

# Suboptimal foraging theory: How inaccurate predictions and approximations can make better models of adaptive behavior

Steven F. Railsback<sup>1,2</sup>

<sup>1</sup>*Lang Railsback & Associates and Department of Mathematics, California State Polytechnic University, Humboldt, Arcata, California 95521 USA*

<sup>2</sup>E-mail: [steven.railsback@humboldt.edu](mailto:steven.railsback@humboldt.edu)

*Abstract.* Optimal foraging theory (OFT) is based on the ecological concept that organisms select behaviors that convey future fitness, and on the mathematical concept of optimization: finding the alternative that provides the best value of a fitness measure. As implemented in, e.g., state-based dynamic modeling, OFT is powerful for one key problem of modern ecology: modeling behavior as a tradeoff among competing fitness elements such as growth, risk avoidance, and reproductive output. However, OFT is not useful for other modern problems such as representing feedbacks within systems of interacting, unique individuals: when we need to model foraging by each of many individuals that interact competitively or synergistically, optimization is impractical or impossible—there are no optimal behaviors. For such problems we can, however, still use the concept of future fitness to model behavior, by replacing optimization with less precise (but perhaps more realistic) techniques for ranking alternatives. Instead of simplifying the systems we model until we can find “optimal” behavior, we can use theory based on inaccurate predictions, coarse approximations, and updating to produce *good* behavior in more complex and realistic contexts. This “state- and prediction-based theory” (SPT) can, for example, produce realistic foraging decisions by each of many unique, interacting individuals when growth rates and predation risks vary over space and time. Because SPT lets us address more natural complexity and more realistic problems, it is more easily tested against more kinds of observation and more useful in management ecology. A simple foraging model illustrates how SPT readily accommodates complexities that make optimization intractable. Other models use SPT to represent contingent decisions (whether to feed or hide, in what patch) that are tradeoffs between growth and predation risk, when both growth and risk vary among hundreds of patches, vary unpredictably over time, depend on characteristics of the individuals, are subject to feedbacks from competition, and change over the daily light cycle. Modern ecology demands theory for tradeoff

behaviors in complex contexts that produce feedbacks; when optimization is infeasible, we should not be afraid to use approximate fitness-seeking methods instead.

*Key words: adaptive behavior, behaviorally mediated ecology, feedbacks, foraging theory, individual-based models, optimization, state- and prediction-based theory, tradeoff decisions*

## INTRODUCTION

Optimal foraging theory (OFT) has been a foundation of mathematical ecology for over a half century, but its relevance seems to decrease as ecologists tackle problems of increasingly realistic complexity. OFT uses two concepts to model individual decision-making: the ecological concept that, due to evolution, behaviors act to convey individual fitness; and the mathematical concept of optimization, which represents decision-making as selecting alternatives that maximize an objective function. The objective function (“fitness measure”) is the link between the two concepts: it is a measure of specific elements of fitness, e.g., growth rate, probability of survival to a future time, expected number of offspring at a future date, expressed mathematically so it can be optimized. (These concepts are also used in optimization theory for other adaptive traits such as life history decisions. In fact, almost everything I say here about OFT also applies to such other theory.) When we consider food intake the sole element of fitness, OFT represents how individuals maximize intake rate (e.g., MacArthur and Pianka 1966; Charnov 1976). For decisions that involve tradeoffs among competing fitness elements, state-based fitness maximization theory (often referred to as state-based dynamic modeling, SDM; Houston et al. 1988; Houston and McNamara 1999; Clark and Mangel 2000) can determine the optimal pathway through a set of behavior alternatives to maximize a fitness measure—such as expected future survival or reproductive output—that represents the combined effects of food intake, predation risk, etc.

