Journal of Economic Behavior and Organization 9 (1988) 25–44. North-Holland

IMPERFECT DECISIONS IN ORGANIZATIONS

Toward a Theory of Internal Structure

Ronald A. HEINER*

Brigham Young University, Provo, UT 84602, USA

Received August 1985, final version received November 1986

I argue that an organization's internal structure systematically depends on how its members use information imperfectly, as distinct from their information also being imperfect. Certain reliability principles are developed to analyze the effects of decision errors: involving the probability of failing to select actions when they are superior to others based on observed information, and the probability of still selecting actions when they are inferior to others based on observed information. A two-stage reliability model is also developed in order to explicitly distinguish between imperfect information and imperfect decisions. The above results imply the need to use rules and procedures to constrain individual decision and information spaces within an organization, and the dynamic flow of information between them, thereby explaining why organizations evolve an internal decision structure in the first place. The analysis is also briefly compared with organization models that incorporate *only* imperfect information; such as 'architecture' theory by Sah and Stiglitz and 'team' theory models by Marschak, Radner and Arrow.

1. Introduction

Elsewhere I introduced a theory of reliability to explain how imperfect information *and* imperfect ability to use information influence behavior [Heiner (1983, 1985b)]. The resulting analysis implies a close link between the scope of information agents can use reliably and the set of actions they can thereby benefit from choosing. In this paper I briefly explore implications of this theory for understanding organizational structure. In this setting reliability theory implies the necessity of using rules and procedures to constrain individual agents' decision and information spaces, and the dynamic flow of information between them. These restrictions themselves affect agents' ongoing experience, and thus what information is sufficiently local or 'familiar' for them to use reliably (thereby affecting how they perceive different aspects of the organization and the attention devoted to them).

*The Institute for Advanced Study at Princeton provided generous support for the initial drafting of the paper during Spring 1985. I also thank Richard Cyert, Wesley Cohen, Joseph Stiglitz and referees for helpful comments and discussion. Of course, I alone am responsible for the paper's content.

0167-2681/88/$3.50 © 1988, Elsevier Science Publishers B.V. (North-Holland)

Standard choice theory ignores these possibilities because it assumes agents use information perfectly (by always selecting actions that maximize expected utility based on observed information). Information may be costly to acquire, but there are no decision errors in using it once it has been observed. Suppose we relax this extreme assumption of no decision errors in responding to observed information. Instead, suppose agents' decision making competence at using information is not necessarily sufficient to always respond optimally, no matter how difficult or complex their decision problems might be (separate from whether there are any costs of observing information in the first place).

Agents then face an additional dimension of uncertainty because there now exists a gap between their *competence* at using information and the *difficulty* of their decision problems (called a C-D gap). Standard choice theory implicitly assumes no C-D gap exists. Consequently, it has never investigated the behavioral implications of widening the gap; that is, of varying agents' decision making competence *relative to* the difficulty of their decision problems. When this happens, agents become progressively worse at imitating optimal decision rules. The objective is to generalize existing theory so that optimal decision rules become limiting cases within a larger set of behavioral possibilities which now opens up for analysis. The generalized theory thus ceases to be at odds with imperfect decisions due to the existence of C-D gaps, but can now itself focus directly on the behavioral symptoms arising from them.

With the above introductory comments, I now proceed to summarize certain key concepts and propositions of reliability theory, and then sketch a few implications for the organization of firms.

2. Reliability principles for modeling imperfect decisions

Let S represent the set of possible *states* of the world and X represent a set of potentially observed *information* whose individual messages are imperfectly correlated with particular states. The set A denotes an agent's decision space of choosable *actions*, where individual actions may represent randomized strategies over a set of more basic acts.

Individual consequences or *outcomes* result from different pairs of actions and states (a, s) taken from the cartesian product $A \times S$. Let p denote a given probability distribution over outcomes (that is, a probability measure over $A \times S$), and the set of all such *outcome-distributions* is denoted P. Each particular outcome-distribution $p \in P$ in general depends on the likelihood of different states arising, the likelihood of receiving different messages conditional on given states occurring, and the likelihood of agents selecting particular actions in response to different messages. A distribution p can thus be conditioned on particular states occurring, messages received, or actions selected.

In particular, let $p_{A' \times S'} \in P$ be the outcome distribution p conditional on agents selecting an action from a set $A' \subset A$ when a state from the set $S' \subset S$ occurs (leaving unspecified what potential messages are observed when $s \in S'$ occurs). Similarly, let $p_{A' \times X'} \in P$ be the outcome distribution p conditional on agents selecting an action from $A' \subset A$ when a message from $X' \subset X$ is observed (leaving unspecified the occurrence of particular states which affect the likelihood of agents observing messages from X').[1]

Let $V(p)$ represent a value function which measures achieved *performance* associated with different outcome distributions $p \in P$; that is, the performance which arises from the statistical relationships between states occurring, messages being received, and agents responding to received messages. V might be a traditional expected utility function or one of the 'non-expected' utility functions of Chew, Machina, Fishburn, etc. [see Machina (1983)]. In order to simplify notation, let $V(p_{K \times L}) = V(K; L)$ for $K \times L$ contained in either $A \times S$ or $A \times X$. Thus, for example, $V(A', X')$ measures performance conditional on agents selecting an action from $A' \subset A$ when a message from $X' \subset X$ is observed; and similarly, $V(A', S')$ measures performance conditional on agents selecting an action from $A' \subset A$ when a state from $S' \subset S$ occurs. In the special case where X' includes all potentially observed messages (so that $X' = X$), then $V(A', X)$ is written simply as $V(A')$.

Next introduce the set β of functions from X into A. Each element $B \in \beta$ is a possible *decision rule* for choosing actions in response to observed messages (i.e., how an agent behaves in responding to information). The usual procedure is to postulate an optimal decision rule (denoted $B^* \in \beta$) which maximizes the 'posterior' attainable performance contingent on observed information; that is, for each $x \in X$, $B^*(x) = a$ if and only if $V(\{a\}, \{x\}) \geqq V(\{a'\}, \{x\})$ for all $a' \in A$. With B^* we can specify those messages for which it is optimal to choose an action a (denoted X_a^*); namely, $X_a^* = \{x \in X: B^*(x) = a\}$. We can also determine the optimal states for choosing action a (denoted S_a^*) by finding those states for which action a performs at least as well as choosing any other action; namely, $S_a^* = \{s \in S: V(\{a\}, \{s\}) \geqq V(\{a'\}, \{s\})$ for all $a' \in A\}$. This is analogous to assuming agents can observe states perfectly and optimally respond according to B^* (i.e., $S_a^* = \{s \in S: B^*(s) = a\}$).

