Introduction

Economics is concerned with efficient ("best," "optimal") use of limited resources. A decision is "economical" if, among available decisions, it is the most desirable one. During the last two or three decades, explicitly economic thinking has been extended beyond its traditional domain, that of the production and distribution of marketable goods. The same general principles that are used to determine efficient sites and production plans for factories and farms, to explain the prices agreed upon by presumably efficient buyers and sellers, or to outline plans for the material development of a country have been applied to broader fields of human decisions and of interactions between men.

The economic concern, the interest in optimal choices, has of course characterized practical arts such as engineering, military planning, and medicine from their very inception. In our generation, this concern has become articulate. Under names such as operations research, cost-efficiency analysis, systems analysis, mathematical programming—clumsy, casual names suggesting their origin in practice rather than in philosophy—complex decision problems are being stated explicitly as such. That is to say, both the constraints of feasibility (limited resources, set of available actions) and the criteria of desirability (ordering of preferences, goals, values) are stated, opening the way to a systematic solution of some maximization problem.

Parallel with these practical developments, economic thinking has also penetrated the very foundations of empirical science. The current logic of inductive inference developed by logicians and mathematicians—notably F. P. Ramsey (1931), B. de Finetti (1937), R. Carnap (1962), and L. J. Savage (1954)—has related a man's "beliefs" about the "probabilities" of uncertain events to his choices between actions that have uncertain outcomes and are viewed as "bets." Moreover, the decision-maker's beliefs will vary depending on his access to information, and his choice of action can be improved by changing his instrument of information. For example, suppose that the results of each possible experiment (or a sampling, a survey) will be used to choose the best possible action; which is then the best possible experiment? Such problems have occupied workers in statistics—beginning with J. Neyman and E. S. Pearson (1928 to 1938) and A. Wald (1950)—and in "adaptive control processes," so called by R. Bellman (1961). In essence, decision theory can be called the "economics of information."

The economic approach has been applied not only to efficient choices of a single person, but also to joint actions of several men. In the theory of games of von Neumann and Morgenstern (1943), further developed by other mathematicians and economists, two or more persons follow generally different desirability criteria ("interests," "preferences"), and are constrained by different feasibility conditions ("rules of the game") or at least attach to them different probabilities; but each player is efficient in a well-defined sense. The problem is to find which arrangements exist (if any) that would be supported by each player's self-interest as he sees it, because no change that he might be able to enforce would better him. Such arrangements, if attained, would be maintained. They are called viable. In this sense, viability can be regarded as a generalization of the optimality requirement, extended from decision-making by a single person to the case of several persons. Corresponding to the choice of the best information instrument, and of the best action based on the information it provides, the viable arrangement between players of a game is implemented by allocating to each certain activities, possibly including the tasks of gathering or communicating specified kinds of information. Viewed in this way, the problem of viable arrangements in games can be conveniently called a theory of efficient organizations. The problem of the optimal information instrument and the optimal decision for a single person can then be regarded as a special (though by no means trivial) case.

The economic theory of teams attacks a middle ground. We study the case in which several persons perform various tasks including those of gathering and communicating information and of making decisions; but they have common, not divergent, interests and beliefs. Hence the optimality requirement is easily defined, just as in the case of a single person. But the single person's problem of optimizing his information instrument and its use is replaced by that of optimizing the allocation of tasks among the members of the team.

We have equated *economical* and *efficient* to denote an arrangement that is most desirable (or, in the general case of organization, one that is viable), under given feasibility constraints. These constraints include the limitations of human capacities for communication and good decision-making (and analogous limitations of inanimate instruments). While our concern is a practical, purposive, prescriptive one, the general solutions we discuss would depend, when applied to any concrete case, on data supplied by workers in descriptive fields—the names of R. Cyert, J. March, and H. Simon come to mind![1]—psychological or sociological data along with those of natural technology. On the other hand, our results may prove of some value to descriptive theorists of human organization by pointing to those data, quantitative and otherwise, that would be of most importance if one wanted to increase the efficiency of a given organization.

Part One of this book, the Prologue, prepares the ground for the theory of teams by discussing the case in which the team consists of one man. Chapter 1 defines his tastes and beliefs, under the condition that he is consistent in the sense of obeying the rules of logic and certain plausible, quasi-logical postulates of decision behavior.

These conditions imply the "expected utility theorem," which is then used throughout the book. Chapter 2 outlines what we have called above the economics of information: the problem of choosing an optimal information instrument. Chapter 3 provides simple illustrations.

In the central Part Two of the book, the team problem is first simplified into the problem of optimal *information structure*: we ask "who should know what?", without inquiring how this allocation of information is achieved. We proceed from some simple illustrative examples in Chapter 4 to the problem, treated in Chapter 5, of finding optimal decision rules for the team members when the information structure is given. In Chapter 6, various important information structures are compared as to the maximum expected payoffs that they can yield. In most of the cases studied in the book, a "static" environment is assumed, but Chapter 7 outlines an approach to the case in which the environment is a process over time, and in which delays intervening between information and decision may, under certain conditions, result in losses.

A given information structure can be generated by a variety of "networks." Each element of a network is characterized by a "task function" transforming the element's inputs into its outputs. The task may be one of observing the environment, of making a decision, of communicating messages about the environment or the decisions, or of performing "final actions" that impinge upon the environment and thus affect the payoff to the team. A network is "implemented" by associating a set of its elements with a physically defined "instrument"—possibly a team member— thus allocating tasks among individuals. Chapter 8 defines the study of networks and analyzes a few simple examples of phenomena that are usually denoted by terms such as coordination, subordination, and so forth. The usual vocabulary is, however, less rich than the variety

1. See, for example, Cyert and March 1963 and March and Simon 1958.

of possible network configurations. The authors regret not having been able to pursue a systematic study of such configurations, to compare the (gross) payoffs they can yield under given payoff and environment conditions, and to attempt, whenever possible, an intuitive interpretation of such results. This is somewhat foreshadowed by the examples given in Chapter 8. It is followed, in Chapter 9, by a rather informal outline of the problems of costs and constraints that must be taken into account when choosing among networks, and thus allocating tasks, optimally. This choice is in itself a complex decision problem; we suppose that it is being solved by an "organizer," a fictitious construct possibly implemented in a physically identifiable "leader" or a leading group of team members. The logic of solving complex problems involves difficult paradoxes, which we meet in an empirical fashion, showing that it may be economical to ignore some information and to treat logical alternatives as if they were so many states of the physical environment. This chapter applies certain previously developed concepts—such as delegation, specialization, and postponement of decision—to the task of organizing.

In Part Three, the Epilogue, we take up again the theme of this Introduction: to specify the place of the economic theory of teams within the general perspective of the problem of efficient organization, and its relation to descriptive theories.