

On the Optimality of Coarse Behavior Rules

RICHARD BOOKSTABER[†] AND JOSEPH LANGSAM[‡]

[†]*Graduate School of Business, Brigham Young University, Provo, Utah 84602, U.S.A. and* [‡]*Department of Mathematics and Statistics, Case Western Reserve University, Cleveland, Ohio 44106, U.S.A.*

(Received 5 November 1982, and in final form 22 February 1985)

Animal behavior can be characterized by the degree of responsiveness it has to variations in the environment. Some behavior rules lead to fine-tuned responses that carefully adjust to environmental cues, while other rules fail to discriminate as carefully, and lead to more inflexible responses. In this paper we seek to explain such inflexible behavior. We show that coarse behavior, behavior which appears to be rule-bound and inflexible, and which fails to adapt to predictable changes in the environment, is an optimal response to a particular type of uncertainty we call extended uncertainty. We show that the very variability and unpredictability that arises from extended uncertainty will lead to more rigid and possibly more predictable behavior.

We relate coarse behavior to the failures to meet optimality conditions in animal behavior, most notably in foraging behavior, and also address the implications of extended uncertainty and coarse behavior rules for some results in experimental versus naturalistic approaches to ethology.

1. Introduction

Animal behavior can be characterized by the degree of responsiveness it has to variations in the environment. Some behavior rules will lead to fine-tuned responses that carefully adjust to the environment. Other rules will fail to discriminate between environmental states, and will lead to coarser, more inflexible responses. There are a number of reasons for inflexibility. For example, there may be limits on neurological capacity, or the complexity of the environment may lead to excessive computational or adjustment costs for a fine-tuned response. This paper argues that coarse behavior might also arise as the optimal response to the environment, that the very nature of uncertainty might make it optimal to follow a behavioral rule that is not fully responsive to environmental cues.

The existence of inflexible behavior rules is apparent in comparing the rule-bound nature of observed behavior with the flexibility predicted by

models of optimal behavior. Models employing optimization techniques have met with a large measure of success in explaining the behavior rules of two decisions that are critical for fitness: foraging (Stephens & Charnov, 1982; Pyke, Pulliam & Charnov, 1977; Charnov, 1976; Cody, 1974) and reproduction (Charnov *et al.*, 1981; McFarland, 1977; Stearns 1976). These models inevitably involve simplification or exclusions, and are intended only as a simplified approximation of the complete optimization problem. But even so, the optimal rule often predicts more variation in behavior or a more extreme response than is observed (Jaeger & Barnard, 1981; Pulliam, 1980; Janetos 1980*a,b*; Krebs *et al.*, 1977; Goss-Custard, 1977; Perrins & Moss, 1975). This is the opposite of what would be expected given the simplified nature of the models. When errors in the predictive accuracy of the model are due to simplifications, we should generally find that observed behavior displays unexplained variation, not unaccountable rigidity.

We present a theory to explain coarse behavior. We show that inflexible behavior arises as a response to uncertainty, that the very variability and unpredictability of the environment leads to more rigid and possibly more predictable behavior.

Animal behavior is governed by a world view that determines how information is perceived and the actions it elicits. This world view may be captured by the biologist within a model of the world that delineates events, determines the probability the events will occur, and evaluates the consequences of alternative courses of action. Of course, the models are not perfect. There will be some simplifications and omissions, which, it is hoped, will lead to only small errors. But in a complex environment, there is also some chance the model will be wrong in a fundamental way and in a way that will lead to a large loss if the model is followed.

There are always some logical possibilities that are beyond any world view; and that impose severe costs to the animal in following the responses dictated by the model. These may include events such as the emergence of a new predator, a one-time climatic or environmental change, or the occurrence of some other unprecedented event—events for which the species has no past experience or “genotypical prior”, events that have never occurred in the course of the species evolution but that it must somehow take into account. By their very nature, such events cannot be delineated and explicitly considered by the model.

We term the uncertainty that is introduced by the possibility of such events *extended uncertainty*. Extended uncertainty is uncertainty about the very structure and assumptions of the model. It reflects the possibility that the environment may have features that are beyond the scope of the model, that there may be events that can occur, or consequences to action that are

taken, which cannot even be delineated, much less assigned probabilities of occurring.

Extended uncertainty is qualitatively different from uncertainty that can be incorporated into the behavioral model: Extended uncertainty cannot be represented by stochastic elements or error terms in the model, since any stochastic term must relate to specific events and be governed by some given set of distributional assumptions.

In this paper we look at the effect extended uncertainty will have on optimal behavior rules. We show that when there is extended uncertainty, the optimal behavior may be to ignore information about the state of the environment, so that the same action will be taken even when a change in the information received dictates, from conventional optimization techniques, that some other action should be taken. That is, in recognizing there is information beyond the scope of the model, the agent will appear to ignore information within the model. Such a behavior rule, which we term a *coarse behavior rule*, will lead to rule-bound behavior, behavior which appears inflexible and which does not adapt to predictable changes in the environment. We explain coarse behavior as the appropriate response to a particular form of uncertainty and not as a response constrained from the optimum by an appeal to some information or adjustment costs (Janetos & Cole, 1981).

Not only may extended uncertainty exist between the animal and the environment, it may also exist between the animal and the biologist modeling the agent's behavior. In this case, extended uncertainty may be thought of as unmodeled uncertainty. There may be events which have such infrequent occurrence that they are either ignored or perhaps never even observed by the biologist. Yet, as we will show, misspecifications of this form can lead to radical qualitative changes in behavior. The model used by the biologist may be very close to reality, so close that errors in the model may not be observable, while the conclusions for behavior arising out of the model may vary greatly from the conclusions that would come out of a true view of the world.

When viewed in this way, the existence of extended uncertainty and its implications for coarse behavior has some bearing on the question of the value of experimental versus naturalistic or ecological methods. In particular, the biologist observing behavior in the laboratory environment may not be capable of taking into account the full nature of the uncertainty under which the animal is laboring. Indeed, the laboratory setting itself, rather than controlling the environment, may actually accentuate the extended uncertainty of the animal—the experimenter may be controlling more than he understands. We would argue that this extended uncertainty will

lead to increasingly coarse behavior rules. This coarseness will naturally lead to the greater regularity in behavior and the manifestations of general principles in the laboratory environment which have been questioned as laboratory artifacts (Rozin, 1976; Johnston, 1981).

The next section of the paper presents the concepts of coarse behavior and extended uncertainty in more detail, and also gives the central theorem of the paper, that extended uncertainty may lead to coarse behavior rules. The following section discusses the extended uncertainty model within the context of the conventional model for behavior under uncertainty, and presents the necessary and sufficient conditions for behavior to be consistent with optimal behavior under extended uncertainty. These conditions provide an extremely valuable tool for using the extended uncertainty model in empirical work.

In sections 4 and 5 we present some useful applications of the extended uncertainty model. We will apply the model to foraging behavior, and to the question of the value of experimental and naturalistic methods in studying animal behavior.

2. Optimal Behavior Under Extended Uncertainty

(A) THE CONVENTIONAL UNCERTAINTY MODEL

The conventional model of decision making under uncertainty can be succinctly described as follows (Ferguson, 1967; Raiffa, 1968):

There is a set of states of nature S . An element s of S completely describes the state of the world. There is a set of actions A that the agent may choose to take. Associated with each action a in A , there is a consequence $C = c(a, s)$. The agent possesses a utility function \tilde{U} defined over the consequences.

The agent observes a random variable X , with the distribution of X depending on the state of nature s in S in a known manner. The agent has a prior probability measure on the states of nature S , and updates the prior based upon an observation of X . Thus an observation of $X = x$ yields a measure $p(\cdot | x)$ on S . In this model the random variable conveys information about the possible states. For example, the variable X might be patch size, day length, or seed type.

The rule d that maps X into A is called a *behavior rule* or a *decision rule*, and the expected utility maximization assumption implies that the agent chooses a rule d to satisfy the canonical problem, for each x , find a $d(x)$ in A such that¹

$$\sum_{s \in S} \tilde{U}(c(d(x), s))p(s | X = x) = \sup_{a \in A} \sum_{s \in S} \tilde{U}(c(a, s))p(s | X = x).$$

In this paper we will be working with uncertainty models which share the same utility function \tilde{U} and which share the same sets C , A , and X . In the canonical problem \tilde{U} appears only in terms of its value on the range of c , so we can replace \tilde{U} by a function U defined on $A \times S$, where $U(a, s) = \tilde{U}(c(a, s))$. That is, we will identify uncertainty models by the four-tuple $\{U, S, c, p\}$. It should be understood that A , C , and X have been suppressed in the notation, and that $U(a, s) = \tilde{U}(c(a, s))$.

(B) COARSE BEHAVIOR RULES

Coarseness is intended to embody the idea of inflexibility of behavior. The concept we wish to convey is that a coarse rule will be less responsive to changes in information. If d_1 and d_2 are two different behavior rules, d_1 is said to be coarser than d_2 if d_1 is constant over all subsets on which d_2 is constant. This condition is equivalent to the existence of a function f such that $f \circ d_2 = d_1$, and is similar to the property that d_2 is a sufficient statistic for d_1 .

Put in more intuitive terms, if d_1 is a coarser behavior rule than d_2 , d_1 will be less responsive to changes in information. Coarse behavior thus appears as the failure to fully discriminate between different sets of information, or the failure to fully react as dictated by the optimal behavior rule. (We will return to examples of this behavior in section 4.)