While OFT has been extremely productive as a framework for individual decisions, it has inherent limitations as a framework for modern ecology. Agrawal et al. (2007) reported the conclusions of a panel convened to identify research priorities for population and community ecology; those priorities included freeing ecology of three traditional assumptions: (1) that effects of multiple factors (e.g., competition and predation) are independent, (2) that traits are uniform and unchanging, and (3) that feedbacks due to interactions can be ignored. The first of these assumptions has been addressed extensively over the past several decades, with OFT—specifically, SDM—as an essential tool. The theoretical and empirical literature on “indirect effects” and “trait-mediated indirect interactions” (e.g., Abrams 1993; Preisser et al. 2005; Werner and Peacor 2003) has established the importance of considering tradeoffs among factors like growth and risk in understanding foraging decisions and their effects on ecology. SDM has been extremely valuable for understanding how such tradeoffs drive individual behavior. The essential characteristic of SDM that makes it powerful for modeling tradeoffs

is its evaluation of expected fitness over an extended future period: when individuals evaluate fitness over a long time, they can make good tradeoffs between food intake (to avoid future starvation and to increase reproductive output) and predation risks (which accumulate over time in a nonlinear way: a small daily risk becomes a low probability of extended survival).

However, OFT is unfortunately dependent on the second and third of the problematic assumptions identified by Agrawal et al. (2007). Addressing trait variability and feedbacks requires modeling systems of diverse, interacting individuals, and OFT cannot represent the decisions of individuals in such systems. Why? Because optimization over a future time period is impractical or impossible when the alternatives available in the future, and their payoffs and risks, are subject to feedbacks from the behaviors of all the other unique individuals. We can model systems of diverse, interacting individuals via individual-based models (IBMs) that represent population or ecosystem ecology as emerging from the behavior, interaction (via competition, etc.), and fates of individuals. In such IBMs, the ecological concept of OFT—that individuals act to increase their fitness over a future time period—remains valid, but the mathematical concept—optimization—is no longer useful. When our models include such realistic complexities as competition among individuals, multiple life stages and trait variability, and variable and dynamic environments, there is no optimal solution: the consequences of any choice depend on the choices and fates of the other individuals and on unknown future environmental conditions. (I illustrate this problem below.) Optimization is feasible only when we simplify away such complexities, but modern ecologists need to tackle the complexities head-on instead of avoiding them.

The inability of traditional OFT to address such realistic complexities has unfortunate consequences. When we use IBMs to address theoretical or management issues without ignoring individual variability, feedbacks, etc., the unusefulness of OFT leads to a general presumption that theory is not relevant to models of useful complexity. My experience with individual-based modelers indicates that this presumption is widespread: when designing models with individual adaptive decision-making, modelers rarely turn to ecological theory and, instead, typically use ad-hoc approaches that have their own limitations. In particular, without the ability to use OFT very few IBMs have successfully modeled tradeoff behaviors, despite their ubiquity and importance.

The question I address is how we can still use OFT's ecological concept that individuals act to increase future fitness when we cannot use its mathematical concept of optimization. The answer is to model like engineers instead of mathematicians. Finding useful solutions to complex problems is the essence of engineering, and engineers use a variety of tools to do so. Instead of simplifying the modeling problems until we can find precise optimal solutions, we can simplify and approximate the solution methods to find *good* solutions to realistically complex problems. Instead of our theory being driven by mathematical concerns, it can be driven by the ecological problems we need to solve.

## THEORY FOR GOOD DECISIONS WHEN OPTIMIZATION IS INFEASIBLE

The traditional optimization approaches to theory have certainly produced many valuable insights and useful models of real systems and problems. But when we need behavior theory useful for managing and understanding real populations and ecosystems, traditional theory is frustrating: it can produce precise answers, but often only to highly approximated and simplified problems.

How can we do theory the other way around: produce approximate answers for more realistic systems and problems? Can we produce theory more useful for complex problems if we seek good instead of optimal behavior—“fitness seeking,” when fitness optimization is impossible? Engineering decision theory provides guidance. Like OFT, engineering decision theory uses a specific mathematical objective function (e.g., minimizing life cycle costs). However, engineers often use simplistic predictions and other approximations to find good solutions for problems that would be impossible to optimize, and, when possible, they update predictions and decisions as new information becomes available. Can theory for foraging and other traits become both more realistic and less constrained if it explicitly represents how organisms base decisions on predictions of limited accuracy and then update those predictions and decisions as they sense new information?

