TRENDS IN THE THEORY OF DECISION-MAKING

IN LARGE SYSTEMS

by

Pravin Varaiya

Memorandum No. ERL-M341

14 April 1972.

ELECTRONICS RESEARCH LABORATORY

College of Engineering
University of California, Berkeley
94720

TRENDS IN THE THEORY OF DECISION-MAKING IN LARGE SYSTEMS

Pravin Varaiya

Department of Electrical Engineering and Computer Sciences
and the Electronics Research Laboratory
University of California, Berkeley, California 94720

Abstract. This paper summarizes the literature dealing with the mathematical aspect of decision-making by a set of agents each of whom has partial information and partial control. All the agents have the same utility function. The literature is classified into 3 classes: team theory, systems whose prototype is the competitive economy, and systems arranged in a hierarchy. An evaluation of the literature is attempted.

Research supported in part by the U. S. Naval Electronics Command, Contract N00039-71-C-0255 and by a Fellowship from the John Simon Guggenheim Foundation.

1. Introduction and Summary.

This paper is a brief summary of the literature dealing with the theoretical and mathematical aspects of decision-making in large systems. There is also an attempt to evaluate the work done so far, and to guess the potential of different lines of investigation. We apologize in advance for any omissions only some of which are forced by considerations of space.

The generally agreed-upon framework for decision theory consists of the following:

1. a set of agents (decision-makers, controllers, actors),

2. for each agent a preference ordering (utility function, performance criterion, pay-off function, reward function),

3. for each agent a set of permissible decisions (actions),

4. each agent makes some observation of the environment (state-of-the-world) in which the agent is operating, we call the mapping from the state-of-the-world to the various observations the information structure,

5. finally each agent has a description or model of the system i.e., the way in which the state-of-the-world and the combined actions of the agent affect consequences.

Each agent's behavior is expressed by a decision rule which is a specification of the action to be taken when a certain observation is made. We borrow from Marschak and Radner [1] the term organizational form to denote a particular choice of information structure and set of decision rules, one for each agent.

Most of the literature with which we are concerned assumes as exogenous to the decision problem being studied the set of agents, their preferences, and the permissible decisions. The focus of the study in the analysis and design of 'good' organizational forms for one of two cases:

Most of the literature with which I am familiar takes as given the set of agents, their preferences, and the permissible decisions, and is concerned with the analysis and design of 'good' organizational forms for one of two cases:

A. many ($\geq$ 2) agents with the same preference ordering,

B. many agents each with a different preference ordering.

The situation envisioned in case B is usually called a game, and we do not deal with it further. There is a third possibility, a single agent with more than one preference ordering. This can arise in dealing with questions of incentives, or with trying to find decision rules for a public agency. We do not deal with this case either.

The literature dealing with case A can be classified into three categories:

a) team theory,

b) organizational forms whose prototype is the competitive economic system, and

c) organizational forms arranged in a hierarchy.

This classification is guided by the differences in areas of application and in the mathematical techniques used. Thus team theory is the most abstract and formal treatment of our subject. The literature classified under class b) deal with decision problems whose structures are specified in considerable detail. The last class contains little abstract theory and is mostly a collection of 'case' studies.

2. Theory of teams. A static team[2] may be formalized as follows. Let S be the set of states-of-the-world. Let Y_i be the space of observations and $\eta_i : S \to Y_i$ the mapping which generates agent i's observation, $i = 1, \ldots, N$. Thus $\eta = (\eta_1, \ldots, \eta_N)$ is the information structure. For each i let D_i be the set of permissible decisions so that the team's decision rule is a function $\delta = (\delta_1, \ldots, \delta_N)$ where $\delta_i = Y_i \to D_i$ is i's decision rule. An organizational form is a pair (η, δ). Let $p(s)$, $s \in S$, be the team's common subjective probability[3] distribution concerning which state-of-the-world will occur. From the model of the system and the team's preference ordering one arrives in the usual way at a function $u(s,d) = u(s,d_1,\ldots,d_N)$, such that the team's objective is to maximize the expected utility

$$\Omega(\eta,\delta;p,u) = \sum_{s \in S} p(s)\, u(s,\delta_1(\eta_1(s)),\ldots, \delta_N(\eta_N(s))). \tag{1}$$

[2]The main reference is [1], which contains an almost exhaustive bibliography.

[3]The common subjective probability is crucial for the theory. Thus many important considerations, arising from differing expectations, cannot be formulated within this framework.

Typically the entire state S is not relevant to u. Rather, there is a (many-to-one) function $z = \zeta(s)$ such that $u(s,d) = \omega(\zeta(s),d)$ for a suitable function ω, and then (1) can be rewritten as (2).

$$\Omega(\eta,\delta;P,\omega) = \sum_{s \in S} p(s)\ \omega(\zeta(s),\delta_1(\eta_1(s)),\ldots,\ \delta_N(\eta_N(s)))$$

$$= \sum_{Z\times Y_1\times..\times Y_N} \omega(z,\delta_1(y_1),\ldots,\ \delta_N(y_N))\ P(\zeta=z,\eta_1=y_1,\ldots,\eta_N=y_N) \tag{2}$$

What sorts of questions can be raised in this formal set-up? At the simplest level, assuming η, P, ω fixed, we can search for a δ which maximizes $\Omega(\eta,\cdot,P,\omega)$, and then there is nothing conceptually different which is not present in the case of a single agent. Let us call such a δ an optimal rule. A more interesting concept is that of a person-by-person satisfactory (pbps) decision rule which is any rule δ^* such that

$$\Omega(\eta,\delta_1^*\cdots\delta_i^*\cdots\delta_N^*;P,\omega) \geq \Omega(\eta,\delta_1^*\cdots\delta_i\cdots\delta_N^*,P,\omega) \text{ for all } \delta_i \text{ and all i.} \tag{3}$$

In words, δ^* is pbps if a unilaterial decision change on the part of any agent will not increase the team's expected utility. If S is finite, then obvious convexity and differentiability conditions will imply that every pbps rule is optimal. Another result is that if p is a Gaussian distribution, if η is linear, and if u is a concave quadratic form, then there are linear pbps rules which are optimal.[4]

[4]This is surprising in that the statemnt is not in general true for dynamic teams (see below).

But these results are of mild interest only. There seem to be three potential useful directions. The first involves considering the information structure η as variable. Let

$$V(\eta;P,\omega) = \underset{\delta}{\text{Max}}\ \Omega(\eta,\delta;P,\omega)$$

be the value of η (relative to (P,ω)). One may then seek to maximize $V(\cdot;P,\omega)$ over a suitable class, or better yet, try to maximize

$$V(\eta;P,\omega) - C(\eta)$$

where $C(\eta)$ is the 'cost' of the information structure η. But both of these questions turn out to be technically not well-posed. The appropriate questions are of the following kind. Given two information structures η, η' find conditions such that

$$V(\eta;P,\omega) \geq V(\eta';P',\omega) \text{ for all } \omega.$$

For the case of a single agent there is an extensive literature[5] addressed to these sorts of issues since they arise in statistical decision theory, information theory, and economics. Extensions of some of these results to the case of teams appear to be quite straightforward. However, the form of the set of all decision rules is such that extensions of most of these results will be difficult.

In many situations the agents of the team take their decisions

---

[5] A very nice survey, together with new results, is found in [2]. [3] is still a basic reference. Some new developements are reported in [4].

sequentially rather than simultaneously. Further, the observation made by an agent depends upon the actions taken by other agents. Such a situation <u>cannot</u> be described as a static team. Consider the following simple formulation.[6] Let S, Y_i, D_i, u be as before. However η_i is now of the form

$$y_i = \eta_i(s;\hat{d}^i) \tag{3}$$

where $\hat{d}^i = (d_1,\ldots,d_{i-1},d_{i+1},\ldots,d_N)$. It turns out that the class of organizational forms that can be placed in this form is much wider than the class arising in static teams. And there are surprises. Suppose s is a Gaussian vector, the η_i are linear, and u is concave and quadratic. Then the optimal decision rule need <u>not</u> be linear [7]. Since the appearance of this 'perverse' result considerable effort has been spent in discovering which information structures imply linearity of the optimal rule. To state a recent result say that j is a precedent of i, $j > i$, if either $\frac{\partial \eta_i}{\partial d_j} \not\equiv 0$ or there exists k such that $j > k$ and $k > i$. The information structure is said to be nested if η_i contains as a component every η_j for every j precedent of i. It has been proved [8] that for nested information structures there exist linear optimal decision rules, and the proof consists in showing that the team problem is equivalent to another static team problem. The literature contains almost no results of a general nature which are relevant to 'truly' dynamic teams.

To motivate the third line of development we hazard a generalization based upon the results in static and dynamic teams concerning

[6] For a more general formulation see [5]. Note that the formulation (3) may lead to an ambiguous situation since y_i may depend on d_j which may depend on d_i which depends on y_i. For a discussion of this point see [6].

the Gaussian, linear, quadratic case. It appears that the computational effort necessary for calculating an optimal rule for a team (two or more agents) is considerably greater than for centralized decision-making that is, when there is only one agent (see e.e. [9]), so that from considerations of design one would use a team decision framework only if decentralization of decision-making were enforced because of physical or institutional constraints.[7] But implicit in some of the literature in team theory is the assertion that decentralized decision-making will result in savings in information-processing and computation. If this assertion is to be treated seriously, we need to study decision processes in teams where the agents communicate some messages to each other (and receive messages from the environment) while at the same time the agents update their decisions in the the light of the new information received.[8] We now mention two examples [13,14] where such processes are studied.[9]

The problem considered in [13] is that of synchronizing (i.e., acheiving a common frequency among) a finite set of geographically separated oscillators. A centralized solution might be to choose one of the oscillators as a 'master clock' and transmit its frequency to the rest. Each one of these would lock onto (tune into) the master clock. The decentralized scheme that was studied in [13] was to transmit

---

[7] This may be a reason for the paucity of applications. Although the subject was introduced in 1955 [10], one finds only two 'applications' papers consistently referenced [11,12].

[8] This is quite different from the usual tatônnement processes one encounters in the economic literature, since the processes we are discussing are occuring in real time. However these are similar to the tatônnement processes where 'recontracting' takes place.

[9] Two examples from economics can be found in [15,16], both compare alternative processes.

to each oscillator the frequencies of some of the rest. Each oscillator locks onto an 'average' of the received frequencies. It is shown that some simple feedback loops around each oscillator can achieve stability in the sense that the oscillators converge to a common frequency. Note that an additional benefit of the scheme is that it is reliable since a break in a transmission link will leave each connected sub-network synchronized.

Reference [14] developes a decentralized scheme for finding the maximum flow of a single commodity through a capacitated network. An agent is located at each node of the network, the agent observes the capacities and the flows in those branches of the network which are incident to that node, and in part controls the flow on these branches. Each agent continuously sends messages to its neighbors, and based upon the received messages updates its decision. The interesting points about the scheme are that (1) the messages are selected from a finite alphabet whose size depends only upon the maximum number of edges incident to a node and not on the size of the network, and (2) the messages need not be synchronized i.e., the order and the time at which messages are transmitted or received are unimportant.

We need many more results of the kind represented in [11-16] so that we can modify the team-theory framework presented earlier in such a way that these problems can be formulated on an abstract level.

3. Organizational Forms Reminiscent of the Competitive Economy. Aside from its social implications, the most celebrated property of the competitive economy derives from the proposition that it achieves an efficient allocation among thousands of agents by transmitting to each agent a

minimal amount of information, namely, the prices of the various commodities in the economy. This has led to a tremendous amount of research directed to devising algorithms for computing optimal decisions in large systems which simulate the competitive economy. With the growth of nonlinear programming the search has broadened to include any algorithm which leads to decentralization in the sense that the algorithm decomposes into many 'parallel' procedures each dealing with a small number of variables the motivation being that such algorithms will permit current computer facilities to solve problems of increasing size. Since there is available an excellent survey of this entire area [17] we will not pursue it any further except for three remarks. First of all, most of these decentralizing schemes ('decomposition algorithms' in nonlinear programming parlance) impose fairly stringent convexity conditions although promising recent developments are overcoming some of these ([18],[19]) at the cost of increased information transmission. Our second point is to emphasize that almost all of these schemes are suitable for static decision problems only.[10] Finally, uncertainty has not been introduced in any interesting way.

4. Organizational Form Arranged in a Hiearchy. We start by assuming that the system is to consist of a number of interacting subsystems each one of which is under the control of one agent. Furthermore, the subsystems are arranged in a hierarchy of levels. The subsystem in each level communicates with several subsystems in the level under it and with one subsystem in the level above. The stratification into levels is not arbitrary.

[10] [20] is one of the exceptions. It is shown that the problem of the optimal control of a dynamic linear system has a structure which makes the Dantzig-Wolfe decomposition principle applicable.

It is implied that the agents at different levels perform different 'tasks' or 'functions.' The agents at the lowest level perform 'routine' tasks while those above take 'long-term' decisions. A significant portion of control engineering designs have the structure of a hierarchical organizational form. However there are no well-articulated design philosophies. There is no mathematical theory of such forms (although [21] is an attempt to formalize the notion of a hierarchy and also gives a useful classification of different designs), and there is no serious attempt to explain why hierarchical forms are worthwhile.[11] Because of the lack of any substantive mathematical theory we merely give an example of a simple hierarchical form due to Minsky and Papert [23].

The perceptron is a pattern recognition device arranged in a two-level hierarchy as follows. The pattern is displayed on a plane divided into a square grid. Each square is either black (1) or white (0), and a pattern corresponds to a particular distribution of black and white squares. Corresponding to every set of n squares there is a lower-level agent who observes these n squares (i.e. observes n binary numbers) and computes a binary function of these n numbers and reports its value to the higher level. Thus if there are N lower level agents, they can be represented by N Boolean functions $\phi_1, \ldots, \phi_N$ where $\phi_i : \{0,1\}^n \to \{0,1\}$. There is only one agent at the higher level. This agent receives the N binary numbers $\phi_1, \ldots, \phi_N$ from the lower level, computes a linear function $\psi = \alpha_1\phi_1, \ldots + \alpha_N\phi_N$ and makes a binary decision using the sign

[11] For a thought-provoking statement explaining the presence of hierarchical structures in the natural and man-made universe in terms of reliability, adaptability and reduction of computational requirements see [22].

of ψ. Thus this agent's decision rule is given by

$$\text{Decision is } \begin{cases} \text{Yes} & \text{if } \Sigma\alpha_i\phi_i > 0 \\ \text{No} & \Sigma\alpha_i\phi_i \leq 0 \end{cases}$$

Every organizational form $\{\phi_1, \ldots, \phi_N, \psi\}$ divides the class S of all patterns into two disjoint jets S_Y, S_N by

$$s \in S_Y \quad \text{if} \quad \psi(\phi_1(s), \ldots, \phi_N(s)) > 0$$

$$s \in S_N \quad \text{if} \quad \psi(\phi_1(s), \ldots, \phi_N(s)) \leq 0.$$

We say that S_Y is the set of patterns recognized by the organizational form. Within this framework it is shown in [23] that there is no organizational form $\{\phi_1, \ldots, \phi_N, \psi\}$ which recognizes exactly the set of all connected patterns. However there is an organization form of order 3 (i.e. $n = 3$) which recognizes the set of all convex patterns. Similar questions and other classes of organizational forms are also considered.

This study [23] suggests that we may make interesting discoveries picking an appropriate class of organizational forms and ask which set of tasks can be accomplished within the chosen class. We end with illustrating this idea with one more example.

Consider a linear system

$$\dot{x}(t) = Ax(t) + Bd(t)$$

where $d(t)$ is the decision vector of dimension N. The ith component of $d(t)$, $d_i(t)$, is under the control of agent i. Suppose that the state vector $x(t)$ is decomposed into N subvectors $x_1(t), \ldots, x_N(t)$, and suppose that the ith subvector $x_i(t)$ is observed by the ith agent who is also in charge of 'controlling' it.[12] Now consider the following class of organizational forms. Suppose all the agents report their observations to a 'managing' agent who chooses a linear feedback control scheme represented by a matrix C so that the other agents now face the task of controlling the system

$$\dot{x}(t) = Ax(t) + BCx(t) + Bd(t).$$

Does there exist a matrix C which permits the task to be accomplished? We can enlarge the class of organizational forms even further to allow systems of the form

$$\dot{x}(t) = Ax(t) + BCx(t) + BKd(t)$$

which means that we are allowing the lower level agents to affect every input. A fairly complete theory for these classes of systems is now available ([25], [26]).

5. A Concluding Remark. Our objective has been to show that researchers in diverse fields are facing the problems of designing organizational forms for decision-making in large systems. The problems that they face are varied and they have devised many ingenious schemes for tackling them. Perhaps the most acute need is to develope an abstract framework within which many of these schemes appear as 'special cases'. At the same time the framework should be sufficiently operationally useful so that some detailed questions of comparisons of organizational forms can be formulated and answered.

---

[12] For example the ith agent's task may be to stabilize x_i at a prespecified target x_i^*. For an application to a problem in economics see [24].

References

1. J. Marschak and R. Radner, The economic theory of teams, (1972), New Haven: Cowles Foundation Monograph 22, Yale University Press.
2. J. Marschak and K. Miyasawa, "Economic comparability of information systems," Int'l. Econ. Rev. 9 (1968), pp. 137-174.
3. D. Blackwell and M. A. Girshick, Theory of games and statistical decisions, (1954) New York: Wiley.
4. C. B. McGuire, "Comparison of information structures," in C. B. McGuire and R. Radner (eds.) Decision and organization (forthcoming).
5. H. S. Witsenhausen, "Separation of estimation and control for discrete time systems," Proc. IEEE 59 (1971) pp. 1557-1566.
6. ________________, "On information structures, feedback and causality," SIAM J. Control 9 (1971) pp. 149-160.
7. ________________, "A counterexample in stochastic optimum control," SIAM J. Control 6 (1968) pp. 131-147.
8. Y.-C. Ho and K. C. Chu, "Team decision theory and information structures in optimal control problems Parts I," IEEE Trans. on Automatic Control AC-17 (1972) pp. 15-21.
9. C.-Y. Chong and M. Athans, "On the stochastic control of linear systems with different informations sets," IEEE Trans on Automatic Control AC-17 (1972) to appear.
10. J. Marschak, "Elements for a theory of teams," Management Science 1 (1955) pp. 127-137.
11. M. Beckmann, "Decision and team problems in airline reservations," Econometrica 26 (1958) pp. 134-145.
12. C. B. McGuire, "Some team models of a sales organization," Management Science 7 (1961) pp. 101-130.

13. I. W. Sandberg, "On conditions under which it is possible to synchronize digital transmission systems," Bell Syst. Tech. J 48 (1969) pp. 1999-2020.

14. R. Lau, R. M. Persiano and P. Varaiya, "Decentralized information and control: a network flow example," IEEE Trans. on Automatic Control AC-17 (1972) to appear.

15. R. Radner, "Allocation of a scarce resource under uncertainty: an example of a team" in C. B. McGuire and R. Radner op.cit.

16. T. A. Marschak, "Computation in organizations: the comparison of price mechanisms and other adjustment processes" in C. B. McGuire and R. Radner op.cit.

17. A. Geoffrion, "Elements of large-scale mathematical programming Part I Concepts, Part II Synthesis of algorithms and bibliography," Management Science 16 (1970) pp. 652-691.

18. A. Camacho, "Externalities, optimality, and informationally decentralized resource allocation processes," Int'l. Econ. Rev. 11 (1970) pp. 318-327.

19. M. Aoki, "Two planning processes for an economy with production externalities," Discussion paper 157, Harvard Inst. for Econ. Res. (1970).

20. G. B. Dantzig and R. M. Van Slyke, "Generalized linear programming," in D. A. Wismer (ed.) Optimization methods for large systems (1971) New York: McGraw-Hill pp. 75-120.

21. M. D. Mesarovic, D. Macko and Y. Takahara, Theory of hiearchical, multilevel systems (1970), New York: Academic Press. (See also the review of this book in IEEE Trans on Automatic Control AC-17 (1972) pp. 280-281.)

22. H. A. Simon, "The architecture of complexity," Proc. Am. Phil. Soc. 106 (1962) pp. 467-482, reprinted in J. A. Litterer (ed) Organizations: systems, control and adaptation Vol II (1969) New York: Wiley pp. 98-114.

23. M. Minsky and S. Papert, Perceptions, (1969) Cambridge, Mass: M.I.T. Press.

24. D. McFadden, "On the controllability of decentralized macroeconomic systems: the assignment problem" in H. W. Kuhn and G. P. Szegö (eds) Mathematical systems theory and economics I (1969), Lecture notes in Operations Research and Mathematical Economics, Number 11, Berlin: Springer-Verlag, pp. 221-240.

25. W. M. Wonham and A. S. Morse, "Decoupling and pole assignment in linear multivariable systems: a geometric approach," SIAM J. Control 8 (1970) pp. 1-18.

26. M. Heymann, "On the input and output reducibility of multivariable linear systems," IEEE Trans. on Automatic Control, AC-15 (1970) pp. 563-569.