With the above notation, we can introduce certain *reliability* concepts. First consider the information potentially used by agents. Its reliability refers to how well the optimal messages for selecting an action distinguish between optimal and nonoptimal states for selecting that action. This is determined by the following conditional probabilities, $r_a^X = p(X_a^* \mid S_a^*)$ and $w_a = p(X_a^* \mid S - S_a^*)$.

[1]For example, in the special case where A, S, X are finite sets, then for each $(a^\circ, s^\circ) \in A' \times S'$, we have $p_{A' \times S'}(a^\circ, s^\circ) = \sum_X (p(a^\circ \mid x) p(x \mid s^\circ) p(s^\circ) / p(A', S'))$; where $p(A', S') = \sum_{A'} \sum_X \sum_{S'} p(a \mid x) p(x \mid s) p(s)$. Similarly, for each $(a^\circ, s^\circ) \in A' \times S$ we have $p_{A' \times X'}(a^\circ, s^\circ) = \sum_{X'} (p(a^\circ \mid x) p(x \mid s^\circ) p(s^\circ) / p(A', X'))$; where $p(A', X') = \sum_{A'} \sum_{X'} \sum_S (p(a \mid x) p(x \mid s) p(s)$. See Heiner (1984) for a precise discussion of the above when $p \in P$ are defined over algebras of both countable and non-countable sets A, S, X.

r_a^X is the chance of the optimal messages being observed when the optimal states for selecting action a occur, and w_a^* is the chance of the optimal messages still being observed when nonoptimal states for selecting action a occur. The ratio $\rho_a^X = r_a^X/w_a^X$ thus measures the accuracy of messages X_a^* in correctly revealing the optimal states for choosing action a without mistakenly arising under nonoptimal states for choosing it. Perfect information means $r_a^X = 1$ and $w_a^X = 0$ for all a.

Now apply the concept of reliability directly to agents' behavior *in responding to information*; namely, how likely are agents to choose actions when optimal messages for doing so are observed without mistakenly selecting them when nonoptimal messages for doing so arise. Thus define $r_a^B = p(B(x) = a \mid X_a^*)$, $w_a^B = p(B(x) = a \mid X - X_a^*)$, and $\rho_a^B = r_a^B/w_a^B$. The ratio ρ_a^B measures the reliability of behavior at responding to the 'right' instead of the 'wrong' messages for choosing action a (analogous to how ρ_a^X measures the reliability of information at signaling the right instead of the wrong states for choosing that action).

The purpose for introducing r_a^B and w_a^B is to incorporate the effects of a C-D gap mentioned in the introduction. In this regard, note a special feature of B^* that is not satisfied by other decision rules in β. Since B^* selects actions if and only if optimal messages for doing so are observed, it implies $r_a^{B^*} = 1$ and $w_a^{B^*} = 0$ for all a (so that $\rho_a^{B^*} = \infty$ for all a). Thus, assuming an optimal decision rule B^* necessarily locks subsequent analysis into only allowing $r_a^{B^*} = 1$ and $w_a^{B^*} = 0$. Conversely, allowing $r_a^B \leqq 1$ and $w_a^B \geqq 0$ permits one to analyze not just the properties of B^*, but also the whole domain of potential decision rules β.

We can thereby investigate when particular decision rules emerge as the predicted solution within the larger domain β instead of having to postulate a particular kind of decision rule such as B^* (or other rules such as 'satisficing', 'framing', 'myopic learning', 'risk aversion', etc.). Reliability concepts thus enable one to continue using existing decision theory tools to analyze the effects of imperfect decision making (i.e., $\rho_a^B < \infty$) in addition to imperfect information (i.e., $\rho_a^X < \infty$). Consequently, these tools need *not* be given up in order to study imperfect decisions, nor must one abstract from agents' true decision-making skills in order to use them.

To complete the necessary notation define the net gain and loss from choosing an action when optimal messages for doing so are observed as compared to nonoptimal ones. $V(A - \{a\}, X_a^*)$ is the performance achieved when other actions in A besides action a are chosen even though action a is the optimal choice given received information. $V(\{a\}, X_a^*)$ is the performance achieved if a is correctly selected when it is optimal given received information. Thus, define $g_a = V(\{a\}, X_a^*) - V(A - \{a\}, X_a^*)$ as the net 'gain' in performance from selecting action a when optimal messages for doing so are observed. Similarly, define $l_a = V(A - \{a\}, X - X_a^*) - V(\{a\}, X - X_a^*)$ as the net 'loss' from selecting action a (compared to selecting other actions besides a) when nonoptimal messages for doing so are observed. Finally, let $\pi_a = p(X_a^*)$ denote the probability of optimal messages for choosing action a (which implies $p(X - X_a^*) = 1 - \pi_a$).

Now consider what happens when agents' reliability ratios ρ_a^B for different actions are bounded instead of infinite, so that they sometimes fail to select actions in response to the appropriate messages (Type I errors because $r_a^B < 1$) and they sometimes choose actions in response to the wrong messages (Type II errors because $w_a^B > 0$). Depending on the relative incidence of these two errors, an agent's performance may or may not increase from selecting particular actions (even when optimal messages for selecting them arise with *positive* probability, $\pi_a = p(X_a^*) > 0$). Suppose in particular, agents start from a 'status quo' behavior pattern where only actions from a set A are selected. When imperfect decisions are involved, we then need to determine whether they will benefit from trying to select additional actions, taking account of the possibility that they may imperfectly use information. To do so, agents' performance from selecting an action a along with other actions must exceed that achieved when they choose only other actions besides a. That is, $V(A)$ must exceed $V(A - \{a\})$. The following theorem specifies when this happens. It was intuitively motivated in my 1983 paper, but without the more precise definitions presented here.

Theorem 1. (The Reliability Condition) For any set A and $a \in A$,

$$V(A) > V(A - \{a\}) \text{ if and only if } \rho_a^B > T_a, \tag{1}$$

where

$$\rho_a^B = \frac{r_a^B}{w_a^B} \text{ and } T_a = \frac{l_a}{g_a} \cdot \frac{1 - \pi_a}{\pi_a}. \tag{2}$$

T_a determines the minimum reliability or 'tolerance limit' (i.e., the minimum size of ρ_a^B) that must be satisfied before agents can benefit from selecting action a in response to information. Inequality (1) compares an agent's actual reliability ρ_a^B at selecting an action with the minimum required reliability T_a. If ρ_a^B exceeds T_a agents will benefit from selecting action a; otherwise, they will benefit from ignoring information about when to do so.[2]

[2]The inequality $\rho_a^B > T_a$ is thus a *diagnostic* condition which tells how performance will be affected by trying to select more actions in response to potential information. It does *not* assume agents are themselves competent to determine when or how it should be satisfied, nor that the very best or 'optimal' methods for satisfying it will necessarily evolve. However, the condition can still be used to analyze agents' behavior even if they have no special competence at applying it themselves, including inability to estimate the probability variables used in the condition (as discussed below in section 3.2). These questions are further discussed in Heiner 1983 (about selection processes sluggishly weeding out inferior performers); 1985a (about the 'tacit' nature of most evolved behavior mechanisms); and 1986 (about the 'unintended' development of social institutions.

By changing notation slightly, we can more easily see an important qualitative result which was also intuitively suggested in my 1983 paper. let $p(a)$ denote the probability that an agent will select action a over all potentially received messages (so that $p(a)=\pi_a r_a^B+(1-\pi_a)w_a^B$), and let $p^*(a)$ be the probability derived from an optimal decision rule $B^*\in\beta$. Since B^* implies $r_a^{B^*}=1$ and $w_a^{B^*}=0$ then $p^*(a)=\pi_a$ for all a (i.e., an optimal decision rule selects action a exactly as often as messages X_a^* are observed). Then substitute $p^*(a)$ for π_a in Condition (1) to obtain,

Corollary

$$V(A)>V(A-\{a\}) \text{ if and only if } \rho_a^B>\frac{l_a}{g_a}\cdot\frac{1-p^*(a)}{p^*(a)}. \qquad (1')$$

Condition (1′) allows the behavior of an optimal decision rule B^* to be compared directly to that of agents whose reliability at using information ρ_a^B is bounded instead of infinite. Note that for positive $l_a/g_a>0$ the required tolerance T_a becomes arbitrarily large as $p^*(a)$ gets smaller and smaller.

Consequently, when the number of actions increases sufficiently, *agents with bounded reliability* ($\rho_a^B<\infty$) *will no longer benefit from trying to imitate the behavior of fully optimizing agents.* Instead, their choices must be limited by processes that systematically restrict behavior away from selecting every action that optimizing agents would choose. Behavior restricted in this fashion [in order to satisfy (1) or (1′)] constitutes *rule-governed* behavior. In general, one can show that optimal decision rules will *not* necessarily approximate the behavior of rule-governed agents except as a limiting case where they become perfectly reliable at using information (i.e., only near the limit where $\rho_a^B\to\infty$ for all a).

I now briefly explore implications of the above analysis for organizational structure.

3. Rule-governed behavior and the existence of organizational structure

The above conclusion (about optimal decision rules deviating systematically from rule-governed behavior) bears directly on remarks made some time ago by Herbert Simon.

> . . . If there were no limits to human rationality administrative theory would be barren. It would consist of a single precept: Always select that alternative, among those available, which would lead to the most complete achievement of your goals. The need for an administrative theory resides in the fact there *are* limits to human rationality.
>
> *Administrative Behavior* (1957), p. 240.

Note how this passage indirectly refers to decision flexibility: fully 'rational' agents (i.e., those with perfect reliability $\rho_a^B = \infty$ at using information) should only be charged with doing whatever is best without any limitations in using any information to select any potentially beneficial action. In contrast, the condition $\rho_a^B > T_a$ implies agents will *not* necessarily benefit from having the option of selecting particular actions (where the answer depends on their reliability at using potentially available information). Thus, as noted at the beginning, the condition implies an intrinsic link between the scope of information agents can use reliably and the set of actions they can thereby benefit from selecting. Moreover, this link is implied *without* assuming any costs of acquiring information or selecting particular actions (such as transactions costs, search costs, asymmetric information, etc.). We can therefore 'endogenously' determine the size and structure of agents' decision and information spaces without introducing other assumptions beyond the core behavioral principles of the theory.[3]

3.1. The necessity of internal structure

In its abstract form the above conclusion does not single out any particular method or process by which agents' decision and information spaces are limited. However, when applied to human organizations, it means their performance will systematically depend on features of internal structure that limit the range of activities over which individual employees and managers have the authority to decide, and the scope of information they are permitted to use or become aware of in the first place (i.e., on the organization's internal authority structure and on the information network used to link separate decisions together).

Within this general picture, consider two particular features. First, some information transmitted in an organization may itself reflect how different individuals behave under various situations. Thus, an employee's reliability at using information ρ_a^B may depend on his or her ability to decipher the strategic aspects of decisions made by others. Second, an organizations's performance may also depend on the process of hiring, promoting, and otherwise shifting the decision roles of its individual members over time.

The latter hiring and promotion policies are important not only because they indirectly motivate workers to administer their assigned responsibilities more reliably (as measured by ρ_a^B for those actions under their jurisdiction), but also because those overseeing them (such as foremen or 'middle'

[3]Information and transactions costs are further discussed in Heiner (1985c and 1986). Analysis is there developed which explicitly combines the effects of costly and imperfect information with imperfectly using it. None of the qualitative results of the analysis depend on whether information costs are introduced or not (for example, agents may not benefit from using more information regardless of how reliable it might be on its own or whether there are any costs of observing it).

managers) may not themselves reliably determine the most reliable workers under any given work environment. For example, if workers are motivated to increase ρ_a^B (given their decision skills and work experience) an organization may realize substantial gains by broadening their zones of discretionary authority (without necessarily determining which workers will turn out to be the most reliable at handling further job authority). On the other hand, poorly motivated workers might deliberately reduce ρ_a^B for decisions under their control. Consequently, their decision flexibility may have to be constrained by rigid rules and procedures (which effectively limit their authority only to small decision and information spaces). This will also reduce the organization's performance compared to that attainable if its members would reliably administer more flexible responsibilities. W. Ouchi's (1981) popularized account of 'American' vs. 'Japanese' management suggests a number of aspects about motivating employees to use information more reliably (toward achieving the organization's collective goals). In particular, more reliable workers (or those motivated to be such) can be given general policies or 'philosophies' to follow, in contrast to more narrowly focused, short-term 'management objectives'. Directing workers and managers toward long-term employment relationships (for example, by using rules that constrain flexibility to grant quick promotions) may be key factors in motivating them to be more reliable.

3.2. *The impact of genuine uncertainty*

Suppose agents cannot infer probability information from their experience. For example, the environment may result from an historical process producing a succession of partially 'unique' events which are extremely difficult to infer from past experience. This kind of uncertainty goes beyond 'risk' where agents are assumed to assign well-defined (subjective) probabilities to potential events. However, regardless of how uncertain the situation appears to agents, their ongoing decisions can still generate well-defined response probabilities r_a^B and w_a^B to potential information. This possibility has a direct basis of empirical support from signal detection experiments in behavioral psychology. In particular, they illustrate how statistically well-defined r_a^B and w_a^B probabilities can arise from people's behavior even in situations where they cannot discern probabilities from their own experience. These 'behavior' probabilities (as distinct from 'subjective' Bayesian probabilities) also shift in quantitatively predictable directions to experimentally controllable parameters that change the complexity or subtlety of the decision task (such as reducing the signal-to-noise ratio until the addition of a signal amid noise is extremely difficult to detect). These statistical regularities are usually displayed in the form of ROC curves (for 'receiver operating characteristic'); see Green and Swets (1974), Heiner (1985a) [and (1985e) for related applications to legal institutions].

The key point is that a well-defined statistical effect on behavior (as measured by the r_a^B and w_a^B probabilities) can arise from any kind of decision uncertainty, including pure 'Knightian' uncertainty which prevents agents from perceiving subjective probabilities to guide their decisions. This is possible even though agents may not themselves be able to calculate the r_a^B and w_a^B probabilities generated from their own behavior [they may even lack the requisite mental equipment to conceive of such probabilities in the first place; see Heiner (1985a)].

Thus, one need not assume agents can themselves determine (or have any mental awareness of) the r_a^B and w_a^B probabilities in order to use them to analyze their behavior. In general, greater uncertainty arising from any source will tend to reduce r_a^B and increase w_a^B, thereby reducing agents' reliability ρ_a^B at using information. When applied to organizations this means additional rules and procedures may be necessary to limit the decision and information spaces of its members to those which enable them to satisfy the reliability condition $\rho_a^B > T_d$. Thus, for example, an organization may display *more* systematic features of internal structure to regulate the ongoing decisions of its members just in the case where uncertainty proceeds beyond risk to include Knightian elements.

4. Combining imperfect information with its imperfect use

Up to now we have talked primarily about the effects of imperfectly using information without explicitly incorporating the effects of imperfect information (so that ρ_a^X is bounded in addition to ρ_a^B being bounded). To do so, define $r_a^{XB} = p(B(x)=a \mid S_a^*)$, $w_a^{XB} = p(B(x)=a \mid S-S_a^*)$, $\rho_a^{XB} = r_a^{XB}/w_a^{XB}$. The ratio ρ_a^{XB} measures agents' reliability at choosing action a (under optimal instead of nonoptimal states for doing so) as *jointly* produced by imperfect information *and* imperfect use of information. ρ_a^{XB} has the following general structure [see Heiner (1984); and the appendix for a brief derivation].

Theorem 2

$$\rho_a^{XB} = \frac{r_a^X(\rho_a^B - 1) + 1}{w_a^X(\rho_a^B - 1) + 1}. \tag{3}$$

Formula (3) implies a direct tradeoff between the reliability of information and agents' reliability at using information: less reliable use of information will reduce their joint reliability at responding to any potential information source (i.e., ρ_a^{XB} necessarily drops below ρ_a^X for bounded ρ_a^B, converging on 1 as $\rho_a^B \to 1$ no matter how large the ratio $\rho_a^X = r_a^X/w_a^X$ might be). The structure of (3) is relatively simple, but [like condition (1) above] it remains valid under general mathematical conditions. We can thus apply it to a wide range of different situations and interpretations. Three applications to organizations are briefly discussed next.

4.1. Better information versus its complexity: A basic tradeoff

What sort of information is needed to predict a subtle, continually changing environment? Such an environment cannot be tracked by a single binary message (such as black–white, hot–cold, etc.). More complex compound messages could be built up from simpler messages, but then agents may have to respond to a rapid succession of compound signals in order to track quickly changing circumstances. Moreover, similar messages must not be confused with each other, nor can atypical or infrequent messages be mistakenly interpreted.

The above suggests that as information better predicts the environment it may itself become more complex and thereby more difficult for agents to use correctly.[4] Stated in terms of the reliability ratios ρ_a^X and ρ_a^B, this means that in order for information to be more reliable its own complexity may increase and thereby reduce agents' reliability in using it (so that higher ρ_a^X causes ρ_a^B to drop). In addition, recall that (3) implies agents' joint reliability ρ_a^{XB} necessarily converges to 1 as $\rho_a^B \to 1$. Consequently, beyond a certain point it will be counterproductive for agents even to try to use more reliable information. This is implied irrespective of whether there are any costs of acquiring more information.[5] Thus, organizations must be structured so as to carefully edit the amount and complexity of information used by its members, as well as how information is dynamically transmitted between them.

4.2. Limited aspiration levels

Aspiration levels are usually viewed as a way of setting attainable performance targets that reduce information processing costs (i.e., lower targets require less information to be used in order to achieve them than required for higher targets). However, at a more basic level such targets can be understood as a feedback process that indirectly regulates the scope of agents' decision and information spaces toward satisfying the reliability condition $\rho_a^B > T_a$ (but without requiring agents to determine ahead of time where the process will lead or what new direction it should take if conditions change). Using such targets can thus itself be viewed as a behavioral regularity which arises from agents' inability to foresee just how to appropriately limit their decision and information spaces.

In particular, consider what happens when agents set successively higher aspiration levels; where achieving them requires agents to deal with larger and more complex decision and information spaces. For an optimal decision

[4]Information complexity can be measured with entropy concepts used in cybernetics [see Heiner (1985b) for a brief statement].

[5]See footnote 3 above.

rule B^* this poses no problem. On the other hand, the information-complexity tradeoff described above [see also Heiner (1985b)] implies agents with imperfect decision skills will in general *not* benefit from trying to achieve higher and higher aspiration levels even if the extra information needed to achieve them is *costlessly* available. That is, beyond a certain point the very attempt to do so will be counterproductive no matter how reliable and readily available additional information might be (recall that ρ_a^{XB} unavoidably drops as ρ_a^B falls toward one). This implication also accords with the long-standing views of Herbert Simon about the need for agents to 'satisfice' rather than striving for maximum potentially attainable performance.

4.3. Selective attention to non-local information

Suppose agents have some ability to learn from their own prior experience in responding to information. In particular, suppose their reliability at using a given information source is improved through repeated use or exposure to it. Agents' past experience will then have a biasing effect on their response to information even when different kinds of messages are otherwise equally reliable. This is a special case of a general principle whereby agents' reliability at using information ρ_a^B drops as it becomes more 'non-local' in some dimension to their immediate experience [see Heiner (1985a)]. Thus, as a particular information source becomes more distant from agents' local experience, ρ_a^B will decrease and thereby cause them to ignore it as ρ_a^{XB} decreases toward one.

Note also that 'ignoring' information doesn't necessarily mean agents are consciously aware of it and still choose to ignore it (especially if the process of becoming aware of information is itself costly). Rather, agents may simply fail to develop a mental 'alertness' to such information in the first place. The latter interpretation refers to the psychology of perception, namely, a tendency to perceive only the more local or 'familiar' aspects of the environment.[6] We can thus regard such psychological tendencies as the indirect symptom of how different kinds of (otherwise perceivable) information affect agent's reliability at using them.

Now apply this to organizations. As already discussed, they must limit

[6]Differential sensitivity to information depending on its localness based on prior exposure or similarity to other familiar messages is the focus of several literatures in experimental psychology and animal behavior. See, for example, the studies of 'exposure effects' in R. Zajonc (1968, 1980), and J. Seamon, N. Brody and D. Kauff (1983); the studies of 'perceptual set and expectancy effects' in U. Neisser (1976); and studies of 'search images' and 'generalization gradients' by D. McFarland (1985) and N.J. Mackintosh (1974). Closely related to these studies is the work of Richard Day on adaptive dynamic search behavior; such as Day (1984) and the references cited therein. Local dynamic search also plays a key role in the 'satisficing' theories of Herbert Simon (1957, 1983), and in Richard Nelson and Sidney Winters' recent book on evolutionary economic change (1982).

each individual agent's scope of decision making authority and information use. Such restrictions will themselves create a localizing frame of reference which determines agents' separate paths of ongoing experience within the organization. This in turn will determine the localness of potentially usable information to different agents, and thereby what information each agent can use reliably. Consequently, there emerges a general pattern of selective attention across the members of the organization. This implication provides a theoretical basis for a recurrent theme in organization theory about selective perception and attention (as illustrated by the following two statements):

> An important proposition in organization theory asserts that each executive will perceive those aspects of the situation that relate specifically to the activities and goals of his department.
>
> Simon (1957), page 309.

> ... We need an attention-focus mechanism that transfers [a person's] demands among three possible states: active set, inactive set, not-considered set.
>
> Cyert and March (1959), page 88.

5. Organization theory with *only* imperfect information

The present analysis is now briefly compared with two applications of conventional choice theory to organizations. The purpose is to show what difference it makes if the analysis is based only on imperfect information (but the reliability of using information ρ_a^B is still assumed to be infinite).

5.1. Architecture theory

First consider the recent work of Joseph Stiglitz and Raaj Sah (1984, 1985a, b) on the 'architecture of economic organization'. In their model an organization must try to distinguish between profitable and unprofitable projects, but only has access to a noisy information variable x about a project's true profitability. The firm is divided into different units that can each sample values of x and make decisions about whether to accept or reject individual projects. A *hierarchy* is defined as a structure in which the different units that screen decisions are arranged in serial order: first one unit decides whether to accept or reject, and then only rejected projects from the first unit are screened by the second unit, and so on. If a project is accepted by one of the successive units it is implemented by the firm; otherwise, projects rejected by all units are discarded. In contrast, 'polyarchic' organization means that individual units screen projects in parallel rather than serially.

Thus, projects may be screened independently by any unit irrespective of whether they have been rejected by earlier units.

With hierarchic organization, Stiglitz and Sah show that optimal screening rules imply all successive units must set the same cutoff point for determining which values of x will lead them to accept or reject a project [page 24 of Stiglitz and Sah (1984a)]. Consequently, there is never a benefit from communicating an earlier observed x to higher units. The best later units can do is to resample to get another value of x and then apply the same cutoff point used by earlier units. This means that although units are arranged in hierarchic fashion, they all independently resample the same information and never communicate their individually sampled information to each other. Hence, despite the assumed serial screening process, there is no corresponding internal structure to the information flow or decision-making between units. They each use the same kind of information to make the same kind of decisions as every other unit.

However, Stiglitz and Sah do focus on two probabilities which can be related to reliability concepts (p_1, the probability of accepting a profitable project and p_2, the probability of accepting an unprofitable project). They interpret these probabilities only in terms of imperfect information, while still assuming optimal Bayesian decision rules B^*; so that p_1 and p_2 are special cases of r_a^X and w_a^X respectively (see for example 1984, pages 26–27; 1984b, page 1; 1985, pages 293, 295). On the other hand, by using the two-stage formula given above in Theorem 2, p_1 and p_2 can now be understood as comprising the joint interplay of *both* imperfect information *and* imperfectly using it (i.e., $p_1 = r_a^{XB}$ and $p_2 = w_a^{XB}$, thereby freeing the analysis from having to assume B^*).[7]

Thus, by allowing each unit's reliability at using information to be bounded instead of infinite, a wider range of possibilities opens up for analysis. In particular, this implies (as already discussed in section 2) that each unit must be limited in its scope of information use and decision authority, thereby requiring a certain pattern of communicating information between units to coordinate the collective decisions of the organization.

As a simple example, consider a bank granting loans to prospective applicants. Suppose some individuals have had more experience with evaluating loan applications than others, and consequently are more reliable in assessing loan information. However, the experienced individuals are relatively scarce within the firm. The bank thus saves costs by giving a number of less experienced persons the authority to judge applications, except for certain ones that must be passed to the more experienced persons (call the

[7]This interpretation also seems consistent with the intent of their work. For example, they title a recent (1985) paper 'Human Fallibility and Economic Organization'. If anything, such 'fallibility' seems more directly tied to using information imperfectly than the fact that people lack perfect knowledge.

more experienced group B and the less experienced group A). In the latter case, group B may not collect its own information, but instead usually evaluates the same information initially obtained by group A.

One possibility is that larger loans have correspondingly larger potential losses relative to interest return if they default soon into the payback period. Such loans may also be more liable to default as the resulting loan payments increase relative to an applicant's income. Both these factors will raise the required reliability for granting loans (i.e., l_a/g_a is higher and π_a is lower, both of which raises T_a). At some point as loans get larger, group B may still satisfy the reliability condition (1) while group A does not. Consequently, an upper bound is placed on group A's authority to grant loans, beyond which authorization requires group B's approval. Note again that there would be no benefit from passing information about large loans to group B if both groups were equally (or perfectly) reliable at using such information.

5.2. Team theory

Next consider the theory of teams spawned by the work of Jacob Marschak (1954, 1955), and more recently by Marschak and Radner (1972), Kenneth Arrow and Radner (1979), etc. A team is an organization composed of a number of 'divisions' linked to a central headquarters (called the 'center'). The center distributes a total resource constraint or 'capital' across the separate divisions, which they each employ using local information about their individual 'environments' to produce a common output. The team's expected output (summed over all its divisions) depends on how the distribution of total capital intermeshes with the particular stochastic features of individual division environments. This in turn depends on the information structure used by the team (i.e., on the amount and type of information transmitted back and forth between the center and the divisions).

A basic objective of team theory thus has been to analyze the effects of different information structures on team performance (each structure corresponding to given sets of information used by the divisions and the center). However, previous work has *assumed* that all members of a team always optimally use their respective information sets regardless of how large or complex they may be (i.e., information structures are varied *independently* of agents' ability to use them correctly).

On the other hand, Theorems 1 and 2 above imply the opposite view; namely, agents' reliability at using information will systematically depend on a team's information structure to the extent that it affects the size and complexity of the information sets its members are assumed to use. *A team's information structure and the reliability of its members at using it are thus intrinsically interrelated (so that the former cannot be varied independently of the latter).*

For example, consider briefly two different information structures: *first*, an 'Incomplete Communication' (or IC) type where the center sends a common message to each division (related to the total amount of capital) and each division sends its individual request for capital to the center (which partially depends on its local environment and the message from the center); *second*, a 'Complete Communication' (or CC) information structure where all divisions communicate everything about their respective environments (and about the center's total capital constraint) to all other divisions. Next expand the team by including more and more divisions linked by a common center. Note that doing so will proportionately increase the center's information set with the number of new divisions for both IC and CC information structures. However, *only* the CC structure requires a similar increase in the divisions' information sets (i.e., in the IC structure, each division must know only its own environment plus a common message from the center, neither of which expands with the number of divisions).

This might make little difference if all divisions were very similar or identical to each other. But suppose there is enough diversity between them so that the information needed to describe them accurately becomes increasingly complex as the number of divisions increases (and more nonlocal to each individual division's ongoing experience). We can then apply the information-complexity tradeoff discussed above in section 5.1, as well as the effects of more nonlocal information discussed in section 5.2. The larger information sets will more reliably indicate which actions will maximize expected output over the whole team (as measured by $\rho_a^X(n)$; where n = the number of divisions). However, their growing complexity and nonlocalness will also reduce team members' reliability at using them to guide selection of these actions (as measured by $\rho_a^B(n)$), thereby at some point reducing their joint reliability $\rho_a^{XB}(n)$ toward 1 as n grows sufficiently large. Note that this tradeoff applies to *all* team members for the CC structure, but *only* to the center for the IC structure.

Now combine the above results with the theorems by Arrow and Radner (1979) and Groves and Hart (1982). They show that (under suitable specification of the IC structure and other regularity conditions) the expected performance of IC will asymptotically approach that of CC as the number of divisions n grows indefinitely. *However, these theorems assume that both information structures are used optimally by all team members independent of the size of* n (so that $\rho_a^B(n) \equiv \infty$ for all n, which in turn implies $\rho_a^{XB}(n) \equiv \rho_a^X(n)$ for all n). When this limiting assumption is relaxed, the expected performance from the IC and CC structures may no longer converge to the same level if decision errors accumulate as n increases. In particular, if one of the two information structures produces relatively fewer decision errors as n increases, then its asymptotic performance will now exceed that achieved by the other structure (since they would otherwise approach the same asymptotic level with no decision errors). Thus, preceding results (about $\rho_a^{XB}(n)$ dropping eventually for all team members for the CC structure, but only for the center for an IC structure) imply that the IC structure will strictly outperform the CC structure as n grows sufficiently large. That is, trying to communicate all information to all members within an increasingly large team will eventually produce *worse* overall performance than allowing only incomplete information to be used.

One possibility for avoiding this result is for agents to edit a larger information set to only those parts they can use reliably. However, this may itself be an extremely difficult decision problem (which may therefore exceed agents' competence to do so optimally; so that a C-D gap also exists for deciding how to self-edit incoming information). This is more likely as the relevant portions of a larger information set are continually shifting over time, or if agents must more rapidly make ongoing decisions in response to incoming information. In either case, if the members of a team have bounded reliability $\rho_a^B < \infty$, then limiting the flow of information between them will at some point raise team performance as their numbers grow sufficiently large (compared to relying only on self-editing as the team expands).

Note also that Grove and Hart's paper does not require the existence of a coordinating 'center' to allocate total resource constraints (see page 1455). Thus, the implication of sufficiently large teams performing better with an IC rather than CC information structure applies to fully decentralized organizations which are *not* linked by any central planning agency.

The above results illustrate how basic conclusions about the relative performance of different kinds of organizations may qualitatively reverse once agents' reliability at using information is not assumed invariant to the size and complexity of their decision and information spaces.[8] This is especially important when very large organizations are involved (such as market institutions which interconnect exchange decisions across an entire economy). The above analysis thus provides theoretical support for ideas intuitively suggested by Hayek (1945, 1979) about the benefits of decentralized market organization in harnessing the productive potential of highly dispersed (and therefore localized) knowledge of exchange and productive opportunities.

[8]These issues also apply to rational expectations modeling. For example, certain key information assumptions and policy conclusions become more robust when imperfect decisions are permitted into the analysis; see Heiner (1985d). Consider also the size and complexity of message spaces in game theory settings (such as the theory of incentives, principle-agent conflicts of interest, etc.). If agents are less reliable at using complex message spaces (especially when they contain messages only infrequently sent by others – say because they would violate 'equilibrium' conditions) then different solution concepts may result than those which assume optimal use of information. For example, agents may use strategies such as *tit for tat* which react only to relatively simple messages from other players (see Robert Axelrod 1984).

6. Summary

I have argued that an organization's internal structure systematically depends on how its members use information imperfectly, as distinct from their information also being imperfect. In order to do so, certain reliability principles were developed to analyze the effects of decision errors. These involved the probability of failing to select actions when they are superior to others based on observed information, and the probability of still selecting actions when they are inferior to others based on observed information. Depending on the relative incidence of these errors, agents may or may not benefit from choosing over larger decision spaces irrespective of whether there are any adjustment costs of shifting between different actions.

A two-stage reliability model was also developed in order to explicitly distinguish between imperfect information and imperfect decisions. It implies a basic tradeoff between information and decision errors. The reason is that better information may itself become more complex, or more distant from the recurrent features of agents' past experience, thereby reducing their reliability at using it. Consequently, it may be counterproductive for agents to use more or better information irrespective of whether there are any costs of acquiring it.

In an organizational setting, the above results imply the need to use rules and procedures to constrain individual members' decision and information spaces, and the dynamic flow of information between them. These restrictions themselves constitute an explanation for why organizations evolve an internal decision structure in the first place (instead of a single agent making all decisions from information pooled over all observed messages). Within this general theme, three further applications were briefly discussed: (1) the need for a more rigid decision structure when organizations face Knightian uncertainty which prevents agents from assigning meaningful subjective probabilities to guide their decisions; (2) the use of limited 'aspiration levels' as a feedback process which regulates the decision and information spaces of individual members toward those that can be used reliably; and (3) the tendency of individual agents to selectively perceive only those events and messages closely related to their recurrent job-related experience (which in turn depends on how their decision and information spaces are internally structured within an organization).

The preceding analysis was also briefly compared with organization models that incorporate *only* imperfect information. For example, 'architecture' theory by Sah and Stiglitz (1984, 1985) also deals with how the decisions of individual agents are organized. However, because all information is used perfectly, there is no corresponding structure to the messages agents observe and the flow of information between them. Other examples are 'team' theory models by Marschak, Radner, Arrow, and others. When imperfect decisions are introduced, implications about the relative performance of different types of information structures may qualitatively reverse, especially when applied to large organizations like an economy-wide network of markets. In particular, a large team may perform better by *not* communicating all information observed by its members to each other (even assuming zero communication costs between all team members).

The paper's results and methodology are part of a larger research effort generalizing standard choice theory so that both imperfect information and imperfect decisions can be formally analyzed. Standard theory assumes agents maximize expected utility based on observed information. Such decisions can be viewed as regulating (but not eliminating) the effects of information errors on behavior. Likewise, a more general theory would analyze how agents must behave in order to regulate the effects of *both* information and decision errors (as well as how these two sources of error interact).

Within this more general setting, imperfect agents will benefit from behavioral mechanisms and rules limiting their flexibility to use (even costlessly) available information or to select potential actions. The above results illustrate how organizations become internally structured in order to regulate the decision errors of their members. In addition, these results also illustrate how regulating imperfect decisions can have systematic behavioral effects beyond those needed to cope with imperfect information alone (and despite the fact that specific decision errors may themselves be erratic and unpredictable). Thus, if our objective is to explain systematic features of institutions (and individual behavior), then analyzing how both information errors *and* decision errors affect behavior may be a more fruitful approach than assuming agents behave as if only the former errors exist.

Appendix

Proofs for Theorems 1 and 2 are here outlined.

Let V_1, V_2, V_3, V_4, equal respectively $V(A-\{a\},X_a^*)$, $V(A-\{a\},X-X_a^*)$, $V(\{a\},X_a^*)$, $V(\{a\},X-X_a^*)$. The definitions of l_a, g_a then imply $l_a=V_2-V_4$ and $g_a=V_3-V_1$. In the special case of a standard expected utility function, the linearity properties of V can be used to expand $V(A)$ and $V(A-\{a\})$ as follows:

$$V(A)=\pi_a[r_a^B V_3+(1-r_a^B)V_1]+(1-\pi_a)[w_a^B V_4+(1-w_a^B)V_2],$$

$$V(A-\{a\})=\pi_a V_1+(1-\pi_a)V_2.$$

Then subtract these expressions and rearrange terms (also recalling the above definitions for g_a, l_a) to yield,

$$V(A)-V(A-\{a\})=\pi_a r_a^B g_a-(1-\pi_a)w_a^B l_a, \tag{A.1}$$

Hence, (A.1) implies $V(a) > V(A-\{a\})$ if and only if $r_a^B/w_a^B > l_a/g_a \cdot (1-\pi_a)/\pi_a$, which is the desired result.

Heiner (1984) proves this result also holds for recent 'nonexpected' utility theories of Machina, Chew, Fishburn, and others, and considers the generality of the probability measures which interrelate the sets A, X, and S.

It remains to outline the derivation of formula (3) for ρ_a^{XB} given Theorem 2. If $s \in S_a^*$ occurs, agents can end up choosing action a either by messages X_a^* 'correctly' signaling S_a^* and agents 'correctly' responding to X_a^* by selecting a; or by messages $X - X_a^*$ occurring instead and agents still responding to $X - X_a^*$ by choosing action a. Thus, $p(B(x) = a | S_a^*) = r_a^{XB} = r_a^X r_a^B + (1 - r_a^X) w_a^B$. Similarly, if $s \in S - S_a^*$ occurs agents may also select action a if messages X_a^* still arise and agents respond to them by selecting action a; or if messages $X - X_a^*$ arise and agents still respond to them by choosing action a. Thus, $p(B(x) = a | S - S_a^*) = w_a^{XB} = w_a^X r_a^B + (1 - w_a^X) w_a^B$. Next, divide the formula for w_a^{XB} into the preceding formula for r_a^{XB}, and rearrange terms to obtain,

$$\frac{r_a^{XB}}{w_a^{XB}} = \frac{r_a^X(r_a^B - w_a^B) + w_a^B}{w_a^X(r_a^B - w_a^B) + w_a^B}. \tag{A.2}$$

Then formula (3) of the text follows immediately by dividing both the numerator and denominator of (A.2) by w_a^B.

References

Arrow, Kenneth and Roy Radner, 1979, Allocation of resources in large teams, Econometrica 47, 361–385.

Axelrod, Robert, 1983, Evolution of cooperation (Basic Books, New York).

Cyert, Richard and James G. March, 1959, A behavioral theory of organizational objectives, in: Haire, ed., Modern organization theory (Wiley, New York).

Day, Richard H., 1984, Disequilibrium economic dynamics, Journal of Economic Behavior and Organization 5, 57–76.

Groves, T. and S. Hart, 1982, Efficiency of resource allocation by uninformed demand, Econometrica 50, 1453–1482.

Hayek, Friedrik, 1945, The use of knowledge in society, American Economic Review 35, 519–530.

Hayek, Friedrik, 1979, Law, legislation and liberty, 3 volumes (University of Chicago Press, Chicago, IL).

Heiner, Ronald A., 1983, The origin of predictable behavior, American Economic Review 73, September 1983.

Heiner, Ronald A., 1983, On reinterpreting the foundations of risk and utility theory, The Institute for Advanced Study, Working paper, August 1984.

Heiner, Ronald A., 1985a, Uncertainty, signal detection and modeling behavior, in: R. Langlois, ed., Economics as a process: Essays in the new institutional economics (Cambridge University Press, New York, 1985).

Heiner, Ronald A., 1985b, Origin of predictable behavior: Further modeling and applications, American Economic Review, May 1985.

Heiner, Ronald A., 1985c, The necessity of imperfect decisions, under revision for Journal of Economic Behavior and Organization, October 1985.

Heiner, Ronald A., 1985d, Rational expectations when agents imperfectly use information, Journal of Post-Keynesian Economics, Fall-Winter 1985.

Heiner, Ronald A., 1986, Uncertainty and rule-governed behavior: On the evolution of legal precedent and rules, Journal of Legal Studies, June 1986.

Kaen, Fred and Robert Rosenman, 1986, Predictable behavior in financial markets: Some evidence in support of Heiner's hypothesis, American Economic Review, March 1986.

Litterer, Joseph, 1980, Organizations: Structure and behavior (Wiley, New York).

McFarland, D., 1985, Animal behavior (Benjamin/Cummings, Menlo Park, CA), 205–221.

Machina, Mark, 1983, The economic theory of individual behavior toward risk: Theory, evidence and new directions (Center for Organizational Efficiency, Stanford University, no. 433), October 1983, Stanford, CA.

Mackintosh, N.J., 1974, The psychology of animal learning (Academic Press, London).

Marschak, Jacob, 1954, Towards an economic theory of information and organization, in: R.M. Thrall et al eds., Decision Processes (Wiley, New York).

Marschak, Jacob, 1955, Elements for a theory of teams, Management Science 1, 127–137.

Marschak, Jacob and Roy Radner, 1972, Economic theory of teams (Yale University Press, New Haven, CT).

Neisser, U., 1976, Cognition and reality: Principles and implications of cognitive psychology (W.H. Freeman, San Francisco, CA). See references to the topics.

Nelson, Richard and Sidney Winter, 1982, An evolutionary theory of economic change (Harvard University Press, Cambridge, CA).

Ouchi, William, 1981, Theory Z (Avon Books, New York).

Radner, Roy, Decentralization and incentives, A.T. & T. Bell Laboratories Technical Report, March 1985.

Seamon, J., N. Brody and D.M. Kauff, 1983, Affective discrimination of stimuli that are not recognized: Effects of shadowing, marking and cerebral laterality, Journal of Experimental Psychology: Learning, Memory and Cognition 9, 544–555.

Simon, Herbert, 1957, Administrative behavior, 2nd ed. (Macmillan: New York).

Simon, Herbert, 1983, Reason in human affairs (Stanford University Press, Stanford, CA).

Stiglitz, Joseph and Raaj Sah, 1984a, The architecture of economic systems: Hierarchies and polyarchies, National Bureau of Economic Research, April 1984; also forthcoming in American Economic Review, September 1986.

Stiglitz, Joseph and Raaj Sah, 1984b, The economics of committees (Department of Economics, Princeton University, Princeton, NJ) September 1984.

Stiglitz, Joseph and Raaj Sah, 1984c, Human fallibility and economic organization, American Economic Review Proceedings, May 1985.

Zajonc R.B.. 1968, The attitudinal effects of mere posture, Journal of Personality and Social Psychology Monographs 9 (Part 2), 1–27.

Zajonc, R.B., 1980, Feeling and thinking: Preferences need no inferences, American Psychologist 35, 151–175.