One approach engineers use to produce good decisions in complex situations is machine learning, a diverse and rapidly evolving family of methods for generating decision algorithms. Methods range from simple neural networks parameterized via genetic algorithms to much more capable methods such as reinforcement learning (Sutton and Barto 2018). These methods work by defining a mathematical decision structure (a neural network; a transition probability function for Markov processes) that is then “trained” by repeatedly exposing it to a set of challenges and using feedback on its success to modify the decision structure. Simple machine learning methods have long been used to model adaptive tradeoff decisions in IBMs (e.g., Strand et al. 2002). Potential applications of reinforcement learning to behavioral ecology are explored by Frankenuis et al. (2019), who describe it as a computational technique for finding good solutions to SDM-like problems too complex for formal optimization. However, machine learning has well-known limitations. First, it requires computer science expertise and effort to develop the decision structure, the training environment, and the method for evaluating success and using feedback to modify the decision structure. Second, these methods are not theoretical in the sense that OFT is: by incorporating the fundamental assumption that behavior acts convey future fitness. Instead, artificially learned behavior only solves the exact problem it was exposed to in training, and we cannot safely assume that such behavior will convey fitness under any other conditions. This second limitation seems critical because the main reason we produce ecological models is to predict responses to novel conditions. And, when we are modeling systems with feedbacks, any change in the population (e.g., the random death of a dominant individual) presents the remaining individuals with a novel situation.

Railsback and Harvey (2020) explore another approach that is more closely based on traditional OFT. “State- and prediction-based theory” (SPT) can be defined as a general method for modeling adaptive decisions that (like SDM)

assumes individuals act to maximize a mathematical measure of future fitness, when (unlike SDM) that fitness measure includes predictions and approximations allowing it to be evaluated even when future conditions are unknown. While SPT is directly based on SDM, it differs in two key ways: SPT accommodates unforeseeable future conditions by assuming individuals update decisions over time, and it uses prediction and approximation to make good decisions when optimization is infeasible. Like SDM, SPT assumes an individual selects the alternative that provides the highest value of a mathematical measure that represents fitness elements such as future survival of predation and starvation. However, SPT does not try to find an “optimal” set of alternatives to use until the time horizon—it does not try to solve behavior over a whole life stage or extended period at once, which requires assuming that future environmental conditions and interactions are fixed and known. Instead, we use SPT to model how individuals update decisions over time as conditions change. Using SPT, an individual evaluates its fitness measure using: (1) explicit predictions of future conditions that are typically simplistic and inaccurate, yet useful; and (2) approximations such as the incorrect assumption that the individual will use one alternative from the present until the time horizon. In other words, SPT represents how an organism decides what to do right now by approximating what might happen to it in the future as a consequence. Repeating the decision over time lets individuals adapt to changes in their environment and in their own state, and is what makes SPT useful when optimization is impossible: individuals make a series of adaptive decisions as their world changes.

In a typical application of SPT, a model organism evaluates each decision alternative (e.g., each patch an animal could forage in) by (a) using the extremely simplistic prediction that growth and risk conditions currently occurring would remain unchanged until the time horizon, (b) assuming incorrectly that it would use the same alternative until the time horizon, and (c) estimating its fitness measure at the time horizon using approximations such as neglecting how its growth and risk depend on size as it grows. (Railsback and Harvey 2020 provide detailed guidance.) The model organism then chooses the alternative offering the highest value of this approximated fitness measure—and repeats the whole decision process every time step as conditions change due to weather, competition with other individuals, its growth, a changing perception of predation risk, etc.

Compared to both SDM and machine learning, SPT has the advantages of simplicity and computational efficiency: its use in an IBM requires neither iterative optimization techniques like dynamic programming nor any of the complex procedures needed to train and use an artificial behavior model. The example applications below illustrate these advantages.

