

costs," and so forth, there is little or nothing new that their elegant mathematics can tell us. Except in special cases, it is very difficult to find and define these quantities. Indeed, no general methods exist for computing even those quantities that we already know exist. Freeman and McFarland (1982) earlier presented and discussed a similar framework. I find the Freeman-McFarland approach much easier to comprehend and use. I believe others feel the same way. Maybe this is because they use a much more standard and comprehensible life history approach.

There are four important features of any unified foraging theory: (1) the often huge *differences in time scales* on which the various processes entering the model occur, (2) the *currency* to be optimized on the various time scales being considered, (3) the *intrinsic state* (e.g., physiological condition and age) of the organism, (4) the *environmental setting* in which the organism finds itself (the extrinsic state; e.g., food abundance, diversity of food types available, presence or absence of competitors and predators). I will briefly discuss these aspects of unified foraging theory.

H & M, referring to their earlier paper (McNamara & Houston 1986), discuss all these features except the different time scales entering the unified foraging theory. This is unfortunate, because day-to-day behavioral decisions are, on the longer time scale of life history optimization, experienced as an instant in time. Hence, we are allowed to assume that a wide variety of environmental features varying over the longer time scale may be treated as if they were constant on the shorter time scale (see, e.g., Stenseth & Maynard Smith, 1984, and Stenseth, 1986, for a discussion of this issue). Elsewhere (Engen & Stenseth, submitted) used this fact in developing a similar unified foraging theory.

H & M explain that the *common currency* to be used for comparing the danger due to predation and gain of energy is the "canonical cost." What is this? According to the authors, it is a quality that measures "the relationship between a momentary decision and the lifetime reproductive success of the animal. It is a direct *measure* of the contribution of an action to the animal's fitness." But why call it a cost when it is said to be a contribution to fitness? In order to understand that, we have to look closely at the mathematical definition of this cost. The optimal behavior always has a zero cost, whereas all other behaviors always have a positive cost. Although this is technically correct, it is confusing and makes H & M's paper difficult to read. One almost gets the feeling that they have twisted the biology under study to fit the requirements for applying a particular tool (dynamic programming). I have discussed this issue elsewhere (Engen & Stenseth, in press).

Another problem is that in order to be able to measure the cost discussed by H & M, we have to be able to measure the fitness. So why not stick to fitness as the common currency? The costs discussed by H & M are only technical ornaments. In evolutionary ecology, fitness or some substitute for fitness is the common currency. In traditional optimal foraging models, the amount of energy (or of other nutritional components) acquired per unit time in the long run is used as the currency because this is assumed (often correctly) to be correlated with fitness.

H & M, as well as Mangel & Clark (1986), demonstrate by their unified foraging theory that a foraging animal should accept a higher level of predation (or other risk factors) when its reserves are low. Engen and Stenseth (submitted), have provided a similar result showing that an old individual should accept a higher risk level than a young individual. Using their approach, the reason for the predicted pattern is very clear: An older individual has far less to lose in residual reproductive value (i.e., the "remaining" expected fitness) than a younger individual; therefore, he ought to choose the more risky (but presumably higher quality) food item. It follows that the optimal diet will depend on the intrinsic state (such as age) of the foraging animal in a predictable manner. Such a result could not

have been derived without a unified theory of optimal foraging that ties together a theory for making optimal behavioral decisions on a short time scale and a theory for optimal life history evolution related to a much longer time scale. This is an example of how optimal decisions depend on the organism's intrinsic state (age). Unfortunately, H & M did not provide similar results indicating how optimal behavior varies with age. This seems strange to me since in a fitness and life history context, age seems the most natural quantity to focus on. I presume this is due to the level of abstraction of their model. At the present stage of development of the unified theory, I prefer specific (and testable) examples instead of very general and abstract ones.