More formally, if D is a class of behavior rules, then we say that d in D is the most coarse rule in D if for each d' in D there is a function $f_{d'} \circ d' = d$. The coarseness of a behavior rule is then a question of a partition of the information set X where X is the disjoint union of sets of the form $d^{-1}\{a\}$ as a ranges over the action set A . The rule d is coarser than d' if the partition of X induced by d' is a refinement of the partition induced by d .

This definition is too strong to use as a necessary condition for one rule to be coarser than another. We may think of this as a minimal condition, but there may be other cases where we would consider one action to be coarser than another. It is possible to partition the information space in other ways that would not fulfill this definition, but may still lead to an acceptable ordering of coarseness. For example, Figs 1(a) and 1(b) give two partitions of the information space into decision rules. Figure 1(a) gives a fine partition, a slight movement in the information leads to a different action being used. Figure 1(b) gives a partition where only two actions are used, the information can vary far more without leading to a change in the action. Yet the partitions of Fig. 1(a) are not a refinement of those of Fig. 1(b), and so we could not apply the above definition to say the decision




FIG. 1. (a) The information space X is partitioned by the decision rule $d(x)$, leading to different actions being taken for slight changes in information. (b) The information space X is partitioned by the decision rule into only two actions. This decision rule may be considered more coarse than the one in (a), but the partition in (a) is not a refinement of it.

rule that maps the information space into actions in Fig. 1(b) is coarser than the decision rule employed in Fig. 1(a). How coarseness is measured depends on the application, the definition is embodied in the model. In the examples later in the paper we will place sufficient structure on the problem to specify various coarse rules in more detail.

We can define behavior rules as being coarse in an absolute sense as well as in a relative sense. The most coarse behavior rule is one that is constant over the subset of X that represents that information set. If there exists a behavior rule d in D such that for F a subset of X , $d(\cdot)|F$ is a constant function, then that behavior rule is at least as coarse as any other element of D over the range F .

To understand the concept of coarseness, it is useful to think of the action space being divided into two parts, state-specific actions and general actions. A state-specific action yields a large utility for a small number of states and a large negative utility for a large number of states. A general action need never yield large positive utilities, but suffers large negative utilities for only a small number of states. As long as information does not enable the agent to pinpoint the actual state s , and thereby yields sufficiently large probabilities of large losses for the state-specific actions, the agent will pick from among the general actions. If these general actions are few in number, and if there is a large subset of information for which general actions are optimal, then the resulting behavior rule will be coarse.

(C) THE CONCEPT OF EXTENDED UNCERTAINTY

The manifest characteristics of the conventional uncertainty model are that the state space S is exhaustive in categorizing all actions that can occur and that the consequence of each action for each state is known. Further,

the associated probabilities are assumed to be known for every state. It might be argued that no model can be usefully developed otherwise. It might seem valueless to include some state which we know nothing about and to which we cannot or will not assign a probability. For how can we say anything about something when we assume at the outset that it is something that we cannot say anything about?

Still, it is not difficult to argue that the world is filled with such events. Any complex environment will have the potential for new and unanticipated occurrences. Furthermore, models are recognized as being less than complete representations of reality. They are governed by a set of specifications and assumptions that might be wrong, and that might exclude some fundamental possibilities. It is hoped that such exclusions will not cause substantial errors, but there is always the probability, in a complex environment, for a surprise which will be disastrous for decisions based on the model, and which, being unanticipated, cannot be included in the model, even in a stochastic form.

We call the uncertainty which arises from such surprises *extended uncertainty*. Extended uncertainty arises from the existence of unanticipated events which may cause substantial errors and losses in the decisions based on the model.

Examples of extended uncertainty might include the introduction of a new and novel predator, a disease that destroys a principle food source (when the food source has always existed in abundance before), or the introduction of a chemical by man into a pristine environment. All these could be unanticipated, and may not be considered explicitly in decision making, much less be given a probability of occurring. Other examples might be the eruption of a volcano for a species that has always been in a geologically stable setting, or a one-time climatic change.

An important characteristic of extended uncertainty is that the animal cannot adapt to the surprise events before its consequences are felt. Adaptation implies the event has been incorporated into the behavioral model (or the model has been adjusted to include the event), and it is then no longer a representation of extended uncertainty. For example, if, as the animal's jungle environment changes to desert, the animal can update his behavioral model to successfully escape serious consequences, the climatic change does not represent extended uncertainty, even though it was a one-time unanticipated event. Once an event occurs and the consequences are realized, an animal may update its behavior to take it into account. The event will then move over into the conventional model. But the realization of the surprise will also increase the animal's awareness that other, yet unanticipated events may also occur. The important characteristic of these unantici-

pated events is that they have important implications for behavior, but cannot be explicitly made a part of the animal's model of behavior. By contrast, a predator which strikes with no warning, but which is known to exist in the environment, would fit in a conventional uncertainty model. The *realization* of the event might be unanticipated, but the state of being attacked by the predator can be delineated, and a probability can be assigned to it.

We can summarize the concept of extended uncertainty as follows: First, there is a non-zero probability of an event occurring. Second, if this event occurs, it will lead to a substantial loss for behavior that is based on the conventional optimization criterion. Third, the event cannot be anticipated or even delineated in the behavioral model (that is, it has zero support in the prior of the model), and there is no opportunity for adaptation before the consequences of the event are realized.

(D) THE COARSENESS OF OPTIMAL BEHAVIOR RULES UNDER EXTENDED UNCERTAINTY

The agent's belief that its world view is incorrect is to some extent dependent upon what it observes. Therefore to each action a and element of information x , the agent is assumed to assign a value $H(a, x)$. This value is a measure of the agent's anxiety in taking action a when x is observed, arising from its belief there is more to the world than it perceives. The agent then seeks to maximize $U(a, s)p(s|x) + H(a, x)$. $H(a, x)$ may itself be a function of $U(a, \cdot)$, the collection of distributions of $p(\cdot|x)$, and the distribution of X given that x is observed.

If $H(a, x)$ is independent of x , and thus can be written as $H(a)$, then it might appear that the utility function could be redefined to be $U(a, s) + H(a)$, and we could approach the question of uncertainty of world view through a respecification of the conventional uncertainty model. However, this is not the case, since $U(a, s) + H(a)$ is not a measure of the utility that is obtained when action a is taken and state s occurs. That value is given by $U(a, s)$. The function $H(\cdot, \cdot)$ is a statement of the uncertainty arising from the belief the model is unaccountably incomplete in its characterization of the world. The model cannot be expanded to include $H(\cdot, \cdot)$. For any expansion or respecification of the model, there is assumed to exist a residual uncertainty on the part of the agent that that model still does not capture. The following lemma shows how the extended uncertainty arising from an uncertain world view can lead to a restriction in the set of actions contained in the agent's behavior rule.

Lemma 1. Given the uncertainty model $\{U, S, c, p\}$, suppose there is a number M such that $\sup_{a,s} U(a, s) \leq M < \infty$. For a real number r , let $A(r) = \{a: \inf_s U(a, s) > r\}$, and suppose for some \bar{r} , $A(\bar{r})$ is not empty. Let ε be a positive number and for some $r < \bar{r}$ let $K = (M - r)/\varepsilon$. If $H(a, x)$ is set equal to $K \inf_s U(a, s)$, then the optimal behavior rule \tilde{d} under extended uncertainty satisfies the property that for all x , $\tilde{d}(x)$ is in $A(r - \varepsilon)$. (All proofs are presented in the appendix.)

Restricting the set of admissible actions will naturally lead to a coarse behavior rule. As we discussed in section 2(B), the most coarse behavior rule is one that is a constant function, one that does not respond at all to changes in information. The following theorem shows that extended uncertainty can lead to this degree of coarseness.

Theorem 1. If for some r and positive ε , $A(r) = A(r - \varepsilon) = \{a_0\}$, then there is an extended uncertainty model $U(a, s) + H(a)$ where the optimal rule is constant.

3. The Relationship Between Optimization Under Extended Uncertainty and Optimization in the Conventional Uncertainty Framework

The behavior rule the agent follows must not only be appropriate for its model, but must also be a good rule for perturbations of that model induced by the unanticipated shifts in world view. Intuitively, we argue that the set of actions that can be good for a broad set of possible models is smaller, and therefore the extended uncertainty will lead to a behavior rule that is coarse. Put another way, the finer the behavior rule, the more it depends on the correctness of the specification of the model. If the agent believes the world may behave according to any one of a number of models, then an observer attempting to explain an agent's actions by only one of these models will find the agent's behavior rule to be invariant to information which, from the observer's perspective, should elicit a more fine-tuned adjustment. While the nature of extended uncertainty precludes the agent from actually delineating the possible models, this way of looking at the relationship between extended uncertainty and coarse behavior is fruitful in relating the optimization problem under extended uncertainty to the optimization problem of the conventional uncertainty framework.

(A) A STATE PREFERENCE FRAMEWORK FOR EXTENDED UNCERTAINTY

We can formalize the concept of extended uncertainty within the state preference framework by representing the state space as a Cartesian product

$S \times T$, where S represents the states as far as they can be delineated in the model, and T represents those states that are unanticipated, and that are outside the specifications of the model. With only a minor loss in generality, we can let $T = \{0, 1\}$ where $t = 0$ is the "no surprise" state where the agent's model is correct, and $t = 1$ is the realization of an unanticipated event.

We will make a distinction between the view held by the agent and the view held by the observer who is modeling the agent's behavior. The *agent* is assumed to recognize the existence of extended uncertainty. That is, the agent recognizes surprises do occur, and that the model or world view it holds may be violated in unanticipated ways. The *observer* shares the same model as the agent when the agent restricts its model to $t = 0$, and fails to recognize the potential for the model's failure. The agent recognizes a $t = 1$ state exists, but has no information about it to add to the model, while the observer explains behavior as if the model he holds is a complete description of the world (once the error terms that are possibly built into the model are considered). In short, the agent shares the same model as the observer, but with a recognition of the possibility of mis-specification. Intuitively, we can view the fine adjustments expected by the observer as an artifact of the observer's insistence that the optimization decision is made strictly on the basis of the world view embodied in the model.

(B) THE DECISION MAKING PROBLEM FROM THE PERSPECTIVE OF THE OBSERVER, AGENT, AND OMNISCIENT PLANNER

An omniscient optimizer, an optimizer who knows the nature of the $t = 1$ states and their associated probabilities, will choose a behavior rule that solves the problem

$$\max_a \sum_s \sum_t p(s, t|x) U(a, (s, t))$$

where $p(s, t|x)$ is the conditional probability for S and T conditional on $X = x$.

The observer labors under the belief that he is proceeding as an omniscient optimizer, but in fact his knowledge is limited to the states $(s, 0)$. He characterizes the optimal behavior rule by solving the problem²

$$\max_a \sum_s p(s|x) U(a, (s, 0)).$$

The observer may be wrong but cannot fully adjust the model for these possibilities. The agent may be wrong but knows it may be wrong. The agent will follow a model that allows for surprise events, even though those

events can be monitored only through their relationship to S , if at all. To illustrate the optimal behavior rule for the agent, imagine there is an omniscient planner who knows the true model described by the complete state space $S \times T$, and knowing this true model must derive a behavior rule to hand down to the agent. (Strictly speaking, the omniscient planner is only "quasi-omniscient", since he is assumed to know the true model, but not which state in the true model will be realized. That is, the omniscient planner can delineate the T 's, and thus has resolved the extended uncertainty, but still faces the conventional uncertainty for the complete model as represented by the states $S \times T$.)

The omniscient planner is clearly in a position to hand down the optimal rule to the agent. There is a problem in allowing the agent to use this rule, however. Doing so will violate the core of the extended uncertainty paradigm, since this rule conveys all the necessary information to eliminate the surprise of the states with $t \neq 0$. To be consistent with the assumptions of extended uncertainty, the rule the omniscient planner hands down to the agent must be restricted so as not to require or contain any information the agent does not already have. That is, the omniscient planner must find a behavior rule that will maximize the expected utility of the agent subject to the constraint that the rule be an insufficient statistic for T , i.e. that it gives up no new information about T to the agent. The omniscient planner can use the full information set to construct the rule, but the resulting rule cannot convey any more information to the agent. A sufficient condition for a behavior rule d to be admissible for an omniscient planner is for there to exist a function $f(a, s)$ such that for all x

$$\sum_s f(d(x), s) p(s|x) \geq \sum_s f(a, s) p(s|x)$$

for all a in A . In particular, a function $f(\cdot)$ used by the omniscient planner might be

$$f(a, s) = \sum_t U(a, (s, t)) p(t|s).$$

The rule of the agent, then, will be to solve the problem

$$\max_a \sum_s \sum_t U(a, (s, t)) p(t|s) p(s|x).$$

If x gives no information about t , i.e. if $p(s, t|x) = p(s|x)p(t|s)$, then the rule arising from the solution to this problem will be the same as the omniscient planner's rule.³ Since the possible variations in T only appear through S , they will be unobservable to the observer.

The constrained rule $d_{\text{opc}}(x)$ will be the optimal rule for the agent given the existence of extended uncertainty. It is the rule the agent would choose

if it could "peek" at the true model in formulating its behavior rule, but then had to act "as if" it had not seen any of the information set implied by $t \neq 0$, i.e. it then had to design the behavior rule it would follow to be consistent with an optimization problem that could have arisen out of the limited information set. The distinction of d_{opc} is that it is the optimal rule relative to the extended state space when the behavior rule must be based only on the limited knowledge given by the agent's information set.

Note that the observed states of the world can have an effect on the perception of extended uncertainty. Since the probability of surprise states T is conditioned on the observed states S , unanticipated changes in the environment will lead to an increase in extended uncertainty.

(C) CONSISTENCY OF BEHAVIOR WITH MODELS OF EXTENDED
UNCERTAINTY: NECESSARY AND SUFFICIENT CONDITIONS

In the conventional uncertainty model, the behavior rule of an expected utility maximizer must satisfy conditions which, at least in principle, are subject to empirical verification.⁴ If the behavior rule of the agent does not satisfy these conditions, is it possible to determine if the behavior rule could have arisen from expected utility maximization in an extended uncertainty model that is a perturbation of the observed conventional uncertainty model? For the concept of extended uncertainty to lead to a positive theory with empirical content, there must be a way of distinguishing those behavior rules that, while inconsistent with expected utility maximization in a conventional framework, still are consistent with expected utility maximization in an extended uncertainty framework. The following two theorems present the necessary and sufficient conditions for a behavior rule to have arisen out of an extended uncertainty model.

For each action a in the range of the behavior rule d , let $R_a = \{x: d(x) = a \text{ and such that for any } a' \neq a, \sum \sum (U(a, s, t) - U(a', s, t))p(s|x)p(t|s) > 0\}$. Let Q_a be the convex hull of the probability measures $\{p(\cdot|x): x \in R_a\}$.

Theorem 2. A necessary condition for d to be consistent with expected utility maximization in an extended uncertainty model where the data x gives no information about t , is that for all a, a' in A with $a \neq a'$, $Q_a \cap Q_{a'}$ is empty.

The intuition behind this theorem is that if an action is optimal for any point in the convex hull of the conditional probability measure, then it must be optimal for all points in the convex hull. Therefore, a' cannot both be in Q_a and also not lead to the same expected utility as the optimal action a . For example, if an action a is optimal when $p(\cdot|x) = p$ and also is optimal when $p(\cdot|x) = p'$, then a necessary condition for the decision rule to be

consistent with an extended uncertainty model is that the action a also optimal for any $p(\cdot|x)$ between p and p' , i.e. for any $p'' = \alpha p + (1 - \alpha)p'$.

The sufficient condition is similar to the necessary condition, but rather than requiring a null intersection for the convex hull of any two actions, the sufficient condition requires the convex hull Q_a and the union of the convex hulls associated with all other actions to be disjoint. For a given behavior rule d , let D_a be the set of x in X with $d(x) = a$, ($D_a = \{x \in X: d(x) = a\}$), let G_a be the closed convex hull of the conditional probability measure $p(\cdot|x)$ on S where x is in D_a , ($G_a = \overline{\text{CONV}}\{p(\cdot|x): x \in D_a\}$), and let \tilde{H}_a be the closed convex hull of the union of all the $G_{a'}$ where a' is not equal to a , ($\tilde{H}_a = \overline{\text{CONV}}\bigcup_{a' \neq a} G_{a'}$).

Theorem 3. A sufficient condition for d to be consistent with expected utility maximization in an extended uncertainty model is that for each a in A , G_a and \tilde{H}_a are disjoint.

The remarkable feature of the necessary and sufficient conditions is that they are, at least in principle, empirically verifiable to the observer, even though the observer is limited to the conventional uncertainty model. That is, the observer can verify that a behavior rule is consistent with optimization within an extended uncertainty model even though he cannot know anything about its nature. The extended uncertainty paradigm thus contains some of the essential properties to lead to a positive theory for the scientist.

4. An Example of the Extended Uncertainty Framework

A number of examples of coarse behavior can be found in the literature on optimal foraging and reproduction. The question of clutch size is an important and widely studied application of the tools of optimization to animal behavior (Lack, 1947, 1948; Klomp, 1970; Stearns, 1976, pp. 12-18). It also is one example of an apparently coarse behavior rule. Perrins & Moss (1975) found that the great tit does not increase its clutch size in years of high food availability to the extent predicted to be optimal. In extreme cases, an optimal brood size of 18 was predicted while the brood size observed in the field averaged only 8.8. While they present a number of possible reasons for changes in observed clutch size being lower than that predicted by their model, such behavior may also be viewed as coarse behavior.

Optimal foraging behavior also gives a number of instances of coarse behavior, where, for example, animals fail to choose exclusively from the food that leads to the maximum nutritional intake. This behavior is illustrated by the failure of the great tit to forage according to 0-1 behavior

(Krebs *et al.*, 1977), and the failure of the red-backed salamander to fully differentiate between small and large flies in diet (Jaeger & Barnard, 1981).⁵

In this section, we will present a hypothetical case study of foraging behavior to give a more detailed illustration of the process of modeling and applying the concepts of coarse behavior and extended uncertainty. We will also provide a numerical example of the use of the necessary and sufficient conditions developed in section 3(c). The purpose of this example is to help provide a bridge between the theoretical development of the paper and the empirical tests and applications of its results.