#### HOW APPROXIMATE APPROACHES CAN BE BETTER

Railsback and Harvey (2020) show that SPT can produce behaviors surprisingly close to optimal in systems simple enough to also use optimization. But the title of this article indicates that this kind of approximate theory can produce “better” models... how can SPT be better than optimization?

The most fundamental reason that approximate decision theory such as SPT can produce better models is that it allows us to unsimplify our representation of

the individuals and their environment: without the need to exactly solve an optimization, we can make the individuals and their habitat more diverse and realistic and still predict behavior and its higher-level consequences. One example identified by Railsback and Harvey (2020) is how starvation and its effects on behavior are represented. Optimization approaches such as SDM often assume that individuals starve if, and only if, their energy reserves fall below a threshold, an unrealistic approximation that greatly simplifies the optimization. But this approximation keeps SDM from reproducing important dynamics such as “top-down” trophic effects: it does not allow individuals the scope to further reduce feeding when predation risk is particularly high (Railsback and Harvey 2013). With SPT, it is easy to use a more realistic assumption that starvation risk increases gradually as energy reserves decline. This unsimplification allows the individuals to hide, and sometimes survive, when predation risk is very high—a behavior that can have major implications to the dynamics (e.g., extinction risk) of at-risk populations. Examples presented below further illustrate how SPT makes it easy to include complexities that would make optimization (or game theory) intractable.

Behavior theory based on approximation and updating can also be better because it allows us to address more realistic and important problems. Perhaps most importantly, it gives us theory useful in models that include feedbacks, especially individual-based population models that represent competition as well as unpredictable habitat dynamics that also make optimization infeasible. Behavior theory useful in IBMs can make IBMs easier to build and less ad hoc, and more likely to include the tradeoff behaviors and feedbacks that we now know are essential to ecology.

Theory that we can implement in complex IBMs is better in a third way: it is more testable. “Testing” conventional theory that assumes a highly simplified ecological system in order to make its mathematics tractable presents a dilemma: what is a “valid” empirical test of theory that optimizes a highly simplified depiction of ecology? Do we use a highly simplified test system that represents the conditions assumed by the theory but not nature, or more natural conditions that violate the theory’s assumptions? Theory that can be implemented in an IBM that contains key complexities of a real system offers not only a way around this dilemma but a far more robust testing approach. Testing theory in the virtual ecosystem provided by an IBM can be more robust by providing an environment that is more realistic than simplified laboratory conditions yet still fully controlled, by facilitating the formulation and contrast of alternative theories in a hypothesis-testing cycle (as illustrated by Railsback and Harvey 2020), and by allowing the theory to be tested against not just a single experiment but against a wide variety of empirical observations at different scales and ecological levels (Grimm et al. 2005; Lorscheid et al. 2019). A number of such theory tests have now been published (e.g., Amano et al. 2006; Railsback and Johnson 2011; Cortés-Avizanda et al. 2014).

Finally, theory such as SPT that assumes organisms have less-than-optimal capabilities is likely a more realistic representation of how organisms actually make behavioral decisions. This kind of theory can also accommodate realistic mechanisms such as limitations on the range or accuracy of sensing environmental

conditions and learning from experience (illustrated below). Neuroscientists, plant physiologists, and others are making rapid progress understanding how organisms actually make adaptive decisions; incorporating this new knowledge into ecology will require bottom-up theory that allows detailed representation of the complex mechanisms individuals actually use.

