COWLES FOUNDATION DISCUSSION PAPER NO. 59b

Economic Theory of Teams*

Chapter 2

J. Marschak and R. Radner

October 29, 1958

* Research undertaken by the Cowles Commission for Research in Economics under Contract Nonr-358(01), NR 047-006 with the Office of Naval Research

Economic Theory of Teams

J. Marschak and R. Radner

CHAPTER 2

ORGANIZATIONAL FORM: INFORMATION AND DECISION FUNCTIONS

1. Rule of action. So far we have discussed the consistency of decisions and emerged with the concept of the expected payoff as our primary tool for the evaluation of actions under uncertainty. Given the (subjective) probability distribution of the states of the environment, the best action is the one with the highest expected payoff. Choice under certainty is a special case, with probability one assigned to one particular state of environment.

We now modify -- and, in a sense, generalize -- the problem in the following way: the individual chooses, not among actions but among rules of action. A rule of action (also called a decision rule, a strategy, or a decision function) is a schedule that determines in advance, for each possible future information, the action that will be taken in response to it. Rules of action (sometimes called "rôles") for the individual members of an organization are essential for the very concept of organization as defined by us at the outset. The search for the best rules of action will be seen to be essential to the economic theory of teams. It is also essential for a realistic theory of single-person decisions since one often has to decide in advance how to respond to each of possible future contingencies. This will be shown in various examples in the present and the next chapter. It will be seen that -- as emphasized by A. Wald [ ] and by von Neumann and Morgenstern [ ] -- this introduction of the possibility of using new information does not alter the basic features of the decision problem. Formally, the rules of action will now play the same role as did the actions themselves in the previous chapter.

The "future" information that enters in the decision rule is to be distinguished from "prior" information. The prior information consists in the description of the set X of possible states x of the environment, the probability distribution φ on X, the set A of alternative actions a, and the payoff function ω. In what follows, information will always mean future information. This will be defined more precisely in Section 2, but we shall first use a familiar example to illustrate the meaning of our concepts.

Consider a firm producing a single commodity for a market in which the price is set by the government at the beginning of each year, at which time the firm sets its production for the coming year. If the goal of the firm is to maximize profits for the coming year then under the usual assumptions of decreasing marginal costs, etc., we get the familiar rule: "Choose the level of production that will make marginal cost equal to price." This rule defines, implicitly, a functional relation between the action -- viz., the quantity a produced -- and the information -- viz., the price y set by the government, say,

$$a = \alpha(y) \ .$$

Thus α is the decision function prescribed by economic theory.

One can, of course, imagine other decision functions for this firm, say the one implied by the rule: "Choose the level of production such that average cost is ninety percent of price." (This decision function typically will not maximize profits, however.) Another decision function, although a somewhat "degenerate" one, is the constant function, namely, the rule that fixes the same level of production year after year, whatever the price happens to be.

In general, let Y be the set of possible informations; then a decision function α is a function from Y to the set A of alternative actions. (In the example of the firm just given, Y is the set of possible prices, and A is the set of alternative levels of production.) An action a depends on the information y thus: $a - \alpha(y)$. The set of alternative decision functions will be denoted by $\{\alpha\}$.

We can now rewrite the payoff $\omega(x,\alpha)$ as $\omega[x,\alpha(y)]$ where x is the generic element of X and y is the generic element of Y. If y_1 and y_2 are elements of Y we shall also say that they are two values of the information variable y. We shall show presently the relation between X and Y.

2. Information. We have denoted by $\omega[x, \alpha(y)]$ the payoff when the true state of the environment is x, and the action is based on information y and on the rule α. Clearly not all imaginable aspects of environment are relevant to the payoff of a particular action. The profit that I may realize on the purchase of domestic natural gas stock may depend on the alternation of seasons but not perceptibly on the rest of the astronomical data, nor on the dynastic changes among Uganda chieftains. Hence in the formulation of the decision problem the detail with which the set X of possible states of environment is described, will depend, at least in part, on the

nature of the payoff function. We shall presently see that it will also depend on the available types of information.*

---

* To avoid certain technical mathematical complications, we will assume that the set X is finite, except in some special applications, in which cases the reader will be warned.