Consider an animal that must decide whether to continue foraging in its present location, or travel to another location to forage. This decision has been treated in a number of papers on stochastic foraging which concentrate on developing optimal stopping rules for foraging in a given patch, (Oaten, 1977, Green, 1980, 1984). The animal is confronted with a trade-off between spending time seeking food and spending time avoiding predators. This trade-off is similar to that used in the stochastic foraging model of McNamara & Houston (1980). We assume the further the animal travels from its current location, the greater the risk from encountering a predator. The animal must trade off the relative security of staying in its current foraging patch, or traveling to another, undepleted, patch with higher risk but with higher abundance of prey. The decision of time spent foraging is then a function of the animal's beliefs concerning prey density and the predator population.

We can summarize this situation in the following model:

States of nature

The states of nature involve the Cartesian product of the density in the current patch relative to other undepleted patches and the density of predators. To simplify the analysis, we look at two possible states for predator density and three possible states for relative prey density.

$$S = \{0, 1\} \times \{1, 2, 3\}$$

where $(n, m) \in S$ has the following interpretation:

$n = 0$ —the predator density is high

1—the predator density is low

$m = 0$ —the current patch has prey density equal to or better than the average of other patches

1—the current patch has prey density that is between 75% and 100% of the average of other patches

2—the current patch has prey density that is less than 75% of the average of other patches.

Actions

The actions that are available to the animal are the distances to travel in searching for prey. We will only consider two possible actions.

$$A = \{0, 1\}$$

where $a \in A$ has the following interpretation:

- $a = 0$ —forage in the present patch
 1 —travel to forage in a new patch.

Information set

The animal may be expected to know the distribution of prey in an undepleted patch and the distribution of prey in the current patch from past experience, and gain an estimate of the density of predators from the number of predators it has viewed over the past month. We will represent the animal's assessment of this information in two measures, forming the two-dimensional information set:

$$X = \{0, 1, 2, \dots, 10\} \times \{0, 1, 2\}$$

where $(x, y) \in X$ is interpreted as follows:

- x = a measure of the predator density, the higher x , the greater the predator density
 y = a measure of relative prey density, the higher y , the greater the prey density.

As with the other specifications in this example, the indexing from 0 to 10 for x and from 1 to 2 for y is arbitrary. Any number of indexes and measures are possible.

Probability function

$P_{(x,y)}(s_1, s_2)$ is the probability that the animal perceives state (s_1, s_2) to occur given that an element (x, y) of the information set is observed.

We will assume the following concerning the probability function $P_{(x,y)}$

1. $P_{(x,y)}(s_1, s_2) = 0$ if $y \neq s_2$.
2. $P_{(x,y)}(s_1, y) = Q_x(s_1)$, where Q_x is a probability function on $\{0, 1\}$.
3. $Q_x(1) > Q_{x'}(1)$ if $x > x'$.

Utility function

The utility function will have the state of nature and the action of the animal as its arguments.⁶

$$W = W(a, (s_1, s_2)).$$

These functions need not be specified completely, but only some properties need to be specified to ascertain the properties of the resultant decision rules.

Based on the functions that are postulated, the optimal decision rule d^* can be computed. The action a will be taken by the optimal decision rule if

$$\sum_{s_1=0}^1 \sum_{s_2=0}^2 W[a, (s_1, s_2)]P_{(x,y)}(s_1, s_2) \geq \sum_{s_1=0}^1 \sum_{s_2=0}^2 W[a', (s_1, s_2)]P_{(x,y)}(s_1, s_2),$$

where $a' \neq a$.

Suppose, as is often the case, empirical studies show that the actual decision rule is given by a strictly determined rule d where $d \neq d^*$. A rule d is called strictly determined if, whenever $d(x) = a$, then the agent strictly prefers action a whenever x occurs. An agent will be observed not to have a strictly determined rule if different occurrences of the same datum x are associated with more than one action. The problem confronted by the biologist is to determine if d is consistent with some extended uncertainty model; that is, if the deviation of behavior from the decision rule that is optimal in a conventional uncertainty model could have arisen from optimal behavior in an extended uncertainty setting. Utilizing the results of the previous section, this question can be answered for this example by checking the following conditions

Condition 1. If there exists x, x', x'' in $\{0, 1, 2, \dots, 10\}$ and y in $\{0, 1, 2\}$ with $x < x' < x''$ such that $d(x, y) = d(x'', y)$ but $d(x', y) \neq d(x, y)$, then d does not arise from an extended uncertainty model.

Condition 2. If for all y in $\{0, 1, 2\}$ and x, x', x'' in $\{0, 1, 2, \dots, 10\}$ with $x < x' < x''$, $d(x, y) = d(x'', y)$ implies that $d(x', y) = d(x, y)$, then there exists an extended uncertainty model with d as the optimal decision rule.

Note that condition 2 is essentially the converse of condition 1. These two conditions are demonstrated in the Appendix.

As an illustration of these conditions, if, for any measure of prey abundance, the animal follows a decision rule of continuing to forage in the current patch when the measure of predator density, x , is equal to 9, and also follows the same decision rule when the measure of predator density measure is equal to 6, then for the decision rule to be consistent with an extended uncertainty model, the animal must also forage in the current patch if the measure of predator density is between 6 and 9.

This illustration can be looked at in terms of the probabilities as well. In this example, there is a higher probability of encountering a predator the higher the measure of predator density. If the animal pursues an action

when the probability of encountering a predator is p and also does so when the probability is p' , then for the decision rule to be consistent with an extended uncertainty model, the animal must continue to pursue that action if the probability of encountering a predator is anywhere in between p and p' .

Conditions 1 and 2 represent the necessary and sufficient conditions for the observed decision rule to be an optimal rule for some extended uncertainty model. That is, if Condition 2 holds, then we cannot dismiss the possibility that the animal is laboring under extended uncertainty, and is employing an optimal decision rule given the extended uncertainty, even though its behavior may appear to be suboptimal from a conventional optimization standpoint.

5. Experimental Intervention versus Naturalistic Observation: The Implications of Extended Uncertainty on Ethological Methodology

(A) THE LABORATORY AS A GENERATOR OF EXTENDED UNCERTAINTY

By its very nature, there is a high potential for the observer to miss the critical element extended uncertainty can play in behavior. This begs the question of what scientific procedure will minimize the errors that can arise from attempting to model behavior when there is extended uncertainty.

Obviously, the best approach is to view the animal from an omniscient perspective, with complete information about the environment. Since what is regarded by the animal as extended uncertainty is conventional uncertainty to the omniscient observer, the observer can follow the constrained omniscient planner's problem described in section 3 to model the animal's behavior.

The surprise events that make up extended uncertainty occur infrequently or without repetition, making it difficult to even approach the perspective of the animal in the field, much less approach the omniscient perspective, without having a deep and intimate understanding of the ecology. And it is likely the ecology that is most complex, and hence is the most difficult to fully understand, is also the one that will display the greatest level of extended uncertainty.

In contrast to field observation, the laboratory setting presents an apparent opportunity to take on the omniscient planner's perspective. It might be argued that the nature of uncertainty is under the control of the experimenter. The usual laboratory experiments for evaluating behavior under uncertainty are careful to filter out any surprise events, and use repeated trials to teach

the subject the distribution of events and their consequences. They therefore appear to fit squarely in the realm of conventional uncertainty.

However, it can clearly be argued that laboratory studies can be constructed which generate extended uncertainty for the subject in a manner that is under the control of the experimenter, thereby assuring his omniscient role. In general, this would be done by introducing one-time surprise events which are outside of the experience and expectations of the subject. An increase in the number of surprises will convey information on the extent of extended uncertainty. The state variable s leads to a probability assessment for the occurrence of surprise events, that is, a change in x will induce a change in $p(s|x)$, which will in turn induce a change in $p(t|s)$. The probability of a surprise state, $T=1$, will increase, and the existence of coarse behavior will increase with it. One might therefore think the experimental setting offers the ideal vehicle first for guarding against the introduction of coarse behavior through extended uncertainty, and second, for testing the hypothesis that extended uncertainty will lead to coarse behavior.

Unfortunately, it might equally be argued that the experimenter, in trying to control all aspects of the environment, is actually controlling more than he can realize. For example, the very nature of the laboratory environment, being so far removed from the natural environment of the subject, and being removed from it in a way that the subject may have never before encountered, (at least in genotypical terms), may serve to heighten its perceptions of extended uncertainty. The tasks and the reinforcements set before the subject, often unrelated to those it faces in nature, may do likewise. The result may be an increase in the tendency toward coarse behavior in the laboratory setting. Indeed, one might wonder if inconsistencies which appear between laboratory results and those in the field (Lea, 1979), or which seem counter to our general perceptions of human behavior (Tversky & Khaneman, 1981) are in part the result of the experimenter failing to match in the laboratory setting the nature of the extended uncertainty which the subject must face in its normal habitat.

A number of phenomenon typical to the experimental setting suggest an increase in coarse behavior in the laboratory. Perhaps the most dramatic illustration is that of experimental neurosis. Experimental neurosis is characterized by a rigidity or insensitivity to information. The "neurosis" is induced by a stimulus which is beyond the subject's control of understanding. For example, giving unescapable shock led dogs to fail to respond to the opportunity to escape when such an opportunity was later presented (Overmier & Seligman, 1967; Overmier, 1968). Dogs given very difficult problems were found to have a marked reduction in their reaction to stimulus (Pavlov, 1928). Appetitive discrimination deteriorated in rats when uncon-

trollable shocks were administered, and in dogs when the stimulus previously associated with rewards and no rewards were administered randomly (Hearst, 1965). These experiments, all cited by Seligman (1975) as examples of learned helplessness, also may serve as examples of extremely coarse behavior resulting from the extended uncertainty induced by the experimental environment. In all of these cases, the stimulus that brought on the inflexible behavior was outside the world view of the subject, in the sense that it was never encountered in nature, and was not internal to any model of the world the animal would have been able to develop. The stimulus in each case was an unanticipated event, divorced from past experience. The resulting inflexibility, failure to respond to stimulus, and lack of discrimination may all be viewed as examples of coarse behavior.

Whether the ability to control the laboratory outweighs the extended uncertainty induced by the laboratory is an empirical question, and one that no doubt will be answered in different ways for different experiments. However, even if experimental intervention does result in an increased sense of extended uncertainty in the subject, and thereby elicits a more coarse response, this result is likely to go unnoticed. This is because coarse behavior is allied with the regularity of behavior which is a premise of the experimental approach to ethology, and may therefore appear to be consistent with the expectations of the experimenter.

For example, the conflict between the naturalistic-ecological method and experimental method in learning theory hinges largely on the question of whether behavior can be described by global theories, theories that reflect a regularity and a reproducibility that carries over from one task to another. The tasks in the laboratory environment inevitably differ from those the animal faces in nature. The experimental setting can therefore only illuminate animal behavior if the principles of learning are not role specific, i.e. if there are regularities of learning that carry over from one task to another. Otherwise, "studying rats in T-mazes is about the surest possible way to learn nothing at all about anything—except perhaps rats in T-mazes" (Schwartz' commentary to Johnston (1981)).

If the laboratory induces extended uncertainty, and this in turn leads to coarse behavior, then the regularity observed in the experimental setting may be no more than a laboratory artifact. This possibility can be tested, since the regularity introduced by coarse behavior should differ in noticeable ways from the regularity of behavior that would be elicited by a global theory of behavior. For example, coarse behavior should lead to rigidity which exceeds what would appear optimal given the amount of information available to the agent. And indeed, at least in learning theory, the regularity in behavior in experiments at times does exceed that predicted by theory.

For example, the well-known experiments by Garcia, Ervin & Koelling (1966) show a regularity that contradicts the principle of equivalence of associability in general process learning theory.

While Theorem 1 shows the essential relationship between extended uncertainty and coarse behavior, in the present context it is useful to suggest ways extended uncertainty can come into the problem that would not be noticeable to the observer. Can extended uncertainty enter at a level so subtle as to be unnoticed by the experimenter, and yet have an effect which would significantly alter behavior from that predicted by the conventional model? How much extended uncertainty is enough to lead to a noticeable restriction in the set of actions, and therefore to a noticeable increase in the regularity of behavior? We will now address these questions more formally.

(B) THE APPARENT REGULARITY OF BEHAVIOR ARISING FROM
UNOBSERVED EXTENDED UNCERTAINTY

Recall the state preference representation of extended uncertainty described in section 3, where the states are defined by the Cartesian product $S \times T$, with $t=0$ being the no-surprise state where the states are consistent with the agent's world view, and $t=1$ is the surprise state, where events unaccounted for in the model occur. The observer is assumed to model behavior according to a conventional uncertainty model which recognizes only the $t=0$ state. This model is a subset of the fuller (though unknown) extended uncertainty model. The observer's model is embedded inside the extended uncertainty model, and here we want to assume the world conforms to the observer's model with all but an arbitrarily small probability. We will construct the relationship between the observer's model and the extended uncertainty model using the following definition

Definition. ϵ Perturbation; Let ϵ be a positive number. The model $\{\tilde{U}, \tilde{S}, \tilde{c}, \tilde{p}\}$ is an ϵ perturbation of the model $\{U, S, c, p\}$ if:

- (1) There exists a set T with 0 an element of T such that $\tilde{S} = S \times T$.
- (2) For a in A and s in S ,

$$c(a, s) = \tilde{c}(a, (s, 0)), \text{ and } \tilde{U}(a, (s, 0)) = U(a, s).$$

- (3) There is a δ , with $\delta < \epsilon$ such that for each x in X and s in S , $\tilde{p}((s, 0) | X = x) > (1 - \delta)p(s | X = x)$.

If an extended uncertainty model is constructed as an ϵ perturbation of the observer's model, then when $t=0$, the true model will look exactly like the observer's model, and furthermore, this will be the case with probability

greater than $(1 - \varepsilon)$. That is, the world may almost always conform to the observer's model, the parameters almost always are consistent with his model.

Embedding the biologist's model inside the extended uncertainty model in this way could lead the biologist to persist in following a mis-specified model. It may be that the $t = 1$ events occur so infrequently that they are never observed by the biologist, or they may be observed, but observed so infrequently that they are discounted as being irrelevant. The events associated with $t = 1$ may have severe consequences when they do occur, and so the ε perturbation is not insignificant in those terms. But the events are so unlikely or infrequent that they may not be considered, the observer may have no empirical evidence to suggest, in the course of his observations, that his model is not correct as stated.

The agent, on the other hand, may have had sufficient experience to know there is some chance the model may fail, but may not know the nature of the potential future failures in enough detail to revise the model to include these potential surprises explicitly. Or, as we discussed with reference to the laboratory setting, the agent may view as extended uncertainty that which appears to the biologist to be well-defined, or even non-stochastic. An important point is that past surprises will lead the animal to do more than simply update the model to take the surprises into account. It will lead the animal to change the framework of the model, to reflect in its behavior an anxiety that further unanticipated events will occur.

We will show that a recognition of extended uncertainty as an ε perturbation of the model held by the agent will restrict the set of actions. First we make the following assumptions and definitions:

Definition. For each a in A , let $m_a = \inf_s U(a, s)$.

Definition. For each a in A , let $M_a = \sup_s U(a, s)$.

Assumption 1. There exists an $M < \infty$ such that $\sup_a M_a < M$.

Assumption 2. There exists an a_0 in A such that for $a \neq a_0$, $m_{a_0} > m_a$.

Assumption 3. $m_{a_0} = 0$.

Assumption 4. There exists a probability distribution $q(\cdot)$ defined on S that can be interpreted as a prior distribution.

Assumption 2 specifies one action, a_0 , as having the greatest lower bound on utility. Assumption 3, which specifies a value for m_{a_0} , is not restrictive, as we can always translate U to make this assumption hold. The action a_0 may be thought of as a general action.

We can categorize a broad set of actions according to worst-case loss. We do so with the following definition.

Definition. $W(r)$. For $r > 0$, let

$$W(r) = \{a: \sup_{\varepsilon > 0} [\varepsilon \cdot q\{s: U(a, s) < -\varepsilon\}] > r\}.$$

$W(r)$ can be thought of as the set of actions of risk greater than r .

We show that for any model the agent may have, we can introduce extended uncertainty by way of an ε perturbation of that model with the perturbed model leading to an optimal behavior rule that restricts itself to those actions not in $W(r)$. This will imply an increase in the regularity of actions in the sense that fewer actions will potentially be chosen by the agent.

Theorem 4. Let r and ε be positive numbers, and let γ be in $(0, 1)$. Then there is an ε perturbation of $\{U, S, c, p\}$ such that if \tilde{d} is an optimal rule for the perturbed model, $\tilde{d}(x)$ is not in $W(r)$ for any x with $p(\cdot | X = x) > \gamma q(\cdot)$.

The condition $p(\cdot | X = x) \geq \gamma q(\cdot)$ can be interpreted as saying that knowing $X = x$ gives limited or weak information about S . This weak information leads to a smaller set of actions being admissible as the optimal action. The behavior rule can be thought of as being coarse in that it is limited to fewer actions. While the result here springs from the nature of information imposed on the conventional uncertainty model, it is closely related to the limitations on information arising from extended uncertainty. Extended uncertainty can be thought of as causing the agent's information to be weak. The reliability of the agent's information diminishes with the possibility that his model is wrong, and the nature of extended uncertainty is such that this unreliability cannot be fed back to revise or update the model.

In the extended uncertainty paradigm, the focus is shifted from the animal itself to the description of the environment. We argue that the essential ingredient in leading to regularity of behavior is the qualitative nature of uncertainty. Namely, an animal should tend toward more regular behavior as the degree of extended uncertainty increases. As a result, we can explain an animal manifesting the apparent regularity of behavior that is consistent with global theories of behavior in one environment—an environment with a high degree of extended uncertainty—while not necessarily doing so in another environment which contains less extended uncertainty.

We will use the development of Theorem 4 as a basis for presenting the broader result that for any conventional model, there is an extended uncertainty model that is an ε perturbation of the conventional model which will yield a behavior rule that is coarse in the strictest sense, e.g. that is constant on an infinite subset of the data. To do so, we first introduce two further assumptions.

Assumption 5. There exists an $\alpha > 0$ such that the complement of the set $G_\alpha = \{a \in A: q\{s: U(a, s) < -\alpha\} > \alpha\}$ is finite.⁷

This assumption is motivated by the fact that for $a \neq a_0$, $M_a < M_{a_0} = 0$, and so $\inf_{s \in S} U(a, s) < 0$. As α gets smaller, the complement of the set G_α gets smaller, decreasing down to a singleton, $\{a_0\}$. The assumption is that as the set of acceptable actions—the equivalent to the complement of the set of actions in $W(r)$ of Theorem 4—becomes smaller it becomes finite.

Assumption 6. There exists a γ in $(0, 1]$ such that $\{x \in X: p(\cdot | X = x) \geq \gamma q(\cdot)\}$ is infinite.

Under these two additional assumptions, we present a corollary to Theorem 5.

Corollary to Theorem 4. For any $\varepsilon > 0$, there is an ε perturbation of the model $\{U, S, c, p\}$ such that there is an optimal behavior rule d for the perturbed model that is constant on an infinite subset of X .

This result is a more structured version of the main theorem of the paper, that an arbitrarily small potential for surprise, as measured by the ε perturbation, can lead to a behavior rule that is coarse. The realization of the potential for a surprise may be thought of as adding more information to the agent's model: Including the recognition that the model may fail in an unanticipated way adds information that is beyond the scope of the conventional model as stated. In this context, the result of the corollary implies that more information—information that leads to uncertainty about the model—may lead to more rigid behavior. Or put another way, having information about what is beyond scope of model may lead the agent to ignore information in the model. The optimal rule for the agent who has the information embodied in the ε perturbation appears coarse to the observer who does not have or who disregards the potential for surprise states. The agent's additional information may make it appear as though it is ignoring information, or possibly is responding in a more general, regular way to the stimuli of its environment.

6. Conclusion

Extended uncertainty is related to model mis-specification. The animal laboring under extended uncertainty has an anxiety that its world view is incomplete. This anxiety means any contemplated action must not only be

beneficial for its model of the world, but also must be likely to be beneficial given the possible but unexplored perturbations of that model. Fine-tuned or highly specialized actions will generally have fewer benefits if the world view of the animal is violated, and as a result a wider set of these actions will be excluded from the animal's repertoire as extended uncertainty increases. The resulting behavior will be coarse; it will appear to be inflexible. In failing to make fine adjustments, it will appear the animal is not fully discriminating or following environmental cues.

The scientist who fails to recognize the extended uncertainty facing the animal will find unaccountable rigidity of behavior. This may lead to a perception the animal is behaving suboptimally, or is displaying a regularity of response that suggests its behavior follows global or general rules.

The necessary and sufficient conditions for a behavior rule to be consistent with optimal behavior under extended uncertainty, presented in section 2(D), can be verified with only a knowledge of the constrained model. It is therefore possible to test whether deviations from apparently optimal behavior can be explained by the existence of extended uncertainty.

If various environments can be ranked according to the degree of extended uncertainty, then the theory would indicate a tendency for an animal to display increasingly coarse behavior as the degree of extended uncertainty increases. Further, the particular coarse behavior displayed should represent a strategy which, although possibly not leading to actions that are optimal according to the conventional model in any state, give actions which do not involve a great loss even in the worst (surprise) state. This ranking of environments is possible both in the laboratory and in the naturalistic setting.

While we have argued the laboratory may intrinsically have extended uncertainty, the degree of that uncertainty can be varied by varying the stimulus and reward away from what the animal can anticipate or what the animal has previously experienced. The experiments we have discussed that lead to experimental neurosis are suggestive of this approach.

In the natural setting, the task of discovering and measuring extended uncertainty is more difficult, since by its very nature it may remain unrevealed to the observer. However, it may be reasonable to assume extended uncertainty increases as a function of the complexity of the environment. With this assumption, the theory would imply coarser or more generalist behavior in more complex and varied environments. For example, the escape mechanism of the cockroach follows a very coarse rule: the cockroach moves in the opposite direction of the gusts of wind that signal the approach of a predator. This rule ignores information about the environment—visual and olfactory cues for example—which it would seem an "optimal warning system" should take into account. However, the very coarseness of the rule

that will lead it to be suboptimal for any given model of predation makes it satisfactory for wide range of unanticipated predators. And indeed the long history of the cockroach and the varied settings in which it has survived suggests this extended uncertainty. By contrast, an animal which has found a well-defined and unvarying niche may follow a specialized behavior rule which depends on the correctness of its world view. This is not to say its environment may not have a high degree of conventional uncertainty, but the nature of possible events and their likelihood will be well known. Thus, some insight into the nature of adaptation and the flexibility of behavior programs may be gained by considering the degree of extended uncertainty facing the agent, and by looking into how extended uncertainty is correlated with the more apparent attributes of the agent's environment.

What are the attributes of the environment that are likely to be related to a high degree of extended uncertainty, and thereby lead to coarse behavior? First, the degree of extended uncertainty will be related to informational instability where the implications of cues vary in unanticipated ways, or where new events occur that have no previous counterpart. As we have already suggested, this instability will be more likely in a complex setting, since there are more parameters to change. It will also be more likely in a dynamic setting, where change and interaction across segments of the environment are frequent. Obviously, complexity and change alone do not guarantee even conventional uncertainty. Over time, the complexity can be modeled, and the change may become predictable. But in such a setting, the chances for the new and novel increase. Second, since coarse behavior is the result of the agent moving to a more restricted set of actions, the more distinct the division between specialist and generalist actions, and the smaller the set of general actions, the more likely coarse behavior will be.

Notes

¹ Here we have made the tacit assumption that S is countable and that the supremum is actually attained. The state space S is taken to be countable for mathematical convenience. The analysis of the paper still applies if S uncountable, but then we must worry about the question of measurability. If S is not countable, then this problem is replaced with

$$E_{X=x}U(\tilde{c}(d(x), s)) = \sup_{\tilde{a}} E_{X=x}U(c(\tilde{a}, s)).$$

That the supremum is obtained is fundamental for an optimal rule d to exist, and we retain this assumption throughout the paper.

² The relationship between $p(s)$ and $p(s, t|x)$ may be taken to be either that $p(s) = p(s, 0)/\sum_t p(s, 0)$, or that $p(s) = \sum_t p(s, t)$. In the event that s and t are

stochastically independent, these two choices are the same. If $p(s)$ is related to $p(s, t)$ by the equation $p(s) = \sum_t p(s, t)$, then $p(s, t) = p(t|s)p(s)$. Then in the extended uncertainty model the agent maximizes $\sum_s \sum_t U(a, s, t)p(t|s)p(s)|X$. For this footnote only, we have denoted $p(\lambda|x)$ by simply $p(\lambda)$.

³ The resulting rule $d_{\text{opc}}(x)$ must satisfy the condition that $d_{\text{opc}}(x_1) = d_{\text{opc}}(x_2)$ whenever $p(s|x_1) = p(s|x_2)$ for all s in S , and must depend on the distribution $\text{dist}(S \times T|X = x)$ only in the form of $\text{dist}(S|X = x)$.

⁴ In particular, let an agent have a utility function $U_0(a, s)$, where a is a control variable (an action in the space of actions A), and s is a realization of a random variable (s is an element of the state space S). A behavior rule d maps the information space X into the action space A . If the agent is an expected utility maximizer, then its behavior rule will satisfy the condition that

$$\sum_s U_0(d(x), s)p(s|x) \geq \sum_s U_0(a, s)p(s|x)$$

for any a in A .

⁵ We have already mentioned that care must be taken in developing the definition of coarse behavior for the particular application. The 0-1 selection rule in foraging provides an excellent example of how the model must define the context of coarseness. At first, it might appear a 0-1 selection rule would be a coarse rule, since it only differentiates between two actions, while a rule that chooses among many combinations of food types would be less coarse, since it allows for a wider variety of food mixes. Any application of a theory requires a careful identification of the variables in the theory and their relationship to the variables in the application. In the foraging example, the states of nature include a full description of the animal's environment, the information the animal receives includes a knowledge of its preferences and its position in the environment, and an action is a complete specification of the sequence of activities that leads to eating the prey. This will include not only the mix of food eaten, but also the procedure for finding the food and eating it. For example, the specification of an action as eating wheat and corn in a two-to-one ratio is not a complete specification of the action in foraging behavior, for it fails to specify how the food was obtained and eaten. When the process of finding the food and the effort in moving to it and eating it is included as part of the action, it is clear that a coarse behavior rule for foraging is not eating all of one food type, but rather is simply eating whatever is immediately available. The action of eating whatever is in front of it will lead the animal to a food mix that is as variable as the distribution of food in its environment. A less coarse behavior rule will be deliberately searching out prey in a mix that deviates from the mix available in nature.

In developing the behavior rule, we must assume all of the behavior rules under consideration select actions relative to the physical orientation of the animal. More formally, we assume there is a function L from the set of information X into a set of orientations Q such that for any of the behavior rules d there exists a function f_d such that $f_d \circ d = L$. Essentially, $L(x)$ gives the orientation of the animal, and the behavior rule is defined relative to that orientation. The function L defines a partition on the space of information X in the obvious way: X is the disjoint union of the sets $L^{-1}\{y\}$ as y ranges over the set Q . By assumption, each behavior rule d in the set D factors through L , so L is in a sense more coarse than any rule d .

For a behavior rule in D , then, we can define coarseness for foraging behavior by the measure

$$K(d) = \sup_{\substack{y \in Q \\ x_1 \in L^{-1}\{y\} \\ x_2 \in L^{-1}\{y\}}} M(d(x_1), d(x_2)),$$

where M is any metric on the space of actions A . (Of course the value of $K(d)$ depends on the choice of the metric, but $K(d) = 0$ for any one metric on A implies $K(d) = 0$ for every metric on A .)

⁶The utility function might be further specified by the arguments $R(a, (s_1, s_2))$, which is a measure of the encounters the animal has with predators when action a is chosen and state (s_1, s_2) occurs, and $L(a, (s_1, s_2))$, which measures the animal's foraging success if it adopts action a and state (s_1, s_2) occurs. Then $W(\cdot)$ could be expressed as

$$W(a, (s_1, s_2)) = W'(R(\cdot), L(\cdot)).$$

⁷Let $A_n = \{a \in A: q\{s: U(a, s) < -1/n\} > 1/n\}$. Then $A_n \subset A_{n+1}$. If a is not in $\bigcup_{n=1}^{\infty} A_n$, then for each n

$$q\{s: U(a, s) < -1/n\} \leq \frac{1}{n},$$

hence

$$q\{s: U(a, s) < 0\} = 0.$$

If for each s in S , $q(\{s\}) \neq 0$, we must have $m_a = m_{a_0}$. Hence, the complements of A_n are decreasing down to a singleton $\{a_0\}$. Unfortunately, this does not impose the condition that for some n , A_n^c is finite. This we must assume.

REFERENCES

- CHARNOV, E. L. (1976). *Theor. Pop. Biol.* **9**, 129.
 CHARNOV, E. L., LOS-DEN HARTOGH, R. L., JONES, W. T. & VAN DEN ASSEM, J. (1981). *Nature* **289**, 27.
 CODY, M. L. (1974). *Science* **183**, 1156.
 FERGUSON, T. S. (1967). *Mathematical Statistics: A Decision Theoretic Approach* New York: Academic Press.
 GARCIA, J., ERVIN, F. R. & KOELLING, R. A. (1966). *Psychonomic Sci.* **5**, 121.
 GOSS-CUSTARD, J. D. (1977). *Anim. Behav.* **25**, 10.
 GREEN, R. F. (1984). *Am. Nat.* **123**, 30.
 GREEN, R. F. (1980). *Theor. Pop. Biol.* **18**, 244.
 HEARST, E. (1965). *J. Exp. Anal. Behav.* **8**, 135.
 JAEGER, R. G. & BARNARD, D. E. (1981). *Am. Nat.* **117**, 639.
 JANETOS, A. C. (1980a). *Behav. Ecol. Sociobiol.* **7**, 107.
 JANETOS, A. C. (1980b). *Foraging Tactics of two guilds of web-spinning spiders*. PhD dissertation, Princeton University.
 JANETOS, A. C. & COLE, B. J. (1981). *Behav. Ecol. Sociobiol.* **9**, 203.
 JOHNSTON, T. D. (1981). *Behav. Brain Sci.* **4**, 125.

- KLOMP, H. (1970). *Ardea* **58**, 1.
- KREBS, J. R., ERICHSEN, J. T., WEBBER, M. I. & CHARNOV, E. L. (1977). *Anim. Behav.* **25**, 30.
- LACK, D. (1947). *Ibis* **89**, 302.
- LACK, D. (1948). *Ibis* **90**, 25.
- LEA, S. E. G. (1979). *Anim. Behav.* **27**, 875.
- LOGAN, F. A. (1965). *J. Comp. Physiol. Psychol.* **59**, 1, 246.
- McFARLAND, D. J. (1977). *Nature* **269**, 15.
- MCNAMARA, J. & HOUSTON, A. (1980). *J. theor. Biol.* **85**, 673.
- OATEN, A. (1977). *Theor. Pop. Biol.* **12**, 263.
- OVERMIER, J. B. (1968). *J. Exp. Psychol.* **78**, 340.
- OVERMIER, J. B. & SELIGMAN, M. E. P. (1967). *J. Comp. Physiol. Psychol.* **63**, 23.
- PAVLOV, I. P. (1928). *Lectures on Conditional Reflexes: The Higher Nervous Activity (Behavior) of Animals, Volume 1*. London: Lawrence and Wishart.
- PERRINS, C. M. & MOSS, D. (1975). *J. Anim. Ecol.* **44**, 695.
- PULLIAM, H. R. (1980). *Ardea* **68**, 75.
- PYKE, G. H., PULLIAM, H. R. & CHARNOV, E. L. (1977). *Q. Rev. Biol.* **52**, 137.
- RAIFFA, H. (1968). *Decision Analysis*. Reading, Mass: Addison-Wesley.
- ROZIN, P. (1976). In: *Progress in Psychobiology and Physiological Psychology, Volume 6*. pp. 245-280. New York: Academic Press.
- SELIGMAN, M. E. P. (1975). *Helplessness: On Depression, Development, and Death*. San Francisco: W. H. Freeman.
- STEARNS, S. C. (1976). *Q. Rev. Biol.* **51**, 3.
- STEPHENS, D. W. & CHARNOV, E. L. (1982). *Behav. Ecol. Sociobiol.* **10**, 251.
- TVERSKY, A. & KHANEMAN, D. (1981). *Science* **211**, 453.

APPENDIX

Proof of Lemma 1. If a is in $A(r)$, which is nonempty by assumption, then $\sum U(a, s)p(s|x) + H(a, x) \geq \sum rp(s|x) + Kr = (K+1)r$. If a is not in $A(r - \epsilon)$, then $\sum U(a, s)p(s|x) + H(a, x) < M + K(r - \epsilon)$. Since $M = r + \epsilon K$, for a not in $A(r - \epsilon)$, $\sum U(a, s)p(s|x) + H(a, x) < r + \epsilon K + Kr - \epsilon K = (K+1)r$.

Proof of Theorem 1. As in the above lemma, set $H(a) = K \inf_s \{U(a, s)\}$. By the above lemma, we find then that for $K = (M - r)/\epsilon$, the optimal rule takes its values on $A(r - \epsilon)$, which contains only a_0 . Thus the optimal rule is a constant function whose only value is a_0 .

Proof of Theorem 2. If the probability measure $\mu(\cdot)$ is in Q_α , then there exists $\alpha_1 \dots \alpha_n$ positive numbers summing to one with

$$\mu = \sum_{i=1}^n \alpha_i p(\cdot | x_i)$$

where each x_i is in R_α . If a' is any action other than a , we have

$$\begin{aligned} & \sum [U(a, s, t) - U(a', s, t)] \mu(s)p(t|s) \\ &= \sum \sum_{s,t} [U(a, s, t) - U(a', s, t)] \alpha_i p(s|x_i) p(t|s) \\ &= \sum_{i=1}^n \alpha_i \sum_{s,t} [U(a, s, t) - U(a', s, t)] p(s|x_i) p(t|s) > 0. \end{aligned}$$

If μ is in $Q_a \cap Q_{a'}$ where $a \neq a'$, then

$$\sum [U(a, s, t) - U(a', s, t)]\mu(s)p(t|s) > 0$$

and

$$\sum [U(a', s, t) - U(a, s, t)]\mu(s)p(t|s),$$

which equals

$$-\sum [U(a, s, t) - U(a', s, t)]\mu(s)p(t|s) > 0,$$

which is impossible.

Theorem 2 does not directly provide an observable test to determine if a decision rule d is consistent with an extended uncertainty model, since R_a and hence Q_a depend on the unknown function $U(a, s, t)$. However Theorem 2 does suggest a simple modification which does provide a meaningful test on the observed rule d .

Let $\tilde{R}_a = \{x: d(x) = a\}$ and suppose

$$\sum_{i=0}^m \alpha_i p(\cdot | x_i) = \sum_{j=0}^k \beta_j p(\cdot | y_j)$$

where $\alpha_i > 0$, $\beta_j > 0$, $\sum \alpha_i = \sum \beta_j = 1$, $d(x_i) = a$, $d(y_j) = a'$, and $a \neq a'$. Then the rule d is consistent with an extended uncertainty model for an agent only if whenever x_i or y_j occurs the agent is indifferent between a and a' .

Proof of Theorem 3. Let H_a be the closed convex hull of \tilde{H}_a and the zero probability measure. If G_a and \tilde{H}_a are disjoint then so are G_a and H_a , so G_a and H_a are disjoint closed compact subsets of R^n where n is the number of states in S . (If S is infinite take the w^* closure and G_a and H_a are both w^* compact.) For each a , there exists a linear functional (w^* continuous) $\bar{\Delta}_a$, and an r_a and s_a real numbers such that

$$\bar{\Delta}_a(\mu) \geq r_a > s_a \geq \bar{\Delta}_a(v)$$

for each μ in G_a and v in H_a . Since $\Delta_a(0) = 0$ and 0 is in H_a , r_a must be positive; thus letting Δ_a equal $(1/r_a)\bar{\Delta}_a$, we can find a linear functional Δ_a such that for each μ in G_a and v in H_a , $\Delta_a(\mu) \geq 1$ while $\Delta_a(v) < 1$. Therefore if μ is in G_a , and $a' \neq a$, we have $\Delta_a(\mu) \geq 1 > \Delta_{a'}(\mu)$. For each a , Δ_a can be identified with a function, $f_a(s)$, where $\Delta_a(\mu) = \sum f_a(s)\mu(s)$, (or $\Delta_a(\mu) = \int f_a(s) d\mu(s)$). Let ε be between 0 and 1. We can let $T = S \cup \{0\}$ (where 0 is assumed disjoint from S) as a set, and we can define $U(a, s, t)$ and $p(t|s)$ as follows:

$$U(a, s, 0) = U_0(a, s)$$

$$p(t=0|s) = 1 - \varepsilon$$

$$U(a, s, t') = U_0(a, s)$$

$$p(t=t'|s) = 0 \text{ if } s \neq t'$$

$$U(a, s, s) = (f_a(s) - (1 - \varepsilon)U_0(a, s))/\varepsilon$$

$$p(t=s|s) = \varepsilon.$$

It follows then that:

$$\begin{aligned} & \sum_s \sum_t U(a, s, t)p(s|x)p(t|s) \\ &= \sum_s \left[U_0(a, s)(1 - \varepsilon) + \left(\frac{\varepsilon[f_a(s) - (1 - \varepsilon)U_0(a, s)]}{\varepsilon} \right) \right] p(s|x) \\ &= \sum_s f_a(s)p(s|x) = \Delta_a[p(\cdot|x)]. \end{aligned}$$

We then find that if $d(x) = a$ and a' is not equal to a , that

$$\sum_s \sum_t U(a, s, t)p(s|x)p(t|s) \geq 1 > \sum_s \sum_t U(a', s, t)p(s|x)p(t|s)$$

so that d is a behavior rule that would have arisen if $U(a, s, t)$ were the agent's utility function.

Proof of Conditions 1 and 2. Condition 1 follows immediately from the discussion after Theorem 2. Condition 2 is demonstrated as follows and is a consequence of Theorem 3. Let $I = \{(x, y) : d(x, y) = 0\}$ and let $J = \{(x, y) : d(x, y) = 1\}$. We need to show that (*) cannot occur:

$$\sum_{(x,y) \in I} \alpha_{(x,y)} P_{x,y} = \sum_{(x,y) \in J} \beta_{(x,y)} P_{x,y} \tag{*}$$

where

$$\alpha_{(x,y)} \geq 0 \quad \beta_{(x,y)} \geq 0 \tag{1}$$

$$\sum_{(x,y) \in I} \alpha_{(x,y)} = \sum_{(x,y) \in J} \beta_{(x,y)} = 1. \tag{2}$$

Suppose (*) occurs. Then for each w in $\{0, 1, 2\}$

$$\sum_{(x,w) \in I} \alpha_{(x,w)} P_{x,w}(\cdot w) = \sum_{(x,w) \in J} \beta_{(x,w)} P_{(x,w)}(\cdot w)$$

and thus

$$\sum_{(x,w) \in I} \alpha_{(x,w)} Q_x = \sum_{(x,w) \in J} \beta_{(x,w)} Q_x.$$

Since $Q_x(0) + Q_x(1) = 1$,

$$\sum_{(x,w) \in I} \alpha_{(x,w)} = \sum_{(x,w) \in J} \beta_{(x,w)}.$$

If $\sum_{(x,w) \in I} \alpha_{(x,w)} \neq 0$, let $\alpha'_{(x,w)} = \alpha_{(x,w)} / \sum_{(x,w) \in I} \alpha_{(x,w)}$ and let $\beta'_{(x,w)} = \beta_{(x,w)} / \sum_{(x,w) \in J} \beta_{(x,w)}$. Then $\sum_{(x,w) \in I} \alpha'_{(x,w)} Q_x = \sum_{(x,w) \in J} \beta'_{(x,w)} Q_x$ with

1. $\alpha'_{(x,w)} \geq 0, \beta'_{(x,w)} \geq 0$ and
2. $\sum_{(x,w) \in I} \alpha'_{(x,w)} = \sum_{(x,w) \in J} \beta'_{(x,w)} = 1.$

By assumption one of the following occurs:

- (a) for all x and x' with (x, w) in I and (x', w) in J , $x > x'$, or
- (b) for all x and x' with (x, w) in I and (x', w) in J , $x < x'$.

If (a) occurs, then for x and x' with (x, w) in I and (x', w) in J , $Q_x(1) > Q_{x'}(1)$ and $\sum_{(x,w) \in I} \alpha'_{(x,w)} Q_x(1) > \sum_{(x,w) \in J} \beta'_{(x,w)} Q_{x'}(1)$. If (b) occurs, then for x and x' with (x, w) in I and (x', w) in J , $Q_x(1) < Q_{x'}(1)$ and $\sum_{(x,w) \in I} \alpha'_{(x,w)} Q_x(1) < \sum_{(x,w) \in J} \beta'_{(x,w)} Q_{x'}(1)$. In either event, we have reached a contradiction. Hence, $\sum_{(x,w) \in I} \alpha_{(x,w)} = 0$ holds for each w in $\{0, 1, 2\}$, $\sum_{w=0}^2 \sum_{(x,w) \in I} \alpha_{(x,w)} = \sum_{(x,y) \in I} \alpha_{(x,y)} = 0$, which is a contradiction, so (*) cannot occur.

Proof of Theorem 4. We construct $\{\tilde{U}, \tilde{S}, \tilde{c}, \tilde{p}\}$ as follows. Let $T = \{0, 1\}$ and thus $\tilde{S} = S \times T$. Define \tilde{c} by $\tilde{c}(a, (s, 0)) = c(a, s)$. Let $h: R \rightarrow R$ be a smooth, C^∞ increasing function with

$$h(x) = \begin{cases} x & \text{for } x \geq 0 \\ \frac{xM}{r\gamma\varepsilon} & \text{for } x \leq -r. \end{cases}$$

Let $\tilde{c}(a, s, 1)$ be a collection of distinct objects disjoint from the set $c(a, s)$. Define \tilde{p} by

$$\tilde{p}((s, t) | X = x) = \begin{cases} (1 - \varepsilon)p(s | X = x) & \text{if } t = 0 \\ \varepsilon p(s | X = x) & \text{if } t = 1. \end{cases}$$

Finally, let \tilde{U} be defined by

$$\tilde{U}(a, (s, t)) = \begin{cases} U(a, s) & \text{if } t = 0 \\ h \circ U(a, s) & \text{if } t = 1. \end{cases}$$

The utility function is continuous and smooth, and it preserves the relative preference ordering of the utility function $U(\cdot, \cdot)$.

It is elementary to verify that $\{\tilde{U}, \tilde{S}, \tilde{c}, \tilde{p}\}$ is an ε perturbation of $\{U, S, c, p\}$. Let x be in X with $p(\cdot | X = x) \geq \gamma q(\cdot)$ and let $a \in W(r)$. Then

$$\begin{aligned} & \sum_s \sum_t \tilde{U}(a, s, t) \tilde{p}((s, t) | X = x) \\ &= \sum_s U(a, s) (1 - \varepsilon) p(s | X = x) + \sum_s h \circ U(a, s) \varepsilon p(s | X = x) \\ &\leq (1 - \varepsilon)M + \varepsilon M + \varepsilon \sum_{s: U(a, s) < -\lambda} h \circ U(a, s) p(s | X = x) \end{aligned}$$

where $\lambda > 0$.

If a is in $W(r)$, then there is a $\lambda > 0$ such that $\lambda q\{s: U(a, s) < -\lambda\} > r$. Since $\lambda q\{s: U(a, s) < -\lambda\} \leq \lambda$, it follows that λ is greater than r , and that

$q\{s: U(a, s) < -\lambda\} > r/\lambda$. If

$$U(a, s) < -\lambda, \text{ then } h \circ U(a, s) < -\lambda M / (r\gamma\epsilon).$$

Therefore, we have

$$\sum_s \sum_t \tilde{U}(a, s, t) \tilde{p}((s, t) | X = x) < M + \epsilon\gamma h(-\lambda) q\{s: U(a, s) < -\lambda\}$$

and thus

$$\sum_s \sum_t \tilde{U}(a, (s, t)) \tilde{p}((s, t) | X = x) < M + \epsilon\gamma \frac{-\lambda M}{r\gamma\epsilon} \frac{r}{\lambda} = 0.$$

However,

$$\sum_s \sum_t \tilde{U}(a_0(s, t)) p((s, t) | X = x) \geq 0$$

since for each s , $U(a_{0,s}) \geq 0$ $s \circ h \circ U(a_{0,s}) \geq 0$. Thus $\tilde{d}(x)$ is not in $W(r)$.

Proof of Corollary to Theorem 4. We first present a lemma.

Lemma. Let F be a finite nonempty set and I an infinite set. If d is a function from I into F , then d is constant on some infinite subset of I .

Proof. I is the disjoint union of a finite collection of sets of the form $d^{-1}(f)$ where f is in F . Therefore for at least one f in F , $d^{-1}(f)$ is infinite.

Proof. The set

$$A_\alpha = \{a: q\{s: U(a, s) < -\alpha\} > \alpha\} \subset W(\alpha^2),$$

since if a is in A_α then

$$\alpha q\{s: U(a, s) < -\alpha\} > \gamma^2.$$

Since $A_\alpha \subset W(\alpha^2)$,

$$W(\alpha^2)^c \subset A_\alpha^c$$

and so $W(\alpha^2)^c$ is finite; furthermore, a_0 is in $W(\alpha^2)^c$, so $W(\alpha^2)^c$ is finite and nonempty. By assumption $\{x \in X: p(\cdot|X=x) \geq \gamma q(\cdot)\}$ is infinite, so by applying the Lemma and Theorem 5 we obtain the desired corollary. Here Λ^c denotes the complement of Λ .