#### EXAMPLES: THEORY PRODUCING *GOOD* BEHAVIOR IN COMPLEX CONTEXTS

Railsback and Harvey (2020) provide examples of SPT-based models over a wide range of complexity. The simplest is a direct adaptation of an early illustration of SDM for optimal foraging under predation risk: the forager patch selection model of Mangel and Clark (1986). In this model, simple foragers select among five locations (patches 1-5) that respectively have daily probabilities of finding food and surviving predation of 0.0 and 1.0 (a refuge), 0.2 and 0.95, 0.35 and 0.85, 0.35 and 0.75 (providing lower survival and no more food than patch 3), and 0.5 and 0.6. Mangel and Clark (1986) used dynamic programming optimization to identify the best sequence of patches to use over a 10-day life stage. However, this optimization was possible only because the model neglects feedbacks such as competition for food or escape cover. Once we assume that the behavior of each forager affects the resources available to the others, the decision becomes more complex and there is no clear optimal solution. Techniques such as stochastic dynamic programming can determine the overall average best series of patches to occupy, but cannot determine what each individual should do as a consequence of how lucky it has been at finding food and what the other foragers chose to do.

The “suboptimal” version of the patch selection model uses the typical assumptions of SPT listed above. Foragers make the prediction that future feeding success and survival probabilities will be the same as they are at the time they make the decision (which is accurate for some versions of the model but not those with feedbacks). Dynamic programming is avoided by foragers using the incorrect approximation that if they select a patch they will stay in it until the time horizon at day 10, but then updating their decision each day considering their current energy reserves and the current state of the patches. Railsback and Harvey (2020) found that this approximate approach produces results—patch selection decisions and survival rates—very close to the optimal solution of Mangel and Clark (Figure 1, panel A). Foragers start in patch 3 and, unless unlucky at catching food, move to the safer patches.

The SPT approach lets us easily add realistic feedbacks and complexities to the model with no change in foraging theory. Food competition and its feedbacks can be represented by assuming that each time a forager captures food the probability of feeding success for subsequent occupants of its patch decreases by 1%. Food competition forces more use of riskier patches and some foragers now use patch 4 (Figure 1, panel B). “Predator swamping” can be represented by assuming risk decreases with the number of foragers in a patch, so survival probability is a feedback of other foragers’ behavior. I simply assumed that each forager evaluates its survival probability in a patch as  $1.0 - \{(20 - 20S)/(20 + n)\}$  where  $S$  is the standard patch survival probability described above and  $n$  is the number of foragers currently in the patch. (The standard survival probability is equivalent to a risk of  $20(1 - S)/20$  foragers being

killed; each forager then increases the denominator, so with 5 foragers in a patch the risk is  $20(1 - S)/25$ .) With this feedback, foragers concentrate in patch 3 instead of using patch 2 (Figure 1 C).

Learning can also easily be represented: I modified the model so foragers can move at most one patch per time step, and learn the predation risk of their current and adjacent patches from experience. Foragers initially assume all patches have a survival probability of 0.9 but then each forager adjusts its estimates of survival probability in its current patch and the two adjacent patches each time another forager gets killed or survives in those patches. This learning process uses Bayesian updating (e.g., Sect. 12.5 of Railsback and Grimm 2019). Patch selection is therefore a function of perceived risk, which is now a feedback of decisions by other foragers. Early in the simulation, foragers disperse more among patches until they learn accurate estimates of survival probability (Figure 1 D). The early use of patch 1 is an artifact of forager movement being limited to one patch at a time: foragers initialized in patch 1 have zero fitness in patches 1 and 2 (neither patch offers sufficient food), so they choose between them randomly.

The original, most complex, and most thoroughly tested application of SPT is for habitat selection theory in our family of stream salmonid IBMs (Railsback and Harvey 2002; Railsback et al. 2021). These models were designed to support river management decisions such as how much flow to release from dams and how to design habitat restoration projects. In them, simulated fish each select among dozens to hundreds of potential foraging patches that differ in both growth rate and survival probability. Growth and survival also depend on the state of each individual (its length and energy reserves), and vary daily as stream flow, temperature, and turbidity change. In the newest version (Railsback et al. 2020), fish adaptively decide whether and where to feed or hide over the circadian light cycle, with growth potential and predation risk varying among dawn, day, dusk, and night as well as over space. The habitat selection theory uses the same simplifying assumptions as the simple forager model does, yet produces realistically good results at multiple scales. Simulation experiments showed that the theory could reproduce a variety of observed habitat selection responses to drivers such as competition, temperature (which drives metabolic demands and hence starvation risk), and predation risk (Railsback and Harvey 2002; Railsback et al. 2005). Further experiments showed that IBMs using the theory could reproduce a variety of population-level observations (Railsback et al. 2002) and trophic interactions such as trait-mediated indirect interactions that are driven by risk-growth tradeoffs (Railsback and Harvey 2013). Many of the realistic dynamics produced in these experiments are driven by feedbacks that are impossible or impractical to model with optimization theory.