Typically, information will give only a partial description of the state of the world; this description can have varying degrees of completeness. For example, instead of getting a complete list of today's closing prices on the New York Stock Exchange, one may get information only about all the gas stock prices; or only about the average of all stock prices; or only about the average gas stock price. Each of these types of information specifies some set of states of the environment within which the true state lies. This set is, of course, the set of all states that have in common the partial description given by the information. Thus, to be told that the average closing stock price was 55, is to be told that the list of prices is among the set of all those lists that have an average of 55.

Thus each information y (an element of Y), is identified with a particular subset of X . The information relevant to the motorist's decision in the traffic may be a signal, green or red. Thus the set Y consists of two elements. The set X of all possible traffic situations is thus partitioned in two subsets; if traffic is in one subset the signal is green, otherwise it is red. This partitioning defines η , the <u>information structure</u>; given η , certain signals (symbols) are assigned to certain subsets of X .

Note that a decision has sometimes to be based on information that reflects aspects of environment which, in fact, do not influence the payoff. For example, in choosing the parts of a country in which a campaign against infantile diseases is most urgent, one may have to base the decision, in the absence of better data, on mortality figures not broken down by age groups. The decision is then made on what is sometimes called "erroneous" information. There is really no need for a special concept of incorrect information. Every item of information is correct with respect to some aspect of the state of the environment although that aspect may be irrelevant to the payoff. Thus, instead of saying that the results of a market survey are "incorrect," we shall say that those results reflect not only the responses of the people interviewed, but also the characteristics of the interviewer, the method of recording the data, etc.; and that, if one ignores this nature of the particular information he can make poor decisions. Sometimes it is convenient to say that the environment variable x consists of a pair of variables: $x = (x_1, x_2)$ where x_1 = true state, and x_2 = the error of observation. See Example F in Chapter 3.

In summary, then, an item of information represents a subset of the states of the environment; and, in the formulation of a decision problem, the states of the environment must be described in sufficient detail to cover, not only those aspects relevant to the payoff function, but also those aspects relevant to the type of information on which the decisions may be based.

3. Information structure. Organizational form. Information can be regarded as the outcome of "information-gathering." A given method of information-gathering applied to the true state of environment x, results in a particular information. For example, consider again our stock price example, where $\underline{x}$, the true state of environment (in its aspect relevant to the payoff) is the complete list of prices; whereas $\underline{y}$ is the average price. To each $\underline{x}$ corresponds a $\underline{y}$; this relation we have already called information structure (it may also be called information function), and denoted by η. Thus $y = \eta(x)$. Each method of information gathering -- the getting as well as processing of data -- is characterized by a particular information structure. Thus, if x still denotes the complete list of prices, then $y' = \eta'(x)$ may be the average price of gas stocks, $y'' = \eta''(x)$ may be the price of a particular gas stock; $y''' = \eta'''(x)$ may give a range of prices that includes this price, and $y'''' = \eta''''(x)$ may be the quotation of this price given by a notoriously inaccurate local broker. We have thus enumerated five different information structures. (A good term might also be "information filters.")

Formally, if Y denotes the set of possible informations, i.e., outcomes of information-gathering, then an information structure is a function η from X, the set of possible states of the environment, to Y.

Since a particular information (signal) y is identified with some subset of the set X, it follows that every information structure is some partitioning of X into an exhaustive family of mutually exclusive subsets, each subset corresponding to a particular information.

This picture of information structure is sometimes helpful when one attempts to visualize the comparison between different information structures. We shall see that making a partition "finer" permits a better choice of decisions and thus may make the corresponding information structure more valuable; whereas making the partition "coarser" results in a less valuable structure. We will return to a more detailed discussion of the comparison of information structures in Sections 7 and 8 of this Chapter.

