii}+s_i, \tag{3}$$

where s_{ii} is agent i's expectation of his own return, and where

$$s_i=\sum_{j\neq i}s_{ij}=\mathrm{E}\{\sum_{j\neq i}\tilde{v}_j|r_i\} \tag{4}$$

is his expectation of the total return to the other agents. The quantity s_i may be interpreted as the expected externality of agent i's choice of action. Like the s_{ij}, it is a function of a_i, r_i, and the other agents' strategies. If agent i knew the other agents' strategies, then providing him a transfer payment s_i would make his objective coincide with the group objective.

Unfortunately, such transfer payments would not balance (add to zero) across all agents. Suppose, however, that agent j contributes an amount t_{ji} toward s_i, where

$$\sum_j t_{ji}=s_i, \qquad t_{ii}=0. \tag{5}$$

This balances the transfer payments. Furthermore, agent j's incentives are unaffected as long as t_{ji} is defined in such a way as to be uninfluenced by agent j's choice of action. This holds, for example, for $t_{ji}=s_i/(n-1)$, and indeed whenever t_{ji} is any function of the same arguments as s_i (the strategies used in defining s_i being fixed throughout).

If such transfers are instituted for every agent i, the net transfer to agent i is

$$u_i=s_i-\sum_j t_{ij}. \tag{6}$$

The resulting transfer scheme balances and makes each agent's objective 'agree' with the group objective, that is, coincide when all other agents follow the team-optimal strategy. This implies that a strategy A_i is optimal for agent i if and only if it is optimal for the group, on the assumption that the other agents use the strategies underlying the definition of the transfers. If, therefore, after finding the team optimum, we determine agent i's transfer payment (6) on the assumption that the other agents' strategies are team-optimal, then agent i will have an incentive to employ a team-optimal strategy also. In (6), the term s_i may be interpreted as an expected externality payment to agent i and $\Sigma_j t_{ji}$ as a balancing payment by agent i.

Specifically, the foregoing transfer scheme and its properties are summarized in the following theorem:

Theorem 1: Expected-externality transfers achieve the team optimum–
Take as given Assumptions 1′, 2′, and 3–7. Find a set of strategies A_i^* achieving the team optimum by maximizing the group objective $\mathrm{E}\{\Sigma_i \tilde{v}_i\}$. Find the expected externalities s_i based on these strategies. Define any set of balancing payments t_{ij} depending only on r_i and a_i with $\Sigma_i t_{ij}=s_j$ and $t_{jj}=0$ for all j. Let the net transfer to agent i be $u_i=s_i-\Sigma_j t_{ij}$. Then, as long as each agent i knows or assumes that the other agents will follow the strategies A_j^*, he can optimize for himself by choosing the strategy A_i^*, and the team optimum will be achieved. This transfer scheme therefore makes the team optimum a Pareto-efficient, Bayesian Nash equilibrium. If the team optimum is unique, then each such A_i^* is a unique optimum for agent i, no agent can deviate without penalty, and the equilibrium is strict.

2.3 Implementability of Incentives Based on Expected Externalities

What knowledge is required to compute and pay the transfers just defined? We are assuming throughout that the central authority knows the joint distribution of the $\tilde{r}_i$ shared by the agents, the actions available to the agents, and the agents' direct return functions. These suffice to determine the team optimum, a set of strategies for achieving it, and transfer rules that provide incentives to use these strategies. The actual transfer payments are in general functions of the agents' actions and observations. We assume throughout that all actions become known to all parties. We wish, however, to allow the agents' observations (and their actual direct returns) to be at least partly unobservable by others. In this case, it is possible that the transfers of Theorem 1 cannot be implemented in practice. We seek conditions under which they can be.

Any part of an agent's observation that becomes known to or can be obtained by the central authority, and any quantity that depends only on such information, will be called *public* or *publicly observable*. It is sometimes convenient to express this in another way. Let z_i be the public portion of agent i's observation, that is, z_i is a publicly observable function of r_i such that any other publicly observable function of r_i is a function of z_i. Then a quantity is publicly observable if, and only if, it is a function of the z_i's.

A transfer rule will be called *implementable* if the information needed to determine all payments that it could ever require is publicly observable. This means that it relies on public information only; no matter what acts the agents choose, all transfer payments are functions only of information that can be obtained by the central authority. Under what conditions are expected-externality transfers implementable? If the expected externality s_i is publicly observable, then the balancing payments t_{ji} can be chosen to be publicly observable (for example, $t_{ji} = s_i/(n-1)$). In any discussion of implementability, we shall take such a choice for granted. This reduces the question to the public observability of the expected externalities s_i. In short, condition (B) below is essentially the definition of (A).

(A) Expected-externality transfers are implementable.
(B) $E\{\tilde{v}_j|r_i\}$ is public for all i and j with $j \neq i$, for all actions a_i, when all agents $j \neq i$ follow the team-optimal strategies A_j^*.

To check implementability in any given situation, of course, one could simply calculate the incentives and see if they are publicly observable. However, general conditions for implementability (listed below) are also useful, especially if they provide some insight.

(C) The distribution of $\tilde{v}_j$ given r_i depends only on z_i (the public portion of r_i), for all i and j with $j \neq i$, for all acts a_i, when all agents $j \neq i$ follow the team-optimal strategies A_j^*.
(D) Each agent's direct return is a function of all agents' actions, his own observation, and only the public portions of the other agents' observations. Each agent's conditional distribution of the other agents' observations given his own observation depends only on the public portion of his own observation.
(E) The public portions of all agents' observations are the same, say z. Each agent's direct return is a function of all agents' actions, his own observation, and z. The agents' observations are conditionally independent given z.
(F) Each agent's direct return is a function only of all agents' actions and his own observation. The agents' observations are independent.

Each of these conditions is stronger than those before it but simpler structurally or intuitively. Others along the same lines could easily be given and might prove useful in particular situations, but these are enough to illustrate the possibilities. Examples could readily be constructed satisfying implementability but not (C), satisfying (C) but not (D), and so on. Condition (D) appears likely to hold in practice, however, in a wide variety of situations. It provides the basis for the title of the following summary theorem, though a weaker Condition (C) could be employed, of course.

Theorem 2: Expected-externality transfers are implementable if the private portions of agents' observations are independent and do not affect each other's direct returns–
Posit Assumptions 1′, 2′, and 3–7 as in Theorem 1. Then if any one of the conditions (C) to (F) holds, implementability, (A) and (B), follows. Moreover, each condition is implied by the next.

Before extending our analysis to two-stage and multistage situations, we should remark that our expected externality transfers are adaptable to a variety of normative criteria which in effect would add or subtract an amount from each player based perhaps on their observations, but not on their actions (see Section 4.2). The level of expected-externality payments provides no ethical baseline in and of itself, just as the effluent charges that would be collected taking zero emissions as the base level – a much simpler problem – may not be the fair fee for a polluter to pay for his action.

3 INCENTIVES LEADING TO A GROUP OPTIMUM WITH ARBITRARILY MANY STAGES OF OBSERVATIONS, ACTIONS, AND SIGNALS

We now consider a general situation in which there is an arbitrary (but fixed and known) number of stages in which observations may be made, acts may be taken, and signals may be sent. We index the stages by the number of stages remaining. Thus stage 0 is the last stage, and stage M, say, is the first stage. All finite sequences of observation, action and signal can be included within our formulation. Now, however, in choosing an action or signal at any stage after the first, each agent can take advantage of what he can learn from the other agents' earlier choices and of any additional observation of his own which he may have made in the meantime. Thus one agent's earlier choices may influence another agent's later ones. This introduces more complicated incentive effects than occur in one-stage situations.

Since signals are structurally equivalent to actions, as mentioned earlier, we shall omit signals from the problem formulation in this more general section. This will simplify the notation and the description of our results. The loss of specificity in dealing with signals will be made up by occasional comments. Furthermore, our results are technically more general, not less, because our theorems require only that no agent can improve the group objective by himself, not that all agents maximize jointly. For example, the team optimum with truthful signals about only a part y_i of each agent i's observation r_i may be inferior to a team strategy conveying more information. The greater information might be conveyed by a truthful signal about r_i or even by an untruthful signal appropriately decoded by the recipient. However, such a signal by one agent cannot improve the team outcome unless another agent interprets it correctly and alters his action accordingly. Thus the team optimum with truthful signals about y_i only has the property that no agent can increase the group objective by himself. This suffices for our theorem to show that this constrained team optimum can be achieved by incentives based on the corresponding expected externalities. A global optimum is not required.

3.1 Problem Formulation

To accommodate the multistage structure, only the first two assumptions need to be changed.

Assumption 1: Agents' observations–
Each agent i ($i=1,\ldots,n$) observes a value r_i^m at each stage m. Stages are numbered from M to 0 according to the number m remaining. The joint distribution of all r_i^m is common knowledge.

Assumption 2: Agents' decisions–
At each stage m, each agent i must choose an action a_i^m depending only on his own observations through stage m and all agents' actions in all previous stages. The sets of available actions depend at most on the same quantities.

The direct returns v_i, the transfers u_i, and the individual and group objectives satisfy Assumptions 3 to 7 as before, but v_i and u_i are now functions of all actions and observations. Note that v_i represents the total direct benefits of agent i, including any benefits received before the terminal stage.

A strategy for agent i is now a vector of act-choice functions $A_i^M,\ldots,A_i^0$ with the domains and ranges indicated in Assumption 2. The domains imply perfect recall. Randomized strategies are included as before by adjoining independently, continuously distributed components to the r's as needed.

Signals are simply acts or components of acts which have no influence on the direct returns v_i, that is, are superfluous or 'inactive' as arguments of the direct return functions, once the type of signal, for example, dimension and precision, is chosen. (If signalling costs are to be included, the choice of type of signal should be modelled as a separate action component.) Our formulation thus allows for a signal which is a possibly untruthful report of any portion of agent i's observation, say of $y_i^m = Y_i^m(r_i^M,\ldots,r_i^m)$, where Y_i^m is a known function. An important special case has $y_i^m = r_i^m$; that is, all information is available for exchange.

Assumption 2 implies that all other agents receive the same report from agent i, but this is merely for convenience. Our results would apply even if each agent were allowed to send a different signal to each other agent. Each agent's report may go to a central agency which then disseminates some or all of the information to everyone.

3.2 Derivation of Incentives in the Multistage Case

The principles for deriving incentives are fairly straightforward, though the mechanics in any particular context may prove complicated. Earlier actions (including signals) will influence not only subsequent actions in a direct manner, but also the magnitude of subsequent incentive payments. Self-interested agents will take these effects on later incentive payments into account. Earlier incentive payments will have to be structured to compensate for them. At the last stage, each agent's expected externality (computed assuming all other agents' earlier decisions were in accord with the team optimum) provides an appropriate incentive, since nothing follows. Appropriate adjustments to the incentive payments for the next-to-last action can then be computed in the light of the expected effect each agent's choice of action has on the balancing payments he will be required to make in support of the other agents' last period incentives. Continuing in this way we can fold back to the first stage and its appropriate incentives.

To see how this all works out, we first consider the situation from agent i's point of view as he decides upon his terminal action a_i^0. He knows all his own observations $r_i^M, \ldots, r_i^0$ and all previous actions of all agents, $a_1^M, \ldots, a_n^M, \ldots, a_1^1, \ldots, a_n^1$. We call this his *stage-0 information* and denote it by q_i^0. Let agent i's expectation of agent j's direct return, in the light of this information, be

$$s_{ij}^0 = \mathrm{E}\{\tilde{v}_j | q_i^0\}. \tag{7}$$

It is a function of agent i's terminal action a_i^0 and his information q_i^0 (including his earlier acts) and the other agents' strategies. This expectation is determined as follows: Given q_i^0, v_j is a function of the other agents' terminal actions and their observations $r^M, \ldots, r^0$. Given their strategies and q_i^0, their terminal actions are specified functions of their observations. Thus v_j becomes a function of their observations. The expectation of v_j given q_i^0, is determined by this functional relationship and by the distribution of the other agents' observations $r^M, \ldots, r^0$, given q_i^0.

As before, our first step is to give agent i a transfer payment

$$s_i^0 = \textstyle\sum_{j \neq i} s_{ij}^0 = \mathrm{E}\{\sum_{j \neq i} \tilde{v}_j | q_i^0\}. \tag{8}$$

This also is a function of his terminal action and information and the other agents' strategies. Giving agent i the transfer payment s_i^0 makes his objective coincide with the group objective at the time of his terminal decision and, consequently, at the time of his earlier decisions as well. Suppose we balance this payment by transfers t_{ji}^0 from other

agents j to i, where $\Sigma_j t_{ji}^0 = s_i^0$ with $t_{ii}^0 = 0$. As long as t_{ji}^0 is a function of the same arguments as s_i^0, it will have no incentive effect on agent j's terminal decision. It will, however, have an incentive effect on agent j's earlier decisions, since these are included in q_i^0. We can offset this incentive effect by a transfer to agent j equal to his expectation of $\tilde{t}_{ji}^0$ at the time of his decision at stage one. The total of these transfer payments to agent j is

$$s_j^1 = \mathrm{E}\{\textstyle\sum_i \tilde{t}_{ji}^0 | q_j^1\}. \tag{9}$$

We balance the transfers s_j^1 by transfers t_{ij}^1 from i to j with $\Sigma_i t_{ij}^1 = s_j^1$, but this introduces further incentive effects which need to be offset. Continuing similarly, for $m = 0, 1, 2, \ldots, M$ successively, we define s_i^0 by (8), let

$$s_i^m = \mathrm{E}\{\textstyle\sum_j \tilde{t}_{ij}^{m-1} | q_i^m\} \text{ for } 1 \leqslant m \leqslant M, \tag{10}$$

and define any transfers t_{ij}^m depending on the same arguments as s_j^m such that

$$\textstyle\sum_i t_{ij}^m = s_j^m \text{ and } t_{jj}^m = 0 \text{ for } 0 \leqslant m \leqslant M. \tag{11}$$

The transfers s_i^m, and hence t_{ji}^m, are functions of agent i's stage-m action a_i^m and information q_i^m, that is, of $a_i^M, \ldots, a_i^m$, $r_i^M, \ldots, r_i^m$, and all $a_j^M, \ldots, a_j^{m+1}$ for $j \neq i$. In taking an expectation given q_i^m, agent i's actions $a_i^M, \ldots, a_i^m$ are decision variables but probabilities are conditional on agent i's observations $r_i^M, \ldots, r_i^m$ and whatever inferences he can derive from the other agents' previous actions $a_j^M, \ldots, a_j^{m+1}$ on the assumption that they are following the strategies used in defining the transfers. Since t_{ij}^M depends only on agent j's initial information r_j^M and action a_j^M, it is uninfluenced by any choice made by any other agent and hence no further offsetting of incentive effects is needed.

Combining the foregoing transfers gives a transfer schedule that balances and makes each agent's objective agree with the group objective as regards his entire strategy, provided all other agents follow the strategies used in defining the transfer. In other words a strategy is optimal for agent i if, and only if, it is optimal for the group, when the other agents use the strategies underlying the definition of the transfers. Now suppose that a given set of strategies is *agent-by-agent team-optimal*, meaning that each agent's strategy maximizes the group objective if the other agents use the given strategies. Then, and this is the basis for our central result, the transfers defined above in terms of these strategies will make the strategies optimal for each agent, assuming that the other agents are also following them. Thus, through the use of appropriate transfers, we can induce agents to follow any agent-by-agent team-optimal set of strategies.

Probably the most important case of an agent-by-agent team optimum is the over-all team optimum when there are no constraints on the types of signals to be sent. Note also that a team optimum with truthful messages about known functions y_i^m of the agents' observations is agent-by-agent team-optimal and hence covered by our results. Of course the team optimum with signals depending only on y_i^m may be dominated by the team optimum with signals conveying more information. If, however, practical considerations would limit the information exchanged by the group acting as a team to y_i^m, then the relevant comparison is to a team exchanging just this information, and this is the team optimum we achieve.

Our theorem on incentive properties of multistage expected externality transfers is

Theorem 3: Multistage expected-externality transfers can achieve any agent-by-agent team optimum–
Take as given Assumptions 1 to 7. Let the act-choice functions A_i^{m*} $(i=1,\ldots,n;\ m=M,\ldots,0)$ form an agent-by-agent team-optimal set of strategies. Use these strategies to define expected externalities s_i^m by (8) and (10) and balancing payments t_{ij}^m satisfying (11) successively for $m=0, 1, \ldots, M$. Let the net transfer to agent i be

$$u_i=\sum_{m=0}^{M}\left(s_i^m-\sum_{j=1}^{n}t_{ij}^m\right). \tag{12}$$

If the agents other than i use the strategies A_j^{m*} then agent i can optimize for himself by choosing A_i^{m*}. This transfer scheme makes the agent-by-agent team optimum a sequential, Bayesian Nash equilibrium. The equilibrium is strict if each agent's strategy uniquely maximizes the group objective when the other agents use the strategies A_j^{m*}. The equilibrium is Pareto-efficient in any set of vectors of strategies within which the group objective is maximized by the agent-by-agent team-optimal strategies A_i^{m*}.

Proof

Although the discussion leading up to Theorem 3 was presented so that

the theorem would follow naturally, the reader may find a proof useful. If all agents except i employ the strategies A_j^{m*}, then whatever strategy agent i employs we have, by (10),

$$E\{\tilde s_i^m-\textstyle\sum_j \tilde t_{ij}^{m-1}|q_i^m\}=0. \tag{13}$$

It follows that $E\{\tilde s_i^m-\sum_j \tilde t_{ij}^{m-1}\}=0$ and therefore, by (12),

$$E\{\tilde v_i+\tilde u_i\}=E\{\tilde v_i+\tilde s_i^0-\textstyle\sum_j \tilde t_{ij}^M\}. \tag{14}$$

From this it follows in turn, similarly, by (8), that

$$E\{\tilde v_i+\tilde u_i\}=E\{\tilde v_i+\textstyle\sum_{j\neq i}\tilde v_j\}-E\{\sum_j \tilde t_{ij}^M\}. \tag{15}$$

The last term is a constant, unaffected by agent i's strategy. Therefore the transfer u_i makes agent i's objective coincide with the group objective. The equilibrium result follows. If the equilibrium is not strict, then some agent can deviate without changing $E\{\tilde v_i+\tilde u_i\}$ and hence his strategy does not uniquely maximize the group objective. Thus unique agent-wise maximization implies strict equilibrium. The Pareto-efficiency statement is obvious.

3.3 Implementability of Incentives Based on Multistage Expected Externalities

Now that the expected-externality approach to designing incentives has been extended to multistage situations, we again face the question of implementability. We assume as before that the balancing payments t_{ji}^m are chosen to facilitate implementability, for instance, to be fixed fractions of the incentive payments s_i^m. Then the multistage expected-externality transfers are implementable if the relevant conditional expectations of the agents are publicly observable. The condition at stage 0 is

(a^0) $E[\tilde v_k|q_i^0]$ is public for $k\neq i$.

At stage m, $1\leqslant m\leqslant M$, the condition is

(a^m) $E[E\{\ldots E(\tilde v_k|\tilde q_{j_0}^0)\ldots|\tilde q_{j_{m-1}}^{m-1}\}|q_i^m]$ is public for $k\neq j_0\neq j_1\neq\ldots\neq j_{m-1}\neq i$.

(These and later conditions are to be understood as holding for all possible values of all indices and unbound variables, and distributions conditional on one agent's information are to be calculated on the assumption that all other agents follow team-optimal strategies.)

As before, we can check the implementability of incentive payments based on expected externalities when there are many stages by simply calculating the incentives and verifying if they are publicly observable. Again, however, it is also useful to have general conditions for implementability, and it is possible to give some that have intuitive content. They can be expressed in terms of the following conditions. The exact relationships among conditions are somewhat more complicated than before, however, and will be deferred to the statement of the theorem.

(b^0) The conditional distribution of $\tilde{v}_k$ given q_i^0 is public for $k \neq i$.

(b^m) ($m \geqslant 1$). The conditional distribution of $E\{\ldots E(\tilde{v}_k|\tilde{q}_{j_0})\ldots|\tilde{q}_{j_{m-1}}^{m-1}\}$ given q_i^m is public for $k \neq j_0 \neq j_1 \neq \ldots \neq j_{m-1} \neq i$.

(c^0) Agent i's conditional distribution, given his terminal stage information q_i^0, of all other agents' terminal actions, agent j's observations at all stages, and the public portions of all observations of all other agents, is public for $j \neq i$.

(c^m) ($m \geqslant 1$) Agent i's conditional distribution, given his information q_i^m at stage m, of all other agents' actions at stage m and the public portions of agent j's observations through stage $m-1$, is public for $j \neq i$.

(d^0) Agent i's conditional distribution, given his terminal-stage information q_i^0, of all other agents' observations at all stages, is public.

(d^m) ($m \geqslant 1$) Agent i's conditional distribution, given his information q_i^m at stage m, of all other agents' observations through stage m and the public portion of agent j's observation at stage $m-1$, is public for $j \neq i$.

(e) Each agent's direct return is a function of all agents' acts, his own observations, and only the public portions of the other agents' observations.

The following multistage implementability theorem generalizes Theorem 2.

Theorem 4: Implementability of multistage expected externality transfers–
The multistage expected-externality transfers of Theorem 3 are implementable if condition (a^m) above is satisfied for all m, $0 \leqslant m \leqslant M$. For all m, condition (a^m) is implied by (b^m). If (e) holds, then (b^0) is implied by (c^0) and (c^0) by (d^0). For $m \geqslant 1$, if (a^{m-1}) holds, then (b^m) is implied by (c^m) and (c^m) by (d^m). Implementability, therefore, follows if either (a^0) or (b^0) or (c^0,e) or (d^0,e) holds and, for each $m \geqslant 1$, either (a^m) or (b^m) or (c^m) or (d^m) holds. Conditions (a^m) – (d^m) are unaffected if 'is public' is replaced by 'depends on agent i's observations only through their public portions', or by 'is a function of agent i's actions through stage M, all other agents' actions through stage $m+1$, and only the public portions of agent i's observations through stage m'. In conditions (b^m) – (d^m), publicness of the conditional distribution, given q_i^m, of the other variables mentioned, is equivalent to conditional independence between agent i's observations through stage m and the other variables, given the public portions of the former.

The proof will be given at the end of the next subsection.

3.4 Interpretation of the Requirements for Implementability

Condition (a^0) is equivalent in the one-stage case to condition (B) of Section 2, and (b^0) to (C). Conditions (d^0) and (e) together are equivalent to condition (D), which was the basis of the title of Theorem 2. That (d^0), (d^m), and (e) suffice for implementability in the multistage case is the result stated (very roughly) in the Introduction.

The full meaning of conditions (a^0) and (b^0) becomes apparent upon reviewing the functional dependencies as was done after (7). Whatever his choice of actions (and signals) at all stages, agent i can calculate a conditional mean and distribution of $\tilde{v}_j$, given his observations $r_i^M, \ldots, r_i^0$ and the actions (and signals) of the other agents at all stages before the last, on the assumption that the other agents are using the designated team-optimal strategies. Conditions (a^0) and (b^0) say that whatever agent i chooses to do at the last stage, and whatever he and the other agents have done previously, this conditional mean and distribution, respectively, depend on agent i's observations $r_i^M, \ldots, r_i^0$

only through their public portions, and hence can be determined by the central authority as well as by agent i.

To see the meaning of conditions (a^1) and (b^1), suppose that agent j has calculated $E\{\tilde{v}_k|q_j^0\}$ as just described and that agent j is using a team-optimal strategy. $E\{\tilde{v}_k|q_j^0\}$ is, of course, a function of q_j^0, that is, of $r_j^M, \ldots, r_j^0$, and all agents' actions through stage 1. Now consider agent i. As he makes his choice at stage 1, he knows only q_i^1, that is, $r_i^M, \ldots, r_i^1$ and all agents' actions through stage 2, and the action and signal he will choose at stage 1. If the other agents follow the team-optimal strategy, for each action that agent i might choose at stage 1, he can calculate a conditional distribution, given q_i^1, of the other agents' actions at stage 1 and $\tilde{r}_j^M, \ldots, \tilde{r}_j^0$. He can therefore calculate the conditional mean and distribution of $E\{\tilde{v}_k|\tilde{q}_j^0\}$. Conditions ($a^1$) and ($b^1$) require that, whatever action agent i may choose at stage 1, this conditional mean and distribution depend on his observations $r_i^M, \ldots, r_i^1$ only through their public portions.

The following lemmas further clarify the meaning of our implementability conditions and the relationships among them. They also underlie the proof of Theorem 4. Lemma 1 shows that each condition can be stated in several forms that sound quite different but are actually equivalent. Lemma 2 provides the basic method of showing that successive conditions are stronger.

Lemma 1: Equivalent conditions for a conditional distribution to be public–
Let v be a function, possibly vector-valued, of d and x, where d is a vector of agent i's decision variables and $\tilde{x}$ is, from agent i's point of view, a random vector whose distribution depends only on d. Let q be agent i's information and let z be the public portion of q. The following conditions on the distribution of $\tilde{v}$ given q, denoted $D(\tilde{v}|q)$, are equivalent.

(i) $D(\tilde{v}|q)$ is public.
(ii) $D(\tilde{v}|q)$ is a function of z and d only.
(iii) $D(\tilde{v}|q)$ depends on q only through z.
(iv) $\tilde{v}$ and $\tilde{q}$ are conditionally independent given z.

Conditions (i) to (iii) are to be understood as holding for all q, and (i), (iii), and (iv) for all d. Components of $\tilde{v}$ which are functions of z and d can be omitted with no effect on the conditions. In (iii), 'on q' can be replaced by 'on $\bar{q}$' if $\bar{q}$ is a function of q such that z is a function of $\bar{q}$. In (iv), $\tilde{q}$ can be replaced by any function of $\tilde{q}$ which, together with z, determines $\tilde{q}$. Decision variables in q are irrelevant except as parts of d. What portions of v and x are public or private is also irrelevant.

Lemma 2: Relationships among conditions of publicity–
In the situation of Lemma 1,

$$D(\tilde{x}|q) \text{ is public} \Rightarrow D(\tilde{v}|q) \text{ is public} \Rightarrow E(\tilde{v}|q) \text{ is public.}$$

If $\tilde{x}$ is substituted for $\tilde{v}$ in conditions (i) to (iv), the resulting conditions are equivalent to one another. If $E(\tilde{v}|q)$ is substituted for $D(\tilde{v}|q)$ in conditions (i) to (iii), the resulting conditions are equivalent to one another.

Proofs

The lemmas are easily proved once the definitions and concepts involved are clearly understood. In Theorem 4, condition (a^0) says that all s_{ij}^0 are public, whence the s_i^0 are implementable. If we choose, for instance, $t_{ij}^0 = s_j^0/(n-1)$, then the t_{ij}^0 will also be implementable, and we will have

$$s_i^1 = E\{\textstyle\sum_{j\neq i} t_{ij}^0 | q_i^1\} = \frac{1}{n-1} \textstyle\sum_{j\neq i} \sum_{k\neq j} E\{E\{\tilde{v}_k|q_j^0\}|q_i^1\}. \tag{16}$$

These quantities s_i^1 are implementable under condition (a^1), whence so are the quantities $t_{ij}^1 = s_j^1/(n-1)$. Continuing similarly, we see that (a^m) for all m implies implementability. The rest of Theorem 4 is proved by applying Lemmas 1 and 2 with careful attention to the variables that are relevant.

3.5 The Case of Common Public Information

If all agents' observations have the same public portions z^m, say, at all stages m, they, of course, become known to all agents immediately. In this case, in conditions (c^m) and (d^m), the only public portion of agent j's observations needed is z^{m-1}, and (c^0) becomes

(c_0^0) Agent i's conditional distribution, given his terminal-stage information q_i^0, of all other agents' terminal action and agent j's observations at all stages, is public for $j \neq i$.

An extreme case is that no observation has a public portion, that is, all information is private; then (c^0) simplifies to (c_0^0) and (c^m) and (d^m) simplify to the following conditions:

(c_0^m) ($m \geq 1$) Agent i's conditional distribution, given his information q_i^m at stage m, of all other agents' actions at stage m, is public.

(d_0^m) Agent i's conditional distribution, given his information q_i^m at stage m, of all other agents' observations through stage m, is public.

3.6 The Case of Prompt Publicity

In Theorem 4, the public portions of each agent's observations may not become known to the other agents until after all decisions have been made. They are merely required to be known to the central authority at the time of implementing the transfers. Of course, what is known, when, and by which agents, is reflected in the observation variables r_i^m. An extreme situation is *immediate publicity*; all public information becomes known immediately to all agents. This is equivalent to assuming that the public portions of all agents' observations are the same at all stages. This case was discussed briefly above.

Another interesting possibility is one we shall call *prompt publicity*: the public portion of each agent's observation at each stage becomes known to all other agents *promptly*, meaning at the next (or same) stage. Since each agent's actions also become known promptly, this is equivalent to the condition that the public portion of q_i^m is included in q_j^{m-1} for all $m \geq 1$ and $j \neq i$. The public portion of q_i^m, which we denote by p_i^m, consists of the public portions of agent i's observations through stage m, all actions through stage $m+1$, and agent i's action at stage m. Another condition equivalent to prompt publicity is that p_i^m is a function of p_j^{m-1} for all $m \geq 1$ and $j \neq i$.

From this characterization, by an elementary chain property of conditional expectations, it follows that, under the condition of prompt publicity,

$$\mathrm{E}\{\mathrm{E}(\tilde{x}|\tilde{p}_j^{m-1})|p_i^m\} = \mathrm{E}(\tilde{x}|p_i^m) \text{ for all random variables } \tilde{x}. \tag{17}$$

By repetition one obtains

$$\mathrm{E}[\mathrm{E}\{\ldots \mathrm{E}(\tilde{v}_k|\tilde{p}_{j_0}^0)\ldots|\tilde{p}_{j_{m-1}}^{m-1}\}|p_i^m] = \mathrm{E}(\tilde{v}_k|p_i^m). \tag{18}$$

Note also that $\mathrm{E}(\tilde{x}|q_i^m)$ is public if and only if

$$\mathrm{E}(\tilde{x}|q_i^m) = \mathrm{E}(\tilde{x}|p_i^m). \tag{19}$$

Conditions (a^m) and (b^m) therefore simplify in Theorem 4 as follows:

(a_1^0) $\mathrm{E}(\tilde{v}_k|q_i^0) = \mathrm{E}(\tilde{v}_k|p_i^0)$ for $k \neq i$.

(a_1^m) $\mathrm{E}(\mathrm{E}\{\tilde{v}_k|p_j^{m-1}\}|q_i^m) = \mathrm{E}(\tilde{v}_k|p_i^m)$ for $j \neq i$ and $m \neq 1$ or $k \neq j$.

(b_1^m) The distribution of $\mathrm{E}(\tilde{v}_k|\tilde{p}_j^{m-1})$ given q_i^m is public for $j \neq i$ and $m \neq 1$ or $k \neq j$.

Theorem 5: Implementability when publicity is prompt–
Under the condition of prompt publicity, in Theorem 4, condition (a^m) holds for all m if, and only if (a_1^m) holds for all m, and (b^m) holds for all m if, and only if, (b_1^m) holds for all m.

The only simplification in the remaining conditions of Theorem 4 produced by prompt publicity is that (c^0) and (e) could refer to the public portions of other agents' terminal observations only and (c^m) could refer to agent j's observations at stages m and $m-1$ only.

If multistage expected-externality transfers are defined with balancing payments t_{ji}^m which are always specified fractions of the expected externalities s_i^m, then each s_i^m will be a sum of specified fractions of terms of the type appearing in condition (a^m) of Theorem 4. If these terms are all public, then the expression in (a^m) equals the left-hand side of (18). Under the condition of prompt publicity, this equals the right-hand side of (18), and s_i^m and t_{ji}^m are therefore expressible as a weighted sum of terms $\mathrm{E}(\tilde{v}_k|p_i^m)$. We will now illustrate this for several cases of interest. We assume implementability and prompt publicity throughout.

Two agents

$s_i^m = t_j^m = \mathrm{E}(\tilde{v}_k|p_i^m)$ for $i \neq j$, where $k = j$ for m even and $k = i$ for m odd.

Cyclical balancing payments

Each agent pays the previous agent's expected externalities: $s_i^m = t_{i+1}^m = \mathrm{E}(\Sigma_j^* \tilde{v}_j | q_i^m)$ where the sum is over all j except the j for which $(m+j-i)/n$ is an integer and where $i+1$ is replaced by 1 when $i=n$.

Equally shared balancing payments

Each agent's expected externalities are paid in equal shares by all other agents:

$$s_i^m = \mathrm{E}(b_{m-1}\tilde{v}_i + b_m \textstyle\sum_{j\neq i} \tilde{v}_j | q_i^m),\ t_{ji}^m = \dfrac{1}{n-1} s_i^m \text{ for } j \neq i, \tag{20}$$

where

$$b_{-1}=0,\ b_0=1,\ b_1=\frac{n-2}{n-1},\ b_2=\left(\frac{n-2}{n-1}\right)^2+\frac{1}{n-1}, \tag{21}$$

$$b_m=\frac{n-2}{n-1}b_{m-1}+\frac{1}{n-1}b_{m-2}=\sum_{k\leq m/2}\binom{m-k}{k}(n-2)^{m-2k}(n-1)^{k-m} \text{ for } m\geq 1. \tag{22}$$

The definition of equal sharing used here does not lead to equal values of every agent i's total balancing payment $\Sigma_{j,m} t_{ij}^m$, since agent i does not contribute to balancing s_i^m. An alternative definition would have $t_{ij}^m = s_j^m/n$ for all i including $i=j$. The same net transfer as before can be achieved while satisfying this alternative definition by a suitable redefinition of the s_i^m. Since each agent i will take into account that he will pay a fraction $1/n$ of his own s_i^m, and thus net only $s_i^m(n-1)/n$, the s_i^m need merely be increased by the factor $n/(n-1)$. Specifically, with s_i^m defined as before, let

$$S_i^m = \frac{n}{n-1} s_i^m \text{ and } T_{ij}^m = \frac{1}{n} S_j^m \text{ for all } i,j,m. \tag{23}$$

Then

$$S_i^m - \textstyle\sum_j T_{ji}^m = s_i^m - \dfrac{1}{n-1}\textstyle\sum_{j\neq i} s_j^m \tag{24}$$

and the net transfers resulting from the S_i^m and T_{ij}^m are the same at each stage as those resulting from the s_i^m and t_{ij}^m. This holds, of course, whether or not publicity is prompt. With prompt publicity, the alternative definition of equal sharing is satisfied by the same formulas as before, (20) to (22), except that every b_m is multiplied by $n/(n-1)$ and $t_{ji}^m = s_i^m/n$ for all i,j including $j=i$.

4 REMARKS, GENERALIZATIONS AND APPLICATIONS

In Sections 1 to 3 we investigated the conditions under which a group can achieve a team optimum through the use of financial incentives. It was possible, of course, to detail only a few sufficient sets of assumptions. Here we sketch some of the most important further generalizations of our results, which relate to the structure of the value function, possibilities for and restrictions on observing and transmitting information, and uncertainty about settling up or the length of the game. Subsections 4.7 and 4.8 address the application of expected-externality methods to some classic problems in collective decision. (See Pratt and Zeckhauser, 1980, for further remarks, in particular about some properties of the equilibrium.) (See equilibrium.)

We begin this section with a perspective on our approach.

4.1 Features and Limitations

The traditional fodder of decentralization discussions has been market processes, public-good provision and other collective decisions, and the allocation of a scarce resource among divisions of a firm. Our analysis begins with the observation that individuals frequently affect one another directly by their actions; no decision by a central authority need be involved. Members of a cartel set production levels, firms and households near a river pollute it and suffer the pollution, workers' efforts influence company profits, and lawyers' care in drawing up contracts affects their clients. The first key element of our analysis is that *agents' actions may directly affect the welfare of other agents.*

In many situations there is more information worth transmitting than individual preferences or production opportunities. A potential polluter may know something about the cost to others of cleaning up his pollutant. Individuals may have a variety of information about an underlying uncertainty of interest to all. They may even observe one another's information. Our second important feature is to allow agents to *possess and signal information about their own, other agents', or a central authority's costs, actions, or information.*

Allowing agents' actions to be guided by previous signals requires an analysis covering at least two stages, with signals in one stage preceding actions in another. Indeed, in reality, when information and signalling guide agents' actions, there are likely to be many rounds of interaction. We capture this central element of the real world in the third significant aspect of our analysis: a *multiple-stage environment* that can accommodate any sequencing within any finite number of stages. If agents could commit themselves to multistage strategies, and if those commitments could be monitored, the multiple stages could be collapsed to one, with incentive payments depending on the agents' strategies. Clearly, however, such commitment and monitoring are usually impossible; thus collapsing is ruled out.[2]

Our multistage setting is reflected in the form of our incentive payments. It also allows us to generalize and provide intuitive meaning to the information structures under which our approach is successful. These can be quite complicated. Informally, having agents signal private information and rewarding them according to their signals enables us to cope with situations in which the central authority cannot observe information that would otherwise be crucial to determining the incentive payments. Table 13.1 identifies the critical features for competitive markets, dominant-strategy mechanisms for demand revelation, and Bayesian collective decision processes in relation to expected-externality methods of incentive-based decentralisation.[3]

Our solution has three primary weaknesses: First, it is not coalition-proof. Here we are in good company. No relatively general mechanism for eliciting honest information from agents, including the competitive market itself, is defensible against coalitions. Secondly, we obtain only a Bayesian Nash equilibrium. But we have already seen that a necessary price of balancing the budget is to give up dominance. Thirdly, we require that it be known how the agents' probabilistic beliefs about each other's information depend on their own information. This is another sacrifice for budget balance. We take mild reassurance that Arrow (1977, 1979) and d'Aspremont and Gérard-Varet (1975, 1979) invoked both this assumption and independence in their major constructive result for a balanced budget.[4]

What is particularly reassuring about our results is that the incentive scheme that can induce agents to co-ordinate actions efficiently and reveal information honestly is motivated by the strong intuit[illegible] notion

Table 13.1 Mechanisms that achieve efficient outcomes

A. *Mechanism*	I *Competitive market with contingent contracts*	II *Demand-revealing dominant strategy processes*		-	III *Bayesian, collective decision*	IV *Incentive-based decentralisation*
B. *References*	Arrow and Debreu (1954) Arrow and Hahn (1971)	Vickrey (1961)	Clarke (1971)	Groves (1973)	d'Aspremont and Gérard-Varet (1975, 1979); Arrow (1977)	Proposed here
C. *Areas of application*	Private decisions under perfect competition	(a) Allocation of standardised commodity in thin market with no technological externalities (b) Sealed-bid auction of indivisible good	Public goods decision	Decisions in decentralised organisations with asymmetric information	Collective decision	Private decisions, collective decisions, or decentralised organisations; technological externalities and asymmetric information permitted
D. *Acts generating externalities taken by*	No one – agents too small to affect market	Buyers and/or sellers, and centre	Centre	Centre	Centre	Agents and centre
E. *Sequence of information receipt, signals and acts*	Arbitrary; discrete, or continuous time collapsed to present through futures markets; markets clear by *tâtonnement*	(a) 1 Agents signal 2 Centre signals price and reward structure 3 Agents act (b) 1 Agents signal 2 Centre allocates and rewards	1 Agents signal 2 Centre decides	1 Agents signal 2 Centre signals reward structure 3 Agents and centre act	1 Agents signal 2 Centre decides	Sequence arbitrary; any finite number of stages of information receipt, signals and acts by agents and by centre
F. *Signal by an agent depends on*	Own preferences, market prices	Own demand or supply schedule	Own demand schedule	Own information and payoff function	Own payoff function	Own information and earlier acts and signals by other agents and centre

Table 13.1

G. *Agent signals information about*	Own demand or supply schedule at the margin	Own demand or supply schedule	Own demand schedule for the public good	Own information and payoff function	Own payoff function	Own payoff function, information shared with some other agents, information relevant to other agents' payoff functions
H. *Act by an agent depends on*	Own preferences, market prices	Own schedule and price and reward structures set by centre	No acts by agents	Own information and payoff function and centre's signal	No acts by agents	Own information and earlier acts and signals by other agents and centre
I. *Permissible information structures*	No unshared information on contingencies on which contracts are based	Only unknown is agent demand or supply	Only unknown is agent demand	Information among agents and centre mutually independent	Agents' information either (a) independent, or (b) in discrete case satisfies conditions for consistency of relevant set of inequalities	Agents' expectations of externalities they generate are known, as when each agent's information is public, signalled or independent of other agents' information, and private portion does not affect another agent's payoff
J. *Incentive transfers defined by*	None needed – voluntary exchange at market prices	Graphical equivalent of agent's addition to others' surplus	Graphical equivalent of marginal cost to others of agent's demand	Agent's contribution to centre's expectation of other agents' profits	(a) integral equal to agent's expectation of others' payoffs minus average of corresponding expectations of other agents (b) existence theorem	Expected externality (including later transfers) generated by action or signal of agent
K. *Properties of solution*	Dominant strategies Balanced budget	Dominant strategies Budget not balanced	Dominant strategies Budget not balanced	Dominant Strategies Budget not balanced	Bayesian Nash equilibrium Balanced budget	Bayesian Nash equilibrium Balanced budget

of paying the *expected externality* that one's signals and/or actions impose on the rest of the group. The externality concept, well established with respect to the direct effects of actions on others' welfare, is less obvious for announcements that influence the actions of the centre or, as in our general setting, influence the later acts and signals of other agents. That the budget can be balanced is also not obvious, but very welcome.

4.2 Generalization of the Value Function

If a function of some or all agents' observations is added to agent i's direct return v_i, it is clear that no choice by any agent will affect this added portion of v_i. Hence this has no incentive effect, and no change in incentive payments is needed. The expected-externality payments will in fact change, but only by functions of the agents' observations. Though such a change does not alter the agents' incentives, it may well affect implementability. Thus it is a real generalization to say that, if the agents' direct return functions v_i differ from the functions $\bar{v}_i$ only by functions of the agents' observations, then the expected externality payments based on the $\bar{v}_i$ will have the same incentive properties for the v_i as those based on the v_i themselves, and will be implementable if, in addition, the $\bar{v}_i$ satisfy the implementability conditions. In short, if implementable incentives exist for some $\bar{v}_i$ differing from the v_i only by functions of the agents' observations, then the same incentives will serve for the v_i.

Another generalization occurs if $\Sigma_i v_i$ is a function only of the agents' observations. Then the group objective is unaffected by the agents' choices, and no incentives are needed. Unfortunately, one cannot combine this possibility with others by a simple process of addition. If $v_i = v_i' + v_i''$ where $\Sigma_i v_i''$ is a function only of the agents' observations, then the group-optimal strategies are the same for v_i as for v_i', but the incentive payments based on v_i' will not serve for v_i because the individual v_i'' may depend on the agents' actions and hence have incentive effects.

These types of generalizations, and any others one might think of, can of course be combined if the problem is completely separable into subproblems, that is, if the agents can be partitioned into subsets and the direct returns of the agents in each subset are determined by the acts and observations of agents in that subset alone. Less trivial types of combinations, and other types of generalization, are a subject for further research.

4.3 Choice of Observations

Our formulation assumed that all agents will make the same observations r_i^m regardless of any decisions by any of the agents. That is, the agents have no choice about the observations they make; they have no choice of experiments. The definition of the expected-externality payments extends immediately to situations in which each agent may choose at each stage among possible quantities to observe; the observations will be revealed to him and perhaps to others at the next stage. The incentive properties of the payments remain the same. It appears that the only change needed in the implementability conditions is that they should hold for all possible choices of observations by the agents, provided the choices are public. If an agent's choice of observation is itself private, the situation is more complicated. We have not yet examined either case in detail.

4.4 Differential Transmission of Information

For convenience and specificity, our formulation requires that all agents be informed about all acts and receive all signals. The derivation of the expected-externality transfers in the multistage case (Section 3) does not require this, however, and in fact is completely general. The conditional expectations (8) and (10) have simply to reflect whatever information agent i actually has at the stage in question. There must, of course, be a process with finitely many stages in which one agent's choices at any stage do not become known to other agents until a later stage, if at all. Implementability in a completely general situation will again require that information observed by an agent but not monitored by the central authority be irrelevant in determining expected externalities, for instance by not affecting any other agent's direct return and meeting an independence condition.

Our formulation can allow for a variety of situations in which there is a choice of how much information to signal and to whom. If there is a cost of transmission to an additional agent, then the team presumably would balance the benefits of better co-ordination against the costs of having an additional receiver. The group can do the same. The easiest

way, given our formulation, would be to include the choice of who receives signals as an agent action.

If there is a cost of having the central authority monitor the information needed to determine the transfer payments, then obviously a team can outperform a group by the magnitude of this cost if the team can transmit the information it needs costlessly. This is not really an appropriate comparison. The question we are addressing in this chapter is whether a decentralized incentive mechanism for a group can reproduce the outcome for the team, disregarding any costs of administering that incentive mechanism. The costs of having the central authority monitor the relevant acts, signals, and observations, are transaction costs, playing a role equivalent to the cost of filling out the checks when making transfer payments.

4.5 Example of Efficient Incentives When Actions Convey Information

Information can be conveyed through actions as well as signals. A group can take advantage of this possibility. This has already been reflected in Section 3 in the incorporation of signals into acts and in the conditioning of agent i's stage-m probabilities on his stage-m information q_i^m. Beyond this, it is even possible for a properly co-ordinated group to gain by 'distorting' its actions to convey information. Our approach can provide incentives supporting a collective strategy of this type provided only that the individual agents' strategies are agent-by-agent optimal for the group.

Let us create a new, slightly more elaborate externality example. The Upstream polluter has traditionally dumped 10 units per period. Pollution remains in the water for one period. The cost to Upstream for dumping an amount $d \leqslant 10$ is $(10-d)^2$.

The Downstream firm now comes on river. Its sensitivity to pollution, s, is unknown; the shared prior probability distribution on s is that it is equally likely to have the value 2 or 20; s is constant from period to period. In each period, Downstream suffers a cost equal to the amount of pollution times its sensitivity. Downstream has a means to reduce its cost. It can decide to clean up the pollution. Cleanup costs 9.2 and removes 50 per cent of the pollution dumped.

In an analogous problem we saw that incentive-based externality payments could achieve the team optimum if Downstream could signal its sensitivity. In this problem, no signals are permitted. Upstream can only draw inference about s from the actions of Downstream. Still the team optimum, where the team is subject to the same communication limitations, is achieved through the use of appropriate incentives.

Let us observe what happens in a two-period world. We start at stage 2 with ten units of pollution in the water. The sequence of actions will now unfold:

Stage 2 – Downstream observes s. It then decides whether or not to undertake cleanup in period 1.
Stage 1 – Upstream decides how much to dump, d.
Stage 0 – Downstream decides whether to clean up in period 2.

With a bit of bookkeeping and a dash of differentiation, we can compute the team optimum:

Stage 2 – Clean up if $s=20$; do not clean up if $s=2$.
Stage 1 – Dump 5 if Downstream did clean up at stage 2. Dump 9.5 if Downstream did not clean up at stage 2.
Stage 0 – Clean up independent of the value of s.

The team-optimal strategy has its interesting aspects. At stage 2, if $s=2$, Downstream does not clean up even though the cost of cleanup, 9.2, is less than the pollution damage avoided, 10. The reason, as we suggested at the outset, is that actions convey information. With signalling not allowed, it is beneficial on net to take an action that yields a lower payoff in and of itself but is more informative. (If Downstream did clean up, Upstream at stage 1 could not distinguish between $s=20$ and $s=2$.) Given that $s=2$, there are two local optima for Upstream's stage 1 dumping decision. The local optimum that assumes cleanup at stage 0 is superior.

If the group is to achieve the team optimum, we must employ the externality incentives for that pair of strategies. Let us start at the end. No incentive is needed at stage 0; Downstream will do whatever is optimal since its action neither conveys information that could inform a future action nor affects the payoff to another party. At stage 1, Upstream 'knows' the value of s from Downstream's action at stage 2. Cleanup implies $s=20$; no cleanup implies $s=2$. Its incentive payment as a function of a dumping decision d is merely the net amount Downstream receives from d. For example, with $s=2$, Downstream, which cleans up at stage 0, receives $-(2d/2+9.2)$: this is the incentive payment to Upstream. Upstream's decision at stage 1 is thus to maximize $-(10-d)^2-2d/2-9.2$. The optimum value of d is thus 9.5. The net payoff to Upstream is -18.45.

A parallel calculation for $s=20$ yields the incentive payment to Upstream at stage 1 of $-(20d/2+9.2)$. When Upstream optimizes, it dumps 5 units; its net payoff is -34.2.

We are now in a position to fold back to stage 2 and to provide appropriate incentives for the Downstream cleanup decision at that stage. It must receive Upstream's net payoff at stage 1, contingent on whether or not there was cleanup. Quite simply, Downstream receives an incentive payment of -34.2 if it cleans up, and of -18.45 if it does not. In response to these incentives, it will choose to clean up when $s=20$, but not when $s=2$. To summarise, the net transfers are:

Payment From Downstream to Upstream	
Cleanup at Stage 2	*No Cleanup at Stage 2*
$34.2-(20d/2+9.2)$	$18.45-(2d/2+9.2)$
$=25-10d$	$=9.25-d$

4.6 Discounting and Chance Termination

The analysis in this paper is directed to problems with a fixed, finite number of stages. Casual observation suggests that many real groups function over a considerable period of time, often without a fixed terminal date. We have not been concerned with the mechanics of discounting. If all agents agree on the appropriate discount rate, we need merely date all consequences and transfers and discount them at the agreed-on rate.

If there is chance termination, we believe that all of our results go through with only minor complications so long as no player has unsignalled private information about the likelihood of termination. If settling up is not possible at the time of termination, then the payments conditional on non-termination will have to be accordingly higher, to make the expected payment equal the expected externality.

In some situations we may need agents to take actions that foster termination because on a discounted expected-value basis they can foresee that the total group payoff will be diminishing. We might mistakenly conclude that problems would arise if post-termination settlement was impossible and if individuals could be 100 per cent sure of terminating. Presumably, they would terminate when their individual payoff would start to decline. The way around that problem is to simply pay them a sufficient transfer (that is, the expected externality from continuation divided by the probability of termination if they act optimally) should the game continue. To compensate, some lump-sum amount could be charged to the individual from the start. Thus, all ex ante Pareto-optimal outcomes are achievable even though individuals can assure themselves of a zero-payoff future by terminating the game.

Situations that never terminate present no problem in principle, we believe, so long as there is discounting. Delaying a payment by a further period merely multiplies the magnitude of that payment by one plus the interest rate. Provided settling up is not delayed indefinitely, individuals can be rewarded on a discounted expected-value basis for behaviour that helps the group. Likewise, we foresee no problems if there is uncertainty about the settling-up date or if, as seems likely, there is only partial settlement in any one period. As with credit cards, there may always be some float. Because many real-world situations persist over long periods of time, we believe consideration of situations without foreseeable termination will be of value.

4.7 Bargaining and Public Goods

Usual bargaining processes do not provide adequate incentives for honest revelation. Moreover, they may break down when a mutually beneficial agreement would be possible. Section 1.8 illustrated how expected externality methods can be employed, as in Chatterjee, Pratt and Zeckhauser (1978), to assure sale of an object whenever it would be efficient.

Consider now a public-goods decision. A project costing c is contemplated. Each agent i must declare a value x_i, which the government will assume is the true value to him of the project. The government will undertake the project if and only if $\Sigma_i x_i > c$. Without loss of generality, we may take $c = 0$. (Let r_i be agent i's value in excess of c/n, and x_i the value he declares in excess of c/n.) If the $\tilde{r}_i$ are independent, the subsidy defined by Theorem 1 turns out to be

$$s_i = \int_{-x_i}^{\infty} y \, dH_i(y) \tag{25}$$

where H_i is the cumulative distribution function of $\Sigma_{j \neq i} \tilde{r}_j$, the total value to the other agents. The subsidy is a single-humped function of x_i, with maximum at $x_i = 0$, that is, when the agent quotes a value equal to the cost per person. It may be derived as follows. Agent i's direc' turn is

$$v_i = r_i S(\textstyle\sum_i x_i) \tag{26}$$

where $S(u) = 1$ if $u > 0$; otherwise $S(u) = 0$. The group optimum clearly occurs when all agents declare $x_i = r_i$. Agent i's expectation of agent j's direct return if the other agents follow the group optimal strategies is ($j = i$ allowed)

$$s_{ij} = \mathrm{E}\{\tilde{r}_j S(x_i + \textstyle\sum_{j \neq i} \tilde{r}_j) | r_i\}. \tag{27}$$

Thus agent i's expected externality is

$$s_i = \mathrm{E}\{\textstyle\sum_{j \neq i} \tilde{r}_j S(x_i + \sum_{j \neq i} \tilde{r}_j) | r_i\}, \tag{28}$$

which is s_i as defined above.

It is easy to verify directly that $s_i(x_i)$ provides the asserted incentive. Agent i's subsidy plus expected direct return, if his value is r_i and he announces x_i, is

$$s_i + s_{ii} = \int_{-x_i}^{\infty} (y + r_i) dH_i(y). \tag{29}$$

Since the integrand is negative for $y < -r_i$ and positive for $y > -r_i$, the integral is maximized when the lower limit is $-r_i$, that is, $x_i = r_i$.

We remark that the subsidy s_i differs only by a constant from agent i's expectation of what he would receive under a demand-revealing scheme (see Section 1.2). Under the usual variant of that scheme, agent i must pay $|y|$, where $y = \Sigma_{j \neq i} x_j$, if y and $x_i + y$ have opposite signs (x_i changes the decision). This scheme has the very strong property of making honesty dominant, but does not balance the budget. To see the relation with s_i, observe that if the other agents are honest, then $\tilde{y}$ has cumulative distribution function H_i and agent i's expected payment is

$$\int_0^{\infty} y dH_i(y) - \int_{-x_i}^{\infty} y dH_i(y) = \int_0^{\infty} y dH_i(y) - s_i. \tag{30}$$

4.8 Application to Collective Decision Problems

The public-goods problem above has been generalized to the choice of level of expenditure, and still further to arbitrary collective decisions with informational decentralization. Even the latter, general problem is still a special case of our model with just one stage.

Specifically, let r_i describe agent i's true preferences and a_i be his report. Let $v_i(a_1, \ldots, a_n; r_i)$ be his value for the collective decision that the central authority will make upon receiving the reports $a_1, \ldots, a_n$. Then, for instance, agent 1's expected externality to agent j is $s_{1j} = E\{v_j(a_1, \tilde{a}_2, \ldots, \tilde{a}_n; \tilde{a}_j) | r_1\}$. Note that this is the expectation of the revealed externality $v_j(a_1, r_2, \ldots, r_n; r_j)$ if all other agents are honest. Again the subsidy $s_i = \Sigma_{j \neq i} s_{ij}$ is agent i's expectation of what he would receive under a demand-revealing scheme. If the $\tilde{a}_i$ are independent, as in Theorem 2(F), (p. 455), each agent's expectation is known to the centre and the subsidies are implementable.

5 CONCLUSION

The goal of this investigation was to determine under what circumstances a group of self-interested agents, who observe private information and send signals that may not be verifiable, and constrained to balance the budget, could do as well as a fully co-operating team with identical communication possibilities. Equivalent performance turns out to be possible in substantially more general circumstances than those heretofore identified. Through a process we call incentive-based decentralization, each agent receives a payment equal to the expected externality conveyed to the group by his actions, including the signals he sends, and contributes to the payments to others to balance the budget. Then self-interested actions support a sequential Bayesian Nash equilibrium that is a group optimum. These results provide a happy complement to earlier methods that make honest revelation a dominant strategy but fail to balance the budget.

For the expected-externality incentive payments defined by our methods to be implementable, it suffices that each element of each agent's information either become public, be signalled by him, or be independent of the information of others; and that his non-public information does not directly affect another agent's value. (For various cases, we give precise sufficient conditions of increasing levels of generality.)

In multistage situations, which have hardly been addressed in the incentive compatibility literature, an agent's payments must take account of an additional factor: the effect his actions and signals have on the contributions he must make – because of budget-balance requirements – towards the incentive payments others receive for their later actions and signals. Fortunately, appropriate payments can be defined, though the computations may be complex, and sequential equilibrium can be maintained. Thus, incentive-based decentralization succeeds in multistage contexts, an important feature for the ultimate practical application of this methodology.

NOTES

1. In a footnote, Arrow related his work to that of d'Aspremont and Gérard-Varet, with which he became familiar only after his own was completed. Here we note that ours relates to both his and theirs in essentially the same way, and that, like Arrow, we learned of our closest ancestry only after producing our offspring, Chatterjee, Pratt, and Zeckhauser (1978), and Pratt and Zeckhauser (1980).
2. Some colleagues, nevertheless, have suggested, sometimes quite insistently, that many stages could be reduced easily to one. Several arguments, any one of which amply suffices to demolish this view, are worth sketching, because they clarify different aspects of the situation and because the view seems so hard to dislodge. First, in the 'normal' (one-stage) form of a multistage problem, an agent's expected externality becomes his ex ante expectation of the externality that his strategy provides to others. This may never become known to the centre even though the expectations needed for our multistage incentives do become known. Thus the normal-form expected-externality incentive may not be implementable when the multistage expected-externality incentive is implementable. Second, even if the normal-form expected-externality incentive is implementable, it will in general be different from the multistage one and might well be less palatable, especially ex post. Third, the multistage mechanism employs expectations conditional on information revealed to or learned by an agent along the way, and exploits the stages so that the incentive effects of budget balancing at each stage are offset at the previous stage. Though its properties could, of course, be verified in normal form, the multistage dynamics provide its motivation and derivation. Could the telescoping formula be discovered within a static framework? No one has done so in the many years since the independent, one-stage case was discovered. Fourth, even if a one-stage reduction of the problem could be found that captured our multistage mechanism as a naturally defined expected-externality mechanism, it seems unlikely to provide a simpler derivation in problems that naturally occur in multiple stages, although it might provide theoretical simplification in the sense of unification. Fifth, the sequential type of equilibrium achieved by our approach in extensive form is in an important and desirable way more restrictive than Bayesian equilibrium in normal form. Indeed, even in the limited one-stage case considered by Arrow, as mentioned earlier, he regarded the difference as significant, and it is far more significant when agents choose acts after receiving messages in one or more stages.
3. As already mentioned, the behaviour we seek to induce need not be fully

efficient, a global group optimum. It suffices that each agent's strategy optimize for the group, given any constraints he and the group are under and given the others' strategies. Furthermore, our incentives make the desired strategy a Bayesian Nash equilibrium for each agent at each stage, not merely initially.

There has been recent interest in how to structure games so that individuals will not drop out as information unfolds. See, for example, Holmstrom and Myerson (1983). In the expected externality formulation, a player's utility may diminish as the game passes through stages. To deter the defection of players, there might be an external enforcer, say the government. Or there might be a bonding mechanism, such as a requirement that each player put on deposit a sum that exceeds the maximum drop in his expected value at any stage in the game.

4. Theorem 6 of d'Aspremont and Gérard-Varet would coincide with our Theorem 2(F) the most special case of our theorem for a single stage, were we to restrict that theorem to situations in which individuals transmit information merely about their own preferences and do not take actions. In the discrete case, d'Aspremont and Gérard-Varet (1979) prove the existence of the required transfers under a weaker condition that unfortunately seems to resist an intuitive interpretation. Arrow's related results are described in subsections 1.3 and 1.7.

REFERENCES

Arrow, K. J. (1951) *Social Choice and Individual Values* (New York: Wiley). 2nd edn, 1963.

Arrow, K. J. (1977). 'The Property Rights Doctrine and Demand Revelation Under Incomplete Information', Discussion Paper 580, November 1977 (Harvard Institute of Economic Research: Harvard University). Published in Michael Boskin (ed.) (1979) *Economics and Human Welfare* (New York: Academic Press) pp. 33–9.

Arrow, K. J. and G. Debreu (1954) 'Existence of Equilibrium for a Competitive Economy', *Econometrica*, 22: 265–90.

Arrow, K. J. and F. Hahn (1971). *General Competitive Analysis* (San Francisco: Holden-Day).

d'Aspremont, C. and L.-A. Gérard-Varet (1975) 'Individual Incentives and Collective Efficiency for an Externality Game with Incomplete Information', CORE Discussion Paper 7519 (Catholic University of Louvain).

d'Aspremont, C. and L.-A. Gérard-Varet (1979) 'Incentives and Incomplete Information', *Journal of Public Economics*, 11: 25–45.

Chatterjee, K., J. W. Pratt and R. Zeckhauser (1978) 'Paying the Expected Externality for a Price Quote Achieves Bargaining Efficiency', *Economics Letters*, 1: 311–13.

Clarke, E. H. (1971) 'Multipart Pricing of Public Goods', *Public Choice*, 11: 17–33.

Green, J. R. and J.-J. Laffont (1979) *Incentives in Public Decision-Making* (Amsterdam: North-Holland).

Groves, T. (1973). 'Incentives in Teams', *Econometrica*, 41: 617–31.

Groves, T. and J. Ledyard (1977) 'Optimal Allocation of Public Goods: A Solution to the "Free Rider" Problem', *Econometrica*, 45: 783–809.

Groves, T. and M. Loeb (1979) 'Incentives in a Divisionalized Firm', *Management Science*, 25: 221–30.

Holmström, B. and R. Myerson (1983) 'Efficient and Durable Decision Rules With Incomplete Information', *Econometrica*, 51: 1799–815.

Hurwicz, L. (1959) 'Optimality and Informational Efficiency in Resource Allocation Processes', in K. J. Arrow, S. Karlin, and P. Suppes (eds.) *Mathematical Methods in the Social Sciences* (1959) (Stanford: Stanford Univ. Press) pp. 27–46.

Hylland, A. (1980). 'Collective Decisions and Individual Incentives: Review of the Literature', Discussion Paper 82D (Kennedy School of Government: Harvard University).

Laffont, J.-J. and E. Maskin (1979) 'A Differential Approach to Expected Utility Maximizing Mechanisms', in J.-J. Laffont (ed.) *Aggregation and Revelation of Preferences* (Amsterdam: North Holland) pp. 289–308.

Loeb, M. (1977) 'Alternative Versions of the Demand-Revealing Process', in T. N. Tideman (ed.), *Public Choice*, 29.

Marschak, J. and R. Radner (1972) *Economic Theory of Teams* (New Haven: Cowles Foundation and Yale University Press).

Pigou, A. C. (1960) *The Economics of Welfare*, 4th edn. (London: Macmillan).

Pratt, J. W. and R. J. Zeckhauser (1980) 'Incentive-based Decentralization: Expected-Externality Payments Induce Efficient Behavior in Groups', Discussion Paper 83D (Kennedy School of Government: Harvard University).

Schultze, C. L. (1977) *The Public Use of Private Interest* (Washington, DC: The Brookings Institution).

Tideman, T. N. (ed.) (1977) 'A Collection of Papers on the "Demand-revealing Procedure",' *Public Choice*, 29.

Tideman, T. N. and G. Tullock (1976). 'A New and Superior Process for Making Social Choices', *Journal of Political Economy*, 84: 1145–59.

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