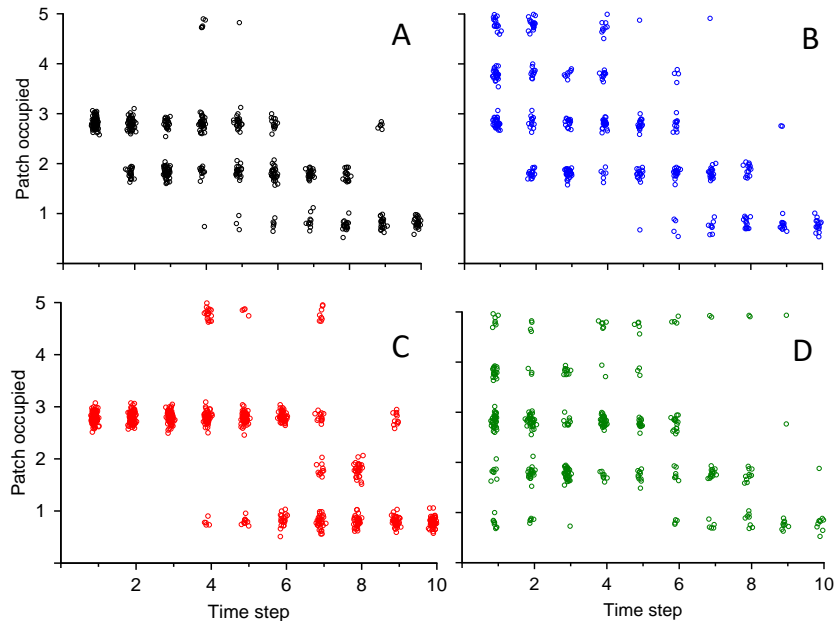


Figure 1. Results of four versions of the SPT patch selection model, all using exactly the same decision theory. Each cloud of dots depicts the number of foragers selecting a patch on the given time step, out of an initial 100 foragers. (Many foragers die each time step.) A: Original version based on the patch selection example of Mangel and Clark (1986). B: With food competition. C: With “predator swamping”: risk decreases with the number of foragers in a patch. D: With limited selection distance and learning of survival probability from predation events.

## CONCLUSIONS

Modelers, psychologists, and computer scientists have developed many ways to represent decision making in complex situations. These range from simple heuristics (e.g., Gigerenzer and Todd 1999), to artificially evolved mathematical structures that reference organisms’ sensing, emotional, and cognitive systems (e.g., Giske et al. 2003, 2013; Budaev et al. 2019), to powerful learning algorithms such as reinforcement learning (e.g., Sutton and Barto 2018). Ecological modelers need to be aware of such approaches and consider their potential usefulness, a topic too broad to address here. Instead, I focus only on methods like SPT that descend directly from, and retain the conceptual basis of, the fitness optimization approaches that are the foundation of theoretical ecology.

My goal here is to expound three concepts. First is simply that we need theory for adaptive tradeoff decisions, such as foraging under risk, that is useful in contexts less simplified than those typically required by traditional, mathematics-focused theory. Especially, we need theory for tradeoff decisions that is useful in IBMs designed to address management problems of real systems. Such contexts typically include complexities such as individual variation, sometimes hundreds of alternatives that individuals choose among, temporal change in both individuals and habitat, and—especially—feedbacks of behavior. Theory is essential for making models with such complexities efficient and coherent (Lorscheid et al. 2019); ad hoc approaches have not not been productive.

