

Townsend R & Harris M. "Allocation Mechanisms, Asymmetric Information, and the Revelation Principle" in *Issues in Contemporary Microeconomics and Welfare*, ed. George R. Feiwel (SUNY Press, 1975), 379-394.

- Rao, R., and D. Rutenberg (1979) 'Preempting an Alert Rival: Strategic Timing of the First Plant by Analysis of Sophisticated Rivalry', *Bell Journal of Economics*, 10: 412–28.
- Reinganum, J. (1981) 'Dynamic Games of Innovation', *Journal of Economic Theory*, 25: 21–41.
- Reinganum J. (1982a) 'A Dynamic Game of R&D: Patent Protection and Competitive Behavior' *Econometrica*, 50: 671–88.
- Reinganum, J. (1982b) 'Patent Races with a Sequence of Innovations' (mimeo).
- Reinganum, J., and N. Stokey (1981) 'Oligopoly Extraction of a Non-renewable, Common Property Resource: The Importance of the Period of Commitment in Dynamic Games' (mimeo).
- Roth, A., and F. Schoumaker (1981) 'Expectations and Reputations in Bargaining: An Experimental Study', *American Economic Review*, 73: 362–72.
- Rubinstein, A. (1979) 'Strong Perfect Equilibrium in Supergames', *International Journal of Game Theory*, 9: 1–12.
- Schelling, T. (1960) *The Strategy of Conflict*, Cambridge, Mass.: Harvard University Press.
- Scherer, F. (1980) *Industrial Market Structure and Economic Performance*, 2nd edn, Chicago: Rand McNally College Publishing
- Schmalensee, R. (1982) 'Product Differentiation Advantages of Pioneering Brands', *American Economic Review*, 72: 349–65.
- Shaked, A., and J. Sutton (1982) 'Relaxing Price Competition Through Product Differentiation', *Review of Economic Studies*, 49: 3–13.
- Simon, H. (1965) *Administrative Behaviour*, New York: Free Press.
- Selten, R. (1978) 'The Chain-store Paradox', *Theory and Decision*, 9: 127–59.
- Smale, S. (1980) 'The Prisoner's Dilemma and Dynamical Systems Associated to Non-Cooperative Games', *Econometrica*, 48: 1617–34.
- Spence, A. (1977) 'Entry, Capacity, Investment and Oligopolistic Pricing', *Bell Journal of Economics*, 8: 534–44.
- Spence, A. (1979) 'Investment Strategy and Growth in a New Market', *Bell Journal of Economics*, 10: 1–19.
- Spence, A. (1981) 'The Learning Curve and Competition', *Bell Journal of Economics*, 12: 49–70.
- Spence, A. (1982) 'Cost Reduction, Competition and Industry Performance' (mimeo), *Econometrica*, 52: 101–22.
- Stigler, G. (1964) 'A Theory of Oligopoly', *Journal of Political Economy*, 72: 44–61.
- Sylos-Labini, P. (1962) *Oligopoly and Technical Progress*, trans. E. Henderson, Cambridge, Mass.: Harvard University Press.
- Tversky, A. (1972) 'Elimination by Aspects: A Theory of Choice', *Psychological Review*, 76: 31–48.
- Tversky, A., and D. Kahnemann (1974) 'Judgement Under Uncertainty: Heuristics and Biases', *Science*, 185: 1124–31.
- Williamson, O. (1975) *Markets and Hierarchies: Analysis and Anti-trust Implications*, New York: Free Press.
- Williamson O. (1981) 'The Modern Corporation: Origins, Evolution, Attributes', *Journal of Economic Literature*, 19 (December): 1537–68.

11 Allocation Mechanisms, Asymmetric Information and the 'Revelation Principle'

MILTON HARRIS
and ROBERT M. TOWNSEND*

INTRODUCTION

The purpose of this chapter is to explain a new approach for predicting both the allocation of resources and the resource allocation mechanism in certain environments in which agents are asymmetrically informed prior to any trading.¹ We illustrate this approach by applying it to a simple pure-exchange environment in which an information asymmetry is present.

The central element of this approach (which is described in more detail below) is to define the concept of an optimal resource allocation mechanism and to characterize such optimal mechanisms and their associated optimal allocations for given economic environments. There are two reasons for approaching the problem in this way, that is, searching for optimal resource allocation mechanisms. The first is simply that one would like to have a theory that *explains* observed mechanisms (or processes or institutional arrangements – we use these terms synonymously). For example, one might wish to explain why auctions are used in certain environments, but not in others.² The second is that one would like to have a theory that explains the final allocation of resources in environments with asymmetric information, and to do so, we argue that one must begin with an explicit consideration of mechanisms. To sketch this argument briefly, let us start with the premise that in any economic environment, the observed allocation of resources actually is achieved by some mechanism. We then assert that in some asymmetric information environments *there are allocations consistent with the resource constraints* (i.e. *technologically feasible*) *which, nevertheless, cannot be achieved by any mechanism.*

* This research has benefited from the support of the National Science Foundation.