Predation is only one aspect of the optimal forager's environment. The presence of competing foragers of the same species or of other species (whose abilities to detect and to reach a food item may differ), represent another aspect. Predictions regarding the optimal decisions by one particular forager living together with other (nonevolving) individuals of the same and other species can be derived (e.g., Engen et al., in press). What has not yet been done, is to investigate how a group of individuals or a guild of species may coevolve their foraging behaviors. To pursue the development of a coevolutionary foraging theory represents one of the greatest challenges in the realm of optimal foraging theory. The work that Houston & McNamara are currently doing to tie together the branch of behavioral ecology discussed in their BBS target article and game theory is a very valuable contribution in that direction.

#### ACKNOWLEDGMENTS

My thinking on optimal behaviors such as foraging has been greatly influenced by Steinar Engen through our joint work on this topic. The ideas presented in this commentary were particularly influenced by long discussions with Brent Danielson. I thank them both.

### Constructing optimal sequences of behavior: Backwards is beautiful, but . . .

William Timberlake

Psychology Department, Indiana University, Bloomington, Ind. 47405

For several reasons I am in sympathy with the approach outlined by Houston & McNamara (H & M). I think that defining optimal sequences of behavior by working backward from a particular outcome clearly related to survival and reproduction is a powerful approach. Specifying possible outcomes at each choice point in terms of probabilities has good potential for fitting the variation in real circumstances and behavior.

I found it particularly important that the proposed procedures force the modeler to consider carefully the many aspects and levels of organismic functioning necessary for a particular instance of survival and/or reproduction. As a result, the technique does not lend itself to general abstract statements about optimal behavior; rather, optimal behavior is defined in terms of a concrete sequence of particular behaviors and organismic measures. Because function is defined in such concrete terms, there is marked potential for developing much closer ties between function and mechanism. This should encourage more careful analysis of phenomena and the inductive development of general theories from well-worked out particulars, as opposed to the tendency to try to subsume as many phenomena as possible under a single mathematical equation.

However, this concreteness, which is such a strength, is also a potential weakness. The value of a particular application will depend on researchers' knowledge and their ability to analyze the potential determinants of the outcome of interest. Dynamic programming is basically a scheduling device that must produce an optimal schedule but only for the events and information

behavioral ecology, what opportunity costs might we pay? For one, we often find that models that gain Realism by adding more parameters or interactions between the parameters (dynamic programming models are nifty; they do things one can't do otherwise (explicitly incorporating the effects of the organism's change of state on its subsequent behavior, as H & M indicate), but they are computer-dependent and don't have the Generality (in my opinion) of an analytical solution. The *method* may be general enough, and if we are lucky the insights from a particular application may have wide applicability, but as the models gain realism they tend to apply to smaller neighborhoods.

The loss of Testability in more Realistic models is, I think, a more serious threat in the case of something like dynamic programming. First, making the optimal solution highly sensitive to an animal's state multiplies the variables that need to be measured and divides the sample sizes available from a given quantum of field (or lab) work. Second, the fact that dynamic optimization models predict a *scatter* of outcomes may seem like a virtue ("Hey - what looks like noise out there may be the music of selection and optimal decision-making!"), but this blade is two-sided, for it becomes harder to judge what a clean falsification of a given prediction might look like. (This may not be much of a problem for those working in the lab - pigeons in Skinner boxes or great tits feeding from conveyor belts - but for anyone working with creatures (including *Homo*) in their natural habitats I think it's a real drawback.) Third, the specific approach taken by H & M depends on defining a "canonical" (fitness opportunity) cost to each course of action. In theory, this is certainly the most General and correct way to proceed, but in practice it's a mighty hard row to hoe.

Let me expand on this last point. Except in limited instances, it is notoriously difficult if not downright impossible to measure the fitness gains and costs of alternative (short-term) behavioral decisions. If we choose our problems carefully (birds foraging in winter - let's see, that eliminates the need to deal with defending territories, attracting mates, provisioning young, most predation . . .), we might make some plausible approximations to the canonical cost of selected activities; but is this being Realistic? And if we can't really measure these costs, such that the common currency promised by the canonical cost dissolves into a set of proxies (predation risk, energy reserves, mating opportunities . . .), where is the virtue in our greater Realism? Perhaps I exaggerate the limitations of the H & M approach; after all, dynamic optimization is still in its infancy in behavioral ecology. But I am skeptical about a method that promises far more than it seems likely to achieve. For example, with respect to the analysis of the foraging/predation trade-off, we would already seem to know - through intuition if not through static optimization models - that declining energy reserves should lead to decreased vigilance. Without operational means of measuring fitness returns from specific alternative mixes of feeding and vigilance, however, what do we really learn from a formal model, and how would we ever test it?

Let me be clear. I am not claiming that greater complexity in the service of greater Realism is always a bad thing, or that dynamic optimization models in particular are inferior to static models. But I am saying that, unlike the message H & M seek to convey, dynamic optimization models are not necessarily superior either. To respond to an example H & M raise: It may be that static rate maximization is "inadequate" for Realistic modeling of the role of risk in foraging decisions. Yet it is possible to capture many of the most important (General) and empirically Testable aspects of risk-sensitive foraging without recourse to dynamic models (Stephens & Krebs 1986: Chapter 6). Furthermore, the real world turns out to be structured such that maximizing the rate of energy gain may often lead to minimizing the risk of shortfall, as shown analytically by Stephens and

Charnov (1982) and through simulation by Winterhalder (1986). If the organism in question is able to take advantage of sharing networks, it is surprising how easily risk can be buffered (Win- uines that complexity actually facilitates some useful simplifications.

To conclude: In science, as in actual decisions, we always have to simplify reality. The question, which has no easy answer, is what to simplify and how much. The choice we make, and its validity, depends crucially on the goals of the research and our available information. In this respect, behavioral ecology still has an awful lot to learn about its subject, information that can best be gained from judicious (but not exclusive) application of simple, static, less Realistic but more Testable and/or General models.

## Houston & McNamara are right, but are they helpful to empiricists?

Nils Chr. Stenseth

Division of Zoology, Department of Biology, University of Oslo, Blindern N-0316 Oslo 3, Norway

The target article by Houston & McNamara (H & M), discussing their version of a unified foraging theory (Mangel & Clark 1986), is not very helpful to empiricists because it does not even vaguely relate to real-world components of fitness (e.g., survival and reproductive schedules) and other quantities that people in the field (especially those in the field!) can relate to; H & M's other relevant papers provide little help. Unfortunately, very few specific interpretations of real-life situations are provided in any of their papers. When they do provide specific examples, the model reduces to rather special cases that do not need the more sophisticated general framework. In addition, most of their examples are only vaguely related to real-life cases. In fact, it is very easy for the reader (and the authors, I believe) to lose sight of the biological situation under study. This is unfortunate because theorists working in the field of the natural sciences need to communicate with empiricists if they are to have some impact; we theorists must always remember that it is the empiricists who have the data against which to compare our theoretical results. There is no question, of course, but that H & M are right. It is unfortunate, however, that people will fail to appreciate it and to value their contribution as much as they ought to. It would have been so much easier if H & M had presented their material as clearly in written form as they do in verbal form.

If progress is to be made, we modelers are faced with a challenge: We must explain our models in a manner that is comprehensible to those in the field and laboratory; technical material should go to appendices leaving the main text written in plain English (or Norwegian). We must also provide examples demonstrating that the new approach actually provides new insights.

H & M tell us that a behavior is optimal *only* if it increases the animal's ultimate fitness more than any other decision. But this we knew. What we need is a way to determine how a behavior made at one instant in time affects the total fitness of the organism. Obviously, this will be different for each type of behavior (e.g., when to breed, how often to breed, what to eat, how much predation risk is worth risking). *When* such behaviors occur in an individual's life may also be important (e.g., before or after having bred once, when young or old, before or after growing to a given body size, and so forth). These needs can be met only through a *mechanistic* model (Schoener 1986) that will show us how observable behaviors relate to life history optimizing. H & M purport to do this; in fact, they do not. Because they resort to nebulous "terminal reward functions," "canonical