The second concept is that ecological theory need not always depend on formal, rigid mathematical frameworks such as optimization and game theory. When ecologists address problems too complex for such frameworks, we should not be afraid to approximate the decision algorithm as well as approximating the model system. Instead of simplifying the system until we can use tidy mathematics, we can use messier mathematics—including approximation and simulation—so we can represent more of nature’s messiness. But messier mathematics can still be theory: general, reusable approaches and algorithms that retain the basic assumptions of OFT while producing useful and testable models of real problems.

Third is that as we learn more about the neurological, biochemical, and other mechanisms used by real organisms to make adaptive decisions, we will need ecological theory that can represent those mechanisms. While the approaches illustrated here are still extreme simplifications, there is evidence that prediction, approximation, and updating (perhaps with learning) are useful not only as modeling abstractions but also as representations of what some organisms actually do (explored, e.g., by Glimcher and Fehr 2013).

The main question raised by these concepts is: What new kinds of theory might be useful, and how do we discover it? What specific predictions and approximations are useful for modeling fitness-seeking behavior in IBMs and elsewhere? Railsback and Harvey (2020) also address these questions. Our experience so far is that the answers can come readily when we focus on specific systems and specific problems—such as how particular populations and communities respond to specific perturbations such as habitat alteration, species introductions, or climate change. When we have a specific system and problem, we can then identify key behaviors and the factors that drive them. We can also identify (from literature or new field studies) observations that guide the development and testing of alternative theories. The cycle of posing alternative theories for individual behavior, falsifying those that do not cause an IBM to reproduce a variety of observations at individual and higher levels, and then refining both theory and observation and starting the cycle over, has already been established as a productive way to find tested, reusable theory for adaptive behavior (Grimm et al. 2005; Grimm and Railsback 2005; Lorscheid et al. 2019). As ecologists apply these methods to specific problems and systems, we are likely to find generalizable categories of approximations and assumptions that are useful for specific classes of behavior that cannot be modeled with optimization theory.

#### LITERATURE CITED

- Abrams, P. A. 1993. Why predation rate should not be proportional to predator density. *Ecology* 74:726-733.
- Agrawal, A. A., D. D. Ackerly, F. Adler, E. Arnold, C. Cáceres, D. F. Doak, E. Post, P. J. Hudson, J. Maron, K. A. Mooney, M. Power, D. Schemske, J. Stachowicz, S. Strauss, M. G. Turner, and E. Werner. 2007. Filling key gaps in population and community ecology. *Frontiers in Ecology and the Environment* 5:145–152.
- Amano, T., K. Ushiyama, S. Moriguchi, G. Fujita, and H. Higuchi. 2006. Decision-making in group foragers with incomplete information: test of individual-based model in geese. *Ecological Monographs* 76:601–616.