For the purposes of this book it will be convenient to give to the pair (η,α) -- the combination of an information structure with a decision rule possible under this structure -- a special name: the "organizational form."

4. Expected payoff reformulated. Since, given the information structure η and the true state of environment x, the information y is determined, $y = \eta(x)$, we can rewrite the payoff again, as follows:

$$\omega(x,a) = \omega[x,\alpha(y)] = \omega[x,\alpha(\eta(x))].$$

Hence, given the true state of environment, the payoff is determined by the information structure, the decision function, and the payoff function. Using now the probability distribution φ, we can write the expected payoff thus:

$$(1) \qquad \sum_x \omega(x,\alpha(\eta(x))\ \varphi\ (x) = \Omega\ (\alpha,\eta;\ \omega,\varphi)$$

say: This quantity depends on the non-controlled conditions ω, φ;

and on the decision function α and the information structure η. In general, these latter two functions are under the control of the decision maker: He has at his disposal more than one pair (α,η) and he will choose that pair which makes the expected payoff U a maximum. This justifies our previous assertion that the problem of choosing the best decision function (and we may now add the best information structure) is formally the same as the simpler one of choosing the best action.

We shall be often able to simplify the discussion by assuming η as given, leaving thus only α to the individual's choice.

We shall again consider the example of a firm (see Section 1 above) but shall now introduce an additional factor into the description of the states of the environment, namely, the price of an important raw material. We shall denote by x_1 the price of product, and by x_2 the price of raw material. Suppose that, at the time the decision about the level of production is to be made, it is not known what price will have to be paid for the raw material during the coming year. Thus, the information variable y is not identical with the state of environment, x. The latter is described by the pair

$$x = (x_1,x_2) \ ;$$

the information y is described by price x_1 alone; the information structure is given by

$$\eta(x) = x_1 \ .$$

Let us assume this to be the only information structure available. Suppose that the cost of producing a quantity a, for given x_2, is

$\kappa(x_2,a)$; and that the firm takes as its measure of utility the net profit; then the payoff function is

$$(2) \qquad \omega(x_1,x_2,a) = x_1 a - \kappa(x_2,a).$$

If the firm decides upon a rule that tells it to produce the quantity $a = \alpha(x_1)$ when it learns that the price of the product is to be x_1, (regardless of x_2) then the payoff for any state (x_1,x_2) is

$$u = \omega(x_1,x_2,a) = x_1\,\alpha(x_1) - \kappa(x_2,\alpha(x_1)) .$$

The prices x_1 and x_2 will have some joint probability distribution $\varphi(x_1,x_2)$ and the decision function α will be evaluated by the expected value of the payoff u just given.

The reader is invited to go over the list of concepts and the Figure 1 at the end of this chapter, and to consider the central and the right-hand part of Figure 1, (neglecting for the moment its left-hand side, dealing with costs of information). Circles are sets; boxes are functions from one set to another. Thus the box "information structure", η, is a function from the set of states of nature, X, to the set of informations, Y. Box "decision rule", α, is a function from Y to the set of actions, A; ω is a function from A and X to the set of payoffs* (set of real numbers). The circle "probability"

---

* They are called here "gross payoffs," as the organizational costs, dealt with on the left-side of the diagram, are still to be deducted.

---

is the set of all non-negative numbers not exceeding 1, and the box "probability function" assigns some such number to each state of nature.

To the sets "probability" and "gross payoff" the operation of "weighted averaging" is applied (not represented by a box), to yield the expected payoff. Note also that the box "decision rules" is itself enclosed in a (dotted) circle; this conveys the fact that a given information rule α is itself an element of a set α of decision rules among which a choice has to be made; the circle is "dotted" to convey that unlike "Nature","Action" etc., the decision rule is a thing to be chosen. Another such thing is the "information structure": the particular box representing a particular information structure η is an element of a set of available information structures, from which a choice is being made.

5. Maximizing Conditional Expectations. A decision function is best if it results in the largest possible expected payoff, i.e., the largest possible value of $E\omega(x,a)$. We shall now show that the nature of $E\omega(x,a)$ enables one to give a more detailed characterization of a best decision function.

First, consider the situation in which the decision-maker finds himself, after he has received the information $\eta(x) = y$. He is about to take an action, $\alpha(y)$, and the consequence of this action is (typically) uncertain, since he knows only that the true state of the environment is one of the (typically) many that could have resulted in the particular information y . In other words, the consequence of the action $\alpha(y)$ is a prospect, in the sense defined in Section 1.8; we shall denote this prospect by π_y .

What are the probabilities associated with the prospect π_y ? They are clearly the conditional probabilities of the states x , given that $\eta(x) = y$. These conditional probabilities $\Pr\left\{x|\eta(x) = y\right\}$, where the vertical bar is, as usual, to be read "given that," are sometimes called "posterior probabilities," i.e., the probabilities of the states x after the information y has been received. The expected utility of the prospect π_y is the <u>conditional expectation</u>

$$(5.1) \qquad E\left\{\omega(x,\alpha(y)) \mid \eta(x) = y\right\} = \sum_x \omega(x,\alpha(y)) \Pr\left\{x|\eta(x) = y\right\} ,$$

The prospect π_y , however, only arises with a certain probability, namely, the probability that the information received is y . Thus the consequence of using the decision function α is a <u>probability mixture</u> of all the various prospects π_y , corresponding to varying y (see Section 1.9). The utility of this over-all prospect is therefore a weighted average of the utilities of the prospects π_y (the weight being, for each value of y , the probability $\Pr[\eta(x) = y]$ that this particular value of y is observed). Hence, to make this average utility as large as possible, the decision-maker must make the utility for each prospect π_y as large as possible, i.e., choose an action that maximizes the expression (5.1).

Thus we have arrived at the following theorem: <u>For α to be a best decision function, it is necessary and sufficient that, for every y, $\alpha(y)$ be an action that maximizes the conditional expected payoff given $\eta(x) = y$</u> .

To illustrate, consider the last example of the production decision problem (Section 4). Here $x = (x_1, x_2)$, $y = \eta(x) = x_1$. For any level of production a , the expected net profit, given output price x_1 , is

$$(5.2) \qquad E\left\{\omega(x,a) \mid x_1\right\} = x_1 a - E\left\{K(x_2,a)|x_1\right\}$$

(see equation (4.2)). Setting the derivative, with respect to a, of this last expression equal to zero, we find that the best value of a, given x, must satisfy*

---

* It is of course assumed that the various conditions necessary for this "marginal analysis" to be valid are satisfied.

$$E\left\{\frac{\partial}{\partial a} K(x_2,a)\,|\,x_1\right\} = x_1 \tag{5.3}$$

In other words, "conditional expected marginal cost must equal price."

6. Cost of Decision. In our definition of a best decision function we have thus far ignored one important factor, the costs of using a decision function. The mere calculation of the action prescribed by a complicated decision function for given information may be a costly procedure. Beyond this, some decision functions may involve greater "administrative" expense than others. Some of the decision costs are fixed once the decision function α is chosen. Others depend on the state of the world, x (in particular they might depend on the information y), and are therefore themselves variables. Thus, in general, we may express decision cost as some function of α and x, $\delta(x,\alpha)$, say. Because of decision costs (and of the costs of information to be discussed in Section 9) it is appropriate to call the expected payoff $E\omega(x, \alpha(y))$ the gross expected payoff. If payoffs as well as decision costs are measured in money, the best decision function would be one that maximizes, not the gross expected payoff but its excess over the decision cost; i.e., it maximizes the quantity $E\omega(x, \alpha(y)) - E\,\delta(x,\alpha)$, where $y = \eta(x)$ and the expectation is taken with respect to x.

Unfortunately, there has been very little theoretical analysis of costs of decision that we are aware of, and, aside from this brief acknowledgement of the importance of these costs, they shall play no role in our theory.

7. The Value of an Information Structure. Once a best decision function has been chosen, given the information structure, nothing more can be achieved in the way of increasing the expected payoff without changing the information structure itself. A different information structure would, of course, typically require a different best decision function, and might possibly result also in a higher expected payoff. Thus one is led to a natural definition of the difference in value between two information structures as the difference in the maximum expected payoffs that can be achieved through their use.

When comparing information structures, there is an abvious "zero." Consider the class of all constant information structures, i.e., the class of all functions η such that information $\eta(x)$ is a constant, independent of the value of x. The partition of the set X of states of the environment that is associated with a constant information structure has only one element, the set X itself. A constant information structure implies, of course, a constant decision function (the most extreme form of routine!). It is clear that a constant information function gives no information at all that has not already been incorporated into the formulation of the decision problem. Therefore it is appropriate to assign to it the value zero. This suggests that one define the absolute value of an information structure as the difference in value between it and a constant information structure. This quantity can never be negative; the set of decision functions available to a person using a non-constant information structure includes all of the decision functions that are available to a person using constant information structure (viz., all constant decision functions), and possibly more. Hence the first person cannot do worse than the second. There is no damage in knowledge!

If we denote the expected payoff for an information structure η and a decision function α by $\Omega(\alpha,\eta)$, -- see equation (4.1) -- then the above definition can be expressed symbolicially thus:

$$(\text{Value of } \eta) = \max_{\alpha} \ \Omega(\alpha,\eta) - \max_{a} \ E\,\omega(x,a).$$

The maximum expected payoff yielded by a given information structure will be denoted by

$$\hat{\Omega}(\eta) = \max_{\alpha} \Omega(\alpha,\eta).$$

8. Fineness of Information Structures. Is it possible for one information structure to be more valuable than another, or at least not less valuable, whatever be the basic payoff function ω ? The answer is yes, and the characterization of such a relationship is provided by the concept of fineness. We shall say that the information structure η' is finer than η'' if the partition of X corresponding to η' is a subpartition of that corresponding to η'', i.e., if every set in the first partition is contained in some set of the second. (Thus η' tells us all η'' can tell, and possibly more besides.)*

* In another terminology, η' is called an extension of η'', and η'' a contraction of η' .

For example, if X is the set of all numbers between 0 and 1, and η' partitions X into ten equal intervals, while η'' partitions X into a hundred equal intervals, η'' is finer than η' (η'' has one digit more!). If X is a set of all pairs (x_1, x_2) of integers and under η' each pair constitutes a subset, while under η'' all pairs with the same value of x_1 constitute a subset, η' is finer than η''. If η' and η'' correspond to the same partition they are said to be equivalent. Clearly, not of every pair η', η'' can it be said that one of them is finer than the other or that they are equivalent: For example, let X be the set of all numbers between 0 and 1; let η' be the partition of X into two sets: "elements of X larger than 1/2", and the "other elements of X"; and let η'' partition X into "elements of X larger than 2/3", "elements of X smaller than 1/3" and "other elements of X." Then η'' is not finer than η' in the sense defined; nor is η' finer than η''; nor are they equivalent. (Thus the relation "finer than" induces only a partial, not a complete ordering of partitions of X).

The significance, for the present question, of the concept of fineness lies in the following theorem: Suppose the probability distribution on X assigns positive probability to every state x in X; then <u>the value of η' is at least as great as the value of η'' for every payoff function ω if, and only if, η' is finer than η''</u>.

The "if" part of this theorem is fairly obvious, for the set of decision functions available to a person using η' essentially includes all of the decision functions available to a person using η'', and possible more, so that the first person cannot do worse than the second. (A special case of this was presented in the preceding section, with η'' a constant function.)

The "only if" part of the theorem is perhaps not quite so obvious. Suppose η' is not finer than η'' nor is η'' finer than η', nor are they equivalent. Then there must exist three states of nature x_1, x_2, x_3 such that, under the partition corresponding to η', x_2 and x_3 but not x_1 are included in the same subset; while under the partition corresponding to η'', x_1 and x_3 but not x_2 are in the same subset:

$$\eta'(x_2) = \eta'(x_3) \neq \eta'(x_1) ,$$

$$\eta''(x_1) = \eta''(x_3) \neq \eta''(x_2) .$$

Now consider a decision problem in which it is sufficiently more important to distinguish between x_1 and x_3 than between x_2 and x_3 (it is easy to construct such problems: see Chapter 3, Example B. Then η' would be more valuable than η'' ; whereas the situation would be reversed in a problem in which it is very important to distinguish between x_2 and x_3 , but not between x_1 and x_3 .

One interesting corollary of this theorem is that it is impossible to define a single measure of the "amount" of information (without regard to the payoff function), such that if one information structure provides a greater amount of information than another, the first will be more valuable than the second, for every payoff function. A measure of the amount of information independent of the payoff function and depending only on the probabilities of the laternative signals y was proposed by C . Shannon [ ]. In the simple case when the signals are finite in number and equiprobable, Shannon's measure is an increasing function (the logarithm) of the number of signals. For example, if x is uniformly distributed over the set X which is an interval, divided into n equal sub-intervals, and if the information $y = \eta(x)$ consists in stating the sub-interval into which the number x falls, then Shannon's measure is the larger, the larger n. Yet this ranking of information structures according to the number of sub-intervals need not coincide with the ranking of the values of those information structures. In our Example B, Chapter 3, the payoff function is such that the value of information is highest when the number of sub-intervals is 2 or any even number; so that is is more valuable to use two equal sub-intervals than to use one hundred and one!

9. Cost of Information Structure. Information, like decision, typically costs something. The information structure that results in the highest expected payoff may involve costs of decision and of information that are so high as to make some other information structure preferable.

One part of information costs is fixed once the information structure is chosen. For example, one may choose to base his future decisions on the outcome of a sample of a fixed size; or on the information of a forecaster who charges a fixed fee regardless of the outcome of each forecast. Another part of information costs is a random variable whose value depends on the actual state of the world. For example, instead of fixing the sample size in advance, one may make it dependent on the outcome of observations whether to continue or -- if they have already been significant enough -- to stop them (sequential sampling). One may arrange with the forecaster to pay him larger sums for those forecasts that prove to be more successful. One may also arrange to pay him a fixed overhead sum plus a success bonus. Thus, in general, total information cost will be a random variable. The cost will depend on the actual state of the world (x) the information structure used (η); it can be expressed as $\gamma(x,\eta)$, say. The expected information cost, $E\ \gamma(x,\eta)$ (where the expectation is taken with respect to x) depends on the functions η and γ and on the probability distribution φ on X .

As in the case of decision cost function δ , little is known about the information cost function γ . It is important to note that, unlike the value of information, the cost of information typically does not depend in a direct way on the payoff function. Like Shannon's amount of information, the cost of information does depend on mathematical properties of the set Y of signals, viz., on the probability distribution over this set. This, in turn, depends on the information structure η and on the probability distribution φ over the set X of states of nature. However, two systems of signals with the same probability distributions may involve different costs. The cost in time that it takes a decision maker to get information on his own, or the money cost that is charged in the market of purchaseable information services, depend on additional factors. The fees for information services depend, for example, on the relative bargaining positions of sellers and buyers of such services, so that, ultimately, the values of a given information structure -- not its value for the particular decision maker alone but for other users of this kind of information as well -- do influence the cost of information. This subject matter has been opened up by Good [ ] and McCarthy [ ], but we shall not pursue it here.

10. Net expected payoff. If the payoffs as well as the costs of decision and information are measured in money, we can formulate the decision problem as that of finding the decision function and the information function that jointly maximize the net expected payoff. This is the difference between the gross expected payoff $E\omega(x,a)$ and the total of expected decision and information costs. The left side of our Figure 1 helps to illustrate this. The maximum net expected payoff is then

$$\max_{\alpha,\eta} \quad E\omega(x,\alpha(\eta(x))) - E\ \gamma(x,\eta) - E\ \delta(x,\alpha)$$
$$= \max_{\alpha,\eta} \quad \Omega^{*}(\alpha,\eta;\omega,\gamma,\delta,\varphi) \ ,$$

say. The right-hand side expression emphasizes the distinction between the givens of the problem (the functions ω, γ, δ, φ) and the controlled variables (the functions α,η). The optimal pair $\hat{\alpha},\hat{\eta}$ depends entirely on the given functions ω, γ, δ and φ. The subject of single-person decision theory can be said to be this: to determine how the properties of the optimal decisions and information functions depend on properties of the givens: the payoff function ω (which, for example, may be linear, non-linear, etc.); the probability distribution φ (which, for example, may be characterized by larger or smaller correlations between environment variables); and the decision and information cost functions, δ and γ.

From the point of view of general utility theory as discussed in Chapter 1, the separation, and subtraction, of "costs" from "gross payoff" is not permissible except in special cases. The decision maker assigns a utility to each outcome of his decisions and of the state of the world. If he has to sacrifice time or other things to achieve certain other things, the "outcome" is the combination of the sacrifices and the achievements. The utility assigned to this combination is not, in general, representable as a difference between some "utility of things achieved" and "disutility of things sacrificed." However, because of its simplicity, the assumption that such separation of achievements from sacrifices is possible has great methodological advantages, at least as an approximation. This approximation may be quite a close one in a society in which there is a market for a great many things including the time of the people who make decisions and convey information. This permits us to put a monetary value on gross payoff as well as on information and decision costs, and consequently on the net payoff.

The problem of decision making, especially in its applications to several-person organizations, is so full of subtle complications that it seems worthwhile to make the assumption of separable achievements and sacrifices in order to throw some light on the problem.

11. Summary of concepts. Our brief sketch of the single-person organization problem is complete, and this seems to be a good place for a list of the concepts, and the symbols denoting them, that will be used in the following chapters. The reader will notice that not all the concepts that have been introduced in this chapter are included in the list. Some of them, such as "prospect" and "outcome function," have already fulfilled their roles as introductory or intermediary ideas.

Symbol	Concept
A	The set of alternative actions, a, available to the decision-maker.
X	The set of (mutually exclusive) states, x, of the environment. The uncertainty about x is expressed by a probability distribution on X .
ω	The payoff function, a real-valued function on (A,X). $\omega(x,a)$ is the payoff, in utility, to the decision-maker when he takes action a, and x is the true state of the environment.
Y	The set of possible outcomes, y, of observation.
η	An information function (or structure): a function from X to Y. $\eta(x)$ is the outcome of observation when x is the true state of the environment.
α	A decision function (or decision rule): a function from Y to A . $\alpha(y)$ is the action prescribed by the decision function α when y is observed.
(α,η)	An organizational form
$\Omega(\alpha,\eta)$	The (gross) expected payoff resulting from the use of the information structure η and the decision function α . $$\Omega(\alpha,\eta) = E\, \omega(x,\alpha[\eta(x)])$$

Symbol	Concept
$\hat{\Omega}(\eta)$	The value of the information structure η . $\hat{\Omega}(\eta) = \max_{\alpha} \Omega(\alpha,\eta) - \max_{a} E\omega(x,a)$
$\gamma(x,\eta)$	The cost of the information structure η .
$\delta(x,\alpha)$	The cost of the decision rule α .

The decision maker should choose that information structure η and decision function α that maximizes the net expected payoff

$$\Omega(\alpha,\eta) - E\,\gamma(x,\eta) - E\,\delta(x,\alpha)$$