That is, we assert that the presence of information asymmetries imposes certain incentive constraints on achievable allocations.³ The point here is that these constraints can only be revealed by an explicit consideration of the processes by which allocations are achieved. Of course, once these constraints are revealed, one can contemplate generating optimal allocations in the usual way as solutions to programming problems that incorporate the constraints. Obviously one way of making sure that information constraints are accounted for in an economic model is simply to analyse the equilibrium of some particular (and explicit) mechanism. This approach is certainly suitable for positive purposes if one believes that the chosen mechanism is a good model of actual arrangements in the environment of interest, and if one is not interested in explaining these arrangements. Yet if one does seek to explain observed arrangements, or if one is interested in normative implications, then one would like to establish that the constraints on allocations of the chosen mechanism cannot be circumvented by some alternative process. This leads us to consider a fairly broad class of available mechanisms.⁴

The approach we describe here uses what has now come to be known as the 'Revelation Principle'.⁵ This principle, or results similar to it, has been developed by Harris and Townsend (1978, 1981), Holmstrom (1978), and Myerson (1979). It can be stated simply as:

The Revelation Principle. Any equilibrium allocation of any mechanism can be achieved by a truthful, direct mechanism.

By a *direct mechanism*, we mean a game in which all agents first simultaneously declare values for whatever parameters they have observed, for example, parameters describing their own tastes, etc. After these 'messages' or 'signals' are sent, some allocation is effected as a function of the declarations of all players. This *allocation rule* is specified in advance. Players in a direct mechanism need not tell the truth about their observed parameters. In a *truthful*, direct mechanism, however, there is an equilibrium in which all players do tell the truth: a truthful mechanism, then, is one in which each player is given an incentive (by the allocation rule) not to lie, provided that he expects all other players to tell the truth.

The power of the Revelation Principle is that it enables one to limit his search for optimal mechanisms to direct mechanisms without fear of ignoring a more complicated mechanism that could have produced a better outcome. Specifying a mechanism can, in general, be quite complicated, involving a specification of what strategies are feasible for each agent at each stage, what each agent knows at each stage, and how the

final allocation depends on the whole history of signals of the agents. A direct mechanism can, however, be completely specified by its allocation rule. This is simply a function from the set of values of observable parameters to the set of feasible allocations.

The Revelation Principle also implies that we can restrict attention to direct mechanisms in which truth-telling is an equilibrium. This imposes a set of constraints on the allocation rules that guarantee that, for each agent and for each value of his observed parameters, it is optimal to tell the truth given that all other agents are telling the truth. These constraints are generally called 'self-selection' or, following Hurwicz (1972), 'incentive compatibility' conditions. Thus, using the Revelation Principle, an optimal mechanism, and its associated equilibrium allocation, can be found by choosing an *allocation rule* that maximizes some social welfare function (e.g. a weighted average of the utilities of the players) subject to technological feasibility conditions and incentive compatibility conditions. In effect, then, one can search for optimal allocations directly. An example of how this is done is given in the third section.⁶

In the second section, we attempt to motivate the general results, primarily the Revelation Principle. Readers are referred elsewhere for proofs. In the third section, we analyse a specific two-person, two-good, pure exchange environment using the results of the second sections. The fourth section provides a summary and conclusion.

GENERAL RESULTS

First let us specify the general type of economic environment to which the results will apply. Suppose there are a finite number of economic agents, say N , indexed by $i = 1, \dots, N$. Further suppose that there is a set A of technologically feasible allocations. An element a of A is a vector that specifies each agent's allocation bundle. The set A will incorporate constraints due to the technology of production and exchange and due to resource availability. In order to introduce asymmetric information into the environment, we shall assume that each agent i may *privately* observe the value of a parameter θ_i which affects his tastes.⁷ We model agent i 's lack of information about the parameters of other agents by assuming that i has a well-defined joint prior distribution over $(\theta_j)_{j \neq i}$. This prior may depend on the observed value of i 's parameter, θ_i . Finally, we denote by $U_i(a, \theta_i)$ the utility of agent i for an allocation a if his parameter has value θ_i . When i 's parameter value is θ_i , i is said to be of *type* θ_i .

Our next step is to define more carefully the concept of a mechanism. In this chapter we will, as in Myerson (1979), define a mechanism to be what game theorists call a 'game in normal form'. A game in normal form is a particular way of formalizing the intuitive notion of a game. This formalization specifies the set of signals each agent can send and an outcome that depends on the signals sent.⁸

To be somewhat more formal, a normal-form game specifies a set of feasible signals S_i for each agent i and an allocation rule F which associates with each vector of signals (s_1, \dots, s_N) in $S = S_1 \times \dots \times S_N$ an allocation $F(s_1, \dots, s_N)$ in A .⁹ We may now define a *mechanism* as any normal-form game, that is, signal sets S_i and allocation rule F .

Our next task is to define what is meant by an equilibrium of a mechanism. This is simply an hypothesis concerning the way we expect players to behave and the outcome that will result. The equilibrium concept we use here is called the Bayesian equilibrium (by Harsanyi, 1967-68) because players' strategies are based on their prior beliefs about the 'types' of the other agents.

The first point to recognize in defining equilibrium is that each player i will choose his signal s_i from his signal set S_i based on the value of his parameter, that is, on his type θ_i . Thus player i 's *strategy* is a function that depends on θ_i and whose value is a signal in S_i . Let us denote i 's strategy by $\sigma_i(\theta_i)$.

The second point to recognize is that i 's optimal strategy, that is, his optimal signal as a function of his type, depends on what signals he believes other players will send. Other players' signals are, in turn, determined by their strategies, σ_j , and their types, θ_j . Therefore, player i 's beliefs about player j 's signal reflect the strategy that i believes j will use and i 's prior on θ_j . In a Bayesian equilibrium, it is assumed that each player i chooses his best strategy *given* the strategies of the other players and *given* i 's prior beliefs about their types.

More formally, a vector of strategies $(\sigma_1^*, \dots, \sigma_N^*)$ is a Bayesian equilibrium of the mechanism defined by the signal sets S_1, \dots, S_N and the allocation rule F if, for each player i and each possible value of his type θ_i , the signal $s_i^* = \sigma_i^*(\theta_i)$ maximizes i 's expected utility given that he believes that each other agent j will be using strategy σ_j^* and given his beliefs about θ_j . In equation form, for each player i and θ_i , $\sigma_i^*(\theta_i)$ solves:

$$\max E_i \{ U_i [F(\sigma_1^*(\theta_1), \dots, \sigma_{i-1}^*(\theta_{i-1}), s_i, \sigma_{i+1}^*(\theta_{i+1}), \dots, \sigma_N^*(\theta_N)), \theta_i] \mid \theta_i \} \quad s_i \text{ in } S_i$$

where $E_i(\mid \theta_i)$ denotes i 's expectation over $\theta = (\theta_1, \dots, \theta_N)$ given that his type is θ_i and using his prior beliefs. Naturally, the equilibrium strategies result in an equilibrium allocation that depends on the actual vector of agent types, θ , that is:

$$a^*(\theta) = F[\sigma^*(\theta)]$$

where $a^*(\theta)$ is the equilibrium allocation and $\sigma^*(\theta) = [\sigma_1^*(\theta_1), \dots, \sigma_N^*(\theta_N)]$.

We are now in a position to motivate part of the Revelation Principle. In particular, we can show how to derive a *direct* mechanism from any given mechanism. Consider a mechanism represented by signal sets S_i ($i = 1, \dots, N$) and allocation rule F . Think of the allocation rule F as being a computer program that uses the signals s_1, \dots, s_N as inputs and produces an allocation $a = F(s_1, \dots, s_N)$ as output. Suppose that $\sigma^* = (\sigma_1^*, \dots, \sigma_N^*)$ is an equilibrium of this mechanism. Instead of having each player compute his optimal signal, based on his type, then feeding this signal into the computer to compute the allocation, suppose we programme the computer to compute signals using σ^* and have each player simply input a value of his parameter. The computer could then use the signals that result from this calculation to compute an allocation using the allocation rule F . This would save the players from computing their optimal signals; the mechanism embodied in the computer program would do it for them. The result is a new mechanism in which players send signals that are interpreted as declared values of their parameters instead of the, possibly much more complicated, signals in the sets S_i . Thus, in the new mechanism, the signal sets are the sets of possible values of the parameters, that is, the new mechanism is a direct mechanism!

Now, what is the relationship between the direct mechanism constructed in the previous paragraph and the original mechanism? In particular, does it yield the same equilibrium outcome $a^*(\theta) = F[\sigma^*(\theta)]$ as the one induced by the equilibrium strategies σ^* of the original mechanism? The answer to this question is, happily yes. To see this, suppose player i believes that, in the new, direct mechanism, all the other players will 'tell the truth', that is, the strategy of agent j is $\psi_j(\theta_j) \equiv \theta_j$ (we use ψ_j to distinguish strategies in the direct mechanism from those of the original mechanism). Now if player i is of type θ_i , and he reports θ_i , then he believes the resulting allocation (as a function of θ) will be:

$$a^*(\theta) = F[\sigma^*(\theta)]$$

just as in the original mechanism. On the other hand, reporting some other value of his parameter, say $\hat{\theta}_i$, would result in the outcome:

$$F[\sigma_1^*(\theta_1), \dots, \sigma_{i-1}^*(\theta_{i-1}), \sigma_i^*(\hat{\theta}_i), \sigma_{i+1}^*(\theta_{i+1}), \dots, \sigma_N^*(\theta_N)]$$

But the expected utility of this outcome is lower than that of $a^*(\theta)$ by definition of σ_i^* , that is, when player i is of type θ_i , $\sigma_i^*(\theta_i)$, not $\sigma_i^*(\hat{\theta}_i)$, maximizes:

$$E_i \{ U_i [F(\sigma_1^*(\theta_1), \dots, \sigma_{i-1}^*(\theta_{i-1}), s_i, \sigma_{i+1}^*(\theta_{i+1}), \dots, \sigma_N^*(\theta_N)), \theta_i] \mid \theta_i \}$$

over s_i . Reporting $\hat{\theta}_i$ in the direct mechanism would be just like *lying to himself* in the original mechanism, that is, acting as if he were some other type. This shows that player i will report truthfully in the direct mechanism, provided he believes that everyone else will also. In game theory language, telling the truth ($\psi_i^*(\theta_i) \equiv \theta_i$) is a Bayesian equilibrium of the direct mechanism. Moreover, as mentioned above, the equilibrium outcome corresponding to these equilibrium strategies is simply $a^*(\theta)$ when the vector of types is θ . This is exactly as in the original mechanism.

The above argument is the essential idea behind the Revelation Principle which we repeat here for convenient reference;

The Revelation Principle. Any equilibrium allocation of any mechanism can be achieved by a truthful, direct mechanism.

As mentioned above, this result is an extremely useful tool if one is searching for an optimal mechanism or simply an optimal allocation in an asymmetric information environment. This is because the Revelation Principle implies that equilibrium allocations of any mechanism must satisfy certain self-selection (or incentive compatibility) constraints. Suppose that $a^*(\theta)$ is an allocation of some mechanism. We know from the Revelation Principle that $a^*(\theta)$ is also the *truthful* equilibrium of a direct mechanism. What does it mean for $a^*(\theta)$ to be the truthful equilibrium of a direct mechanism? It means first of all that the direct mechanism has an allocation rule G which gives each player an incentive to reveal his type truthfully provided everyone else behaves similarly. Formally, we

must have:

$$E_i \{ U_i [G(\theta_1, \dots, \theta_N), \theta_i] \mid \theta_i \} \\ \geq E_i \{ U_i [G(\theta_1, \dots, \theta_{i-1}, \hat{\theta}_i, \theta_{i+1}, \dots, \theta_N), \theta_i] \mid \theta_i \} \quad (11.1)$$

for each i , θ_i , and $\hat{\theta}_i$

Condition (11.1) states that for any player i and any two values of his type θ_i and $\hat{\theta}_i$, if his true type is θ_i , then he prefers the allocation associated with his reporting θ_i to the one associated with his reporting $\hat{\theta}_i$, provided all other players are reporting truthfully. This last caveat is embodied in the fact that the expectations in (11.1) are taken with respect to player i 's prior beliefs about the *true* θ_j 's. Equation (11.1) is almost the self-selection condition we seek.

The second step in the argument is to recall that if everyone reports truthfully, the direct mechanism represented by G will result in the allocation $a^*(\theta)$, for any vector of types, θ . Thus we must have:

$$G(\theta) \equiv a^*(\theta)$$

Substituting a^* for G in (11.1) gives the self-selection (SS) conditions:

$$E_i \{ U_i [a^*(\theta_1, \dots, \theta_N), \theta_i] \mid \theta_i \} \geq \\ E_i \{ U_i [a^*(\theta_1, \dots, \theta_{i-1}, \hat{\theta}_i, \theta_{i+1}, \dots, \theta_N), \theta_i] \mid \theta_i \} \\ \text{for all } i, \theta_i, \text{ and } \hat{\theta}_i$$

Thus the (SS) conditions are satisfied by the outcome of *any* mechanism. These conditions in effect become constraints on technologically feasible allocations as noted in the introduction to this chapter.

Finally, note that the direct mechanism whose truthful equilibrium allocation implements the original allocation a^* is simply the direct mechanism whose allocation rule $G = a^*$. This makes it trivial to construct a direct mechanism that implements a given allocation a^* , provided, of course, that a^* satisfies the (SS) conditions. Moreover, one can search for an *optimal* mechanism by searching for an optimal allocation in the space of allocations that satisfy the (SS) conditions using standard mathematical programming techniques.

AN EXAMPLE

In the remainder of the chapter we focus on a simple, pure risk-sharing example consisting of two agents, one consumption good and two states of nature.¹⁰ Agent 1, the informed agent, is presumed to know the true probability that each state will occur, and this is known by the uninformed agent. Agent 2, the uninformed agent, has a prior distribution over these probabilities, and this prior is known by the informed agent. Both agents are assumed to know all other aspects of the environment.

Endowments of the good for each agent for each state are exogenously fixed. Let x_s denote the total endowment of the good in each state s , $s = 1, 2$ with $x = (x_1, x_2) > 0$. These define the Edgeworth box $B = [0, x_1] \times [0, x_2]$, (see Figure 11.1). The endowment of agent 2 is represented by $e = (e_1, e_2)$, a point in the interior of the Edgeworth box. The endowment of agent 1 is $x - e$. Similarly, given a point $c = (c_1, c_2)$ in the Edgeworth box, the consumption bundle of agent 2 is c and that of agent 1 is $x - c$.

Each agent i , $i = 1, 2$, has a von Neumann–Morgenstern utility function u_i defined for all non-negative consumption w where u_i is twice continuously differentiable, $u_i'(w) > 0$, $u_i''(w) < 0$ for all $w > 0$ and $u_i'(0) = \infty$. Each agent has as objective the maximization of his expected utility.

Thus if agent i knew the probability of state 1 to be θ , he would evaluate the bundle (w_1, w_2) consisting of w_j units of the good if state j occurs ($j = 1, 2$) by taking the expectation:

$$\theta u_i(w_1) + (1 - \theta)u_i(w_2)$$

Any point c in the Edgeworth box, B , can be evaluated by each agent i in this way given the probability of state 1, θ . Thus, for any point c in B , and any $0 \leq \theta \leq 1$, let:

$$U_1(c, \theta) = \theta u_1(x_1 - c_1) + (1 - \theta)u_1(x_2 - c_2)$$

$$U_2(c, \theta) = \theta u_2(c_1) + (1 - \theta)u_2(c_2)$$

The functions U_1, U_2 then define the preferences of the two agents over points in the Edgeworth box, B , given some common value θ for the probability of state 1.

The curve labelled C in Figure 11.1 is the set of allocations that would be optimal if both agents were informed as to the true value of the parameter θ . Under our assumptions, this curve is independent of the actual




FIGURE 11.1 Basic concepts

value of θ so long as both agents agree on its value. Thus C is called the consensus contract curve. Formally:

$$C = \{ c \in B \mid u_2'(c_1)/u_2'(c_2) = u_1'(x_1 - c_1)/u_1'(x_2 - c_2) \}$$

It should be emphasized however that preferences do depend on the parameter θ . Geometrically, for each value of θ there corresponds a family of indifference curves for each agent. We assume that both agents believe correctly that this parameter can have only one of two values, θ_1 or θ_2 , with $0 < \theta_1 < \theta_2 < 1$. There are then two possible families of indifference curves for each agent, with the 'steeper' curves being associated with $\theta = \theta_2$ (see Figure 11.2). Thus for agent 1, for example, the consumption bundle c in Figure 11.2 is preferred to c' if $\theta = \theta_2$ and conversely if $\theta = \theta_1$.¹¹

We assume now that the actual probability of state one, θ , is drawn from a known distribution defined on $\{\theta_1, \theta_2\}$. We denote the probability that $\theta = \theta_1$ by p and the probability that $\theta = \theta_2$ by $1 - p$. Agent 2, the uninformed agent, is assumed not to know θ initially but to have the prior distribution given by p . Agent 1, the informed agent, observes the actual realization of θ . Again, both agents are assumed to know all other aspects of the environment.




FIGURE 11.2 Self-selection

In order to motivate the self-selection result, that certain allocations are not achievable in this environment because of the informational asymmetry, consider the following mechanism. Agent 1 (the informed agent) is asked to name a value of θ , either θ_1 or θ_2 . If he names θ_1 , then some allocation c' on the contract curve C is effected, and if he names θ_2 , then some allocation c'' also on C is effected. This mechanism has some *a priori* appeal since both c' and c'' are full-information Pareto optimal, that is, both are on C . But suppose that $c' \neq c''$, for example as shown in Figure 11.1. In this case, since it is impossible for agent 2 to require agent 1 to name the true value of θ , agent 1 will always claim $\theta = \theta_2$, even if $\theta = \theta_1$ (since $x - c'' > x - c$). Thus the allocation $c = c'$ if $\theta = \theta_1$ and $c = c''$ if $\theta = \theta_2$ is not achievable by this mechanism. (The equilibrium allocation is $c = c''$ for either value of θ .) As shown in the previous section, no matter how complicated we make the mechanism, allocations like:

$$c = \begin{cases} c^1 & \text{if } \theta = \theta_1 \\ c^2 & \text{if } \theta = \theta_2 \end{cases}$$

are *not* achievable in this environment *unless* agent 1 prefers c^1 to c^2 if $\theta = \theta_1$ and vice versa if $\theta = \theta_2$. This condition, which is illustrated in Figure 11.2, is the self-selection condition for this example. Clearly, the allocation $c = c'$ if $\theta = \theta_1$ and $c = c''$ if $\theta = \theta_2$ does not satisfy the self-selection condition and is therefore not achievable.

The remainder of this section is devoted to characterizing an optimal allocation mechanism and its equilibrium allocation for the example.

In searching for an optimal mechanism, we will use the Revelation Principle and consider only direct mechanisms. Thus we will simply search over allocation rules $G(\theta)$. Moreover, we will consider only such rules that result in truth-telling behaviour by agent 1, namely those that satisfy the (SS) constraints of the second section. This procedure will, of course, *not* result in ignoring mechanisms that perform better than the ones in the class we consider. Finally, for purposes of exposition, we will concentrate on finding a mechanism that maximizes the utility of the *uninformed* agent, agent 2, subject to the constraint that the informed agent, agent 1, be willing to participate (i.e. be no worse off than in autarky).¹²

These considerations lead us to characterize an optimal mechanism as an allocation rule $F(\theta) = [c_1(\theta), c_2(\theta)]$ which solves the following maximization problem:

$$\max pU_2[F(\theta_1), \theta_1] + (1-p)U_2[F(\theta_2), \theta_2] \quad (11.2)$$

subject to:

$$F(\theta_j) \text{ in } B \quad j = 1, 2 \quad (11.3)$$

$$U_1[F(\theta_j), \theta_j] \geq U_1(e, \theta_j) \quad j = 1, 2 \quad (11.4)$$

$$U_1[F(\theta_1), \theta_1] \geq U_1[F(\theta_2), \theta_1] \quad (11.5)$$

$$U_1[F(\theta_2), \theta_2] \geq U_1[F(\theta_1), \theta_2] \quad (11.6)$$

The objective function of this problem is simply the expected utility of the uninformed agent for the allocation F , assuming that the informed agent reveals θ truthfully (i.e. reveals θ_1 with probability p). This assumption is justified by imposing the self-selection constraints, (11.5) and (11.6). Constraint (11.5) guarantees that agent 1 prefers to report $\theta = \theta_1$ when in fact θ is θ_1 . Similarly, (11.6) guarantees that agent 1 prefers to report $\theta = \theta_2$ when in fact θ is θ_2 . Constraint (11.3) simply imposes techno-

logical feasibility, and constraints (11.4) guarantee that agent 1 is no worse off than in autarky for either value of θ (we are assuming that agent 1 can choose whether or not to play after having observed θ). Constraints (11.4) are often called 'individual rationality' constraints.

A solution of the problem (11.2)–(11.6) for a particular value of the endowment e is shown in Figure 11.3 (the proof may be found in Harris and Townsend, 1978). The exact location of $F^*(\theta_1)$ along the consensus contract curve C depends on agent 2's prior beliefs about the probability that $\theta = \theta_1$, namely p . As agent 2 becomes more certain that $\theta = \theta_1$, $F^*(\theta_1)$ moves along C towards point A_1 in Figure 11.3. Point A_1 is, of course, the allocation that gives agent 2 the most utility when $\theta = \theta_1$ subject to the individual rationality constraint (11.4) for $j=1$. When $p=1$, $F^*(\theta_1) = A_1$. Note that $F^*(\theta_2)$ is at the intersection of the θ_1 -indifference curve of agent 1 through $F^*(\theta_1)$ (dashed line in Figure 11.3) and the θ_2 -indifference curve of agent 1 through the endowment, e . Thus as p approaches 1 (agent 2 is certain that $\theta = \theta_1$), $F^*(\theta_2)$ approaches e (no trade if θ actually turns out to be θ_2). As p approaches 0 (agent 2 is certain that $\theta = \theta_2$), $F^*(\theta_1)$ both converge to point A_2 in Figure 11.3. Point A_2 is the bundle that maximizes agent 2's utility if he




FIGURE 11.3 Optimal allocation

knew that $\theta = \theta_2$, subject to individual rationality for agent 1. Thus we see that optimal allocations when information is asymmetric depend on agents' prior beliefs about the values of parameters they cannot observe.

Finally note that the optimal direct mechanism involves agent 1 declaring a value for θ , either θ_1 or θ_2 . The allocation is then $F^*(\theta_j)$, as in Figure 11.3, if he declares $\theta = \theta_j$, for $j = 1, 2$. Other mechanisms could, however, also be used to accomplish the same allocation. One such scheme is for agent 2 to start by offering agent 1 any menu of bundles c^1 and c^2 , then letting agent 1 choose which one is to be imposed. Obviously, agent 2 would offer $c^j = F^*(\theta_j)$ for $j = 1, 2$, so that this is just the direct mechanism in a thin disguise. One advantage of this version, however, is that the form of the optimal mechanism (signal sets and allocation rule) does *not* depend on the prior beliefs of agent 2. This is not the case for the optimal direct mechanism whose allocation rule F^* , depends on p . Perhaps this is why the alternative scheme has some intuitive appeal.¹³

CONCLUSIONS

In this chapter, we have argued that, to analyse the allocation of resources in certain types of environments with asymmetric information, one must first consider the process by which allocations are achieved. We have presented a methodology for such analyses, and applied this approach to a specific, abstract environment characterized by asymmetry of information between two agents.

In the remainder of these conclusions, we take up several issues not considered previously. In the process, we suggest some directions for further work. First, in our approach, as well as in the signalling literature (see, for example, Spence, 1974; Riley, 1979; Rothschild and Stiglitz, 1976; Wilson, 1977) the prior of the uninformed agent plays a key role in determining the allocation of resources. This is not the case, however, in some other attempts at devising mechanisms to allocate resources under asymmetric information, most notably in the public goods literature (see, for example, Groves and Ledyard, 1977). We explore this difference further in Harris and Townsend (1981).

Second, note that in the example of the previous section, at the end of the optimal direct mechanism, the actual allocation will not be full-information optimal (i.e. on C) if in fact $\theta = \theta_2$. At the time at which this allocation is effected, however, *both* agents are fully informed that $\theta = \theta_2$. Thus, *at that time* there will be gains to further trade (i.e. given *ex post* information). Certainly it is optimal to agree *ex ante* that such *ex post*

gains to trade will not be exploited, but in some circumstances such an agreement will be impossible to enforce (since both agents would like to violate it *ex post*). When modelling environments without suitable enforcement possibilities, it may be necessary to impose a constraint that the final allocation be Pareto optimal with respect to the final information structure.

Third, we have not considered the question of from where does an optimal mechanism come; how is it that agents adopt a particular allocation process. In some situations, it may be appropriate to assume that one of the agents has enough 'bargaining power' to impose the mechanism of his choice. If this agent has no private information, one can simply proceed as in the example of the previous section (see also Harris and Raviv, 1981a and 1981b). If the mechanism choice is made by an agent with private information, the choice itself may reveal some of this information. This is a much more difficult problem which has been taken up by Myerson (1982) and Holmstrom and Myerson (1981).

Finally, the approach outlined in this chapter is applicable only when all private information is known at the outset. One would conjecture that some form of the Revelation Principle will apply generally to situations in which private information is revealed over time.¹⁴

NOTES

1. By asymmetric information, we mean that some agents are *better* informed about some aspects of the environment and this fact is known to other agents.
2. See Harris and Raviv, 1981a, 1981b, and Myerson, 1981, for a start in this direction.
3. This assertion is motivated in the second section; see also Harris and Townsend, 1978, 1981, and Myerson, 1979.
4. The class is broad enough to include most imaginable mechanisms. One of course could restrict attention to a few obvious mechanisms, but that restriction would be counter to the spirit of our chapter – to make as few exogenous restrictions as possible. Our results are 'strong' to the extent that the class we consider is 'large'.
5. We believe the name 'Revelation Principle' was coined by Roger Myerson.
6. See also Harris and Townsend, 1981, and Harris and Raviv, 1981a, 1981b, for additional examples. The literature on contracts, information and incentives is, by now, replete with both implicit and explicit applications of the Revelation Principle. As this is not intended to be a survey chapter, we have made no attempt to provide references to this literature.

7. It is possible to allow one individual to observe several parameters and/or several individuals to observe the same parameter without affecting the results. See Harris and Townsend, 1981.
8. A 'signal' could actually be a whole sequence of functions specifying at each stage what message to send as a function of the history of the game to that stage. In this way sequential games can be modelled as normal-form games. Generally, in using the normal form, one does not spell out the sequential aspects explicitly. This is done in the 'extensive form' version of the game. See Friedman, 1977, for further discussion of normal and extensive form games. The results presented below can be proved also for extensive-form mechanisms; see Harris and Townsend, 1981.
9. Strictly speaking, we should allow for random allocation rules. For expository purposes, we shall ignore both random strategies and random allocation rules. The results of this section have been proved in this case by Myerson, 1979. In general, optimality may require random allocation rules (but not random strategies), as is argued in Prescott and Townsend, 1982, for example. Random allocation rules are not needed, however, for the example of the the third section, as shown in Harris and Townsend, 1978.
10. This example has been taken from Harris and Townsend, 1978. Several interesting mechanisms for this example are also discussed there.
11. Notice that we are changing slightly our notation here relative to the previous section. Since there is only one parameter in this example, we simply call it θ . Instead of using the subscript to refer to the agent whose type is given by the parameter, agent 1 in this case, we use the subscript to denote a particular value of the parameter θ . We hope that no confusion results.
12. The arguments in Harris and Townsend, 1978, establish that one can generate the entire class of optimal allocations in this way, if one varies the endowment parametrically. Also the particular optimality concept employed here is not critical to the argument. The Revelation Principle can be applied to environments with diverse optimality criteria.
13. The alternative mechanism is sequential, however, thereby falling outside the class of mechanisms considered here. Again, the reader is referred to Harris and Townsend, 1978, for a more general treatment.
14. See Townsend, 1982, for a proof of this conjecture in a particular model.

REFERENCES

- Friedman, J. (1977) *Oligopoly and the Theory of Games*, Amsterdam: North-Holland.
- Groves, T., and J. Ledyard (1977) 'Optimal Allocation of Public Goods: A Solution to the "Free Rider Problem"', *Econometrica*, 45 (4) (May): 783–810.

- Harris, M., and A. Raviv (1981a) 'Allocation Mechanisms and the Design of Auctions', *Econometrica*, 49 (6) (November): 1477-99.
- Harris, M., and A. Raviv (1981b) 'A Theory of Monopoly Pricing Schemes with Demand Uncertainty', *American Economic Review*, 71 (3) (June): 347-65.
- Harris, M., and R. M. Townsend (1978) 'Allocation Mechanisms for Asymmetrically Informed Agents', working paper no. 35-76-77, GSIA, Carnegie-Mellon University (March 1977, revised September 1978).
- Harris, M., and R. M. Townsend (1981) 'Resource Allocation Under Asymmetric Information', *Econometrica*, 49 (1) (January): 33-64.
- Harsanyi, J. C. (1967-8) 'Games with Incomplete Information Played by "Bayesian" Players (Parts I-III)', *Management Science*, 14: 159-82, 320-34, 481-502.
- Holmstrom, B. (1978) 'On Incentives and Control in Organizations', unpublished PhD dissertation, Graduate School of Business, Stanford University.
- Holmstrom, B., and R. Myerson (1981) 'Efficient and Durable Decision Rules with Incomplete Information', Discussion paper no. 495, CMSEMS, Northwestern University (September).
- Hurwicz, L. (1972) 'On Informationally Decentralized Systems', in C. B. McGuire and R. Radner (eds), *Decision and Organization* (Amsterdam: North-Holland, ch. 14).
- Myerson, R. (1979) 'Incentive Compatibility and the Beginning Problem', *Econometrica*, 47 (January): 61-74.
- Myerson, R. (1981) 'Optimal Auction Design', *Mathematics of Operations Research*, 6 (February): 58-73.
- Myerson, R. (1982) 'Mechanism Design by an Informed Principle', Discussion paper no. 481, CMSEMS, Northwestern University.
- Prescott, E., and R. M. Townsend (1982) 'Pareto Optima and Competitive Equilibria With Adverse Selection and Moral Hazard', working paper, GSIA, Carnegie-Mellon University, forthcoming in *Econometrica*.
- Riley, J. (1979) 'Informational Equilibrium', *Econometrica*, 47 (2) (March): 331-60.
- Rothschild, M., and J. Stiglitz (1976) 'Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information', *Quarterly Journal of Economics*, 90 (November): 620-49.
- Spence, A. M. (1974) 'Competitive and Optimal Responses to Signals: An Analysis of Efficiency and Distribution', *Journal of Economic Theory*, 7 (March): 296-332.
- Townsend, R. M. (1982) 'Optimal Multiperiod Contracts and the Gain from Enduring Relationships Under Private Information', *Journal of Political Economy*, 90 (6) (December): 1166-86.
- Wilson, C., (1977) 'A Model of Insurance Markets with Incomplete Information', *Journal of Economic Theory*, 16: 176-207.

12 Production Functions, Transactions Costs and the New Institutionalism

VICTOR P. GOLDBERG*

INTRODUCTION

For most of the postwar period, economic theory has focused on the analysis of impersonal markets. In the past few years, however, there has been a resurgence of interest in the role of institutions in the allocation process: why does some behaviour take place within firms and not within markets? Why are long-term contracts used instead of spot markets? What determines the structure of long-term contracts? How does the internal organization of a firm affect its performance? Why are some workers compensated by piece-rates, others by hourly wages, and still others by annual salary? What are the effects of seniority provisions or of a legal prohibition of termination of employment contracts at will? Does the structure of employment contracts have an influence on macroeconomic variables? What are the effects of alternative tort liability systems on accident rates? And so forth.

In this chapter I want to consider two concepts - production functions and transactions costs - that have been used and abused in developing the New Institutionalism. If we are to rely on them at all in our exploration of the causes and effects of economic institutions, it will be necessary to subject them to careful scrutiny. My reading is that the transactions cost concept in particular has proved to be misleading and unhelpful and that it would be best if we simply abandoned it. However, the terminology is probably too deeply entrenched for this cold turkey approach

* This author would like to thank the following for comments on an earlier draft: Moshe Adler, Avner Ben-Ner, Ronald Coase, Douglass North, Joseph Ostroy, David Teece, and Oliver Williamson.