- Budaev, S., C. Jørgensen, M. Mangel, S. Eliassen, and J. Giske. 2019. Decision-making from the animal perspective: bridging ecology and subjective cognition. *Frontiers in Ecology and Evolution* 7:164.
- Charnov, E. L. 1976. Optimal foraging, the marginal value theorem. *Theoretical Population Biology* 9:129–136.
- Clark, C. W., and M. Mangel. 2000. *Dynamic state variable models in ecology*. Oxford University Press, New York.
- Cortés-Avizanda, A., R. Jovani, J. A. Donázar, and V. Grimm. 2014. Bird sky networks: how do avian scavengers use social information to find carrion? *Ecology* 95:1799–1808.
- Frankenhuis, W. E., K. Panchanathan, and A. G. Barto. 2019. Enriching behavioral ecology with reinforcement learning methods. *Behavioural Processes* 161:94-100.
- Gigerenzer, G., P. M. Todd, and the A. R. Group. 1999. *Simple heuristics that make us smart*. Oxford University Press, New York.
- Giske, J., S. Eliassen, Ø. Fiksen, P. J. Jakobsen, D. L. Aksnes, C. Jørgensen, and M. Mangel. 2013. Effects of the emotion system on adaptive behavior. *The American Naturalist* 182:689-703.
- Giske, J., M. Mangel, P. Jakobsen, G. Huse, C. Wilcox, and E. Strand. 2003. Explicit trade-off rules in proximate adaptive agents. *Evolutionary Ecology Research* 5:835-865.
- Glimcher, P. W., and E. Fehr (eds.). 2013. *Neuroeconomics: Decision making and the brain*. Academic Press.
- Grimm, V., and S. F. Railsback. 2005. *Individual-based modeling and ecology*. Princeton University Press, Princeton, New Jersey.
- Grimm, V., E. Revilla, U. Berger, F. Jeltsch, W. M. Mooij, S. F. Railsback, H.-H. Thulke, J. Weiner, T. Wiegand, and D. L. DeAngelis. 2005. Pattern-oriented modeling of agent-based complex systems: lessons from ecology. *Science* 310:987–991.
- Houston, A., C. Clark, J. McNamara, and M. Mangel. 1988. Dynamic models in behavioural and evolutionary ecology. *Nature* 332:29–34.
- Houston, A. I., and J. M. McNamara. 1999. *Models of adaptive behaviour: an approach based on state*. Cambridge University Press, Cambridge.
- Lorscheid, I., U. Berger, V. Grimm, and M. Meyer. 2019. From cases to general principles: A call for theory development through agent-based modeling. *Ecological Modelling* 393:153–156.
- MacArthur, R. H., and E. R. Pianka. 1966. On optimal use of a patchy environment. *American Naturalist* 100:603–609.
- Mangel, M., and C. W. Clark. 1986. Toward a unified foraging theory. *Ecology* 67:1127–1138.
- Preisser, E. L., D. I. Bolnick, and M. F. Benard. 2005. Scared to death? The effect of intimidation and consumption in predator-prey interactions. *Ecology* 86:501-509.
- Railsback, S. F. and V. Grimm. 2019. *Agent-based and individual-based modeling: a practical introduction*, 2nd edition. Princeton University Press, Princeton, New Jersey.

- Railsback, S. F., and B. C. Harvey. 2002. Analysis of habitat selection rules using an individual-based model. *Ecology* **83**:1817–1830.
- Railsback, S. F., and B. C. Harvey. 2013. Trait-mediated trophic interactions: is foraging theory keeping up? *Trends in Ecology & Evolution* **28**:119–125.
- Railsback, S. F., and B. C. Harvey. 2020. Modeling populations of adaptive individuals. Princeton University Press, Princeton, New Jersey.
- Railsback, S. F., B. C. Harvey, R. H. Lamberson, D. E. Lee, N. J. Claasen, and S. Yoshihara. 2002. Population-level analysis and validation of an individual-based cutthroat trout model. *Natural Resource Modeling* **15**:83–110.
- Railsback, S. F., B. C. Harvey, J. W. Hayse, and K. E. LaGory. 2005. Tests of theory for diel variation in salmonid feeding activity and habitat use. *Ecology* **86**:947–959.
- Railsback, S. F., B. C. Harvey, and D. Ayllón. 2020. Contingent tradeoff decisions with feedbacks in cyclical environments: testing alternative theories. *Behavioral Ecology* **31**:1192–1206.
- Railsback, S. F., D. Ayllón, and B. C. Harvey. 2021. InSTREAM 7: Instream flow assessment and management model for stream trout. *River Research and Applications* **37**:1294–1302.
- Railsback, S. F., and M. D. Johnson. 2011. Pattern-oriented modeling of bird foraging and pest control in coffee farms. *Ecological Modelling* **222**:3305–3319.
- Strand, E., G. Huse, and J. Giske. 2002. Artificial evolution of life history and behavior. *American Naturalist* **159**:624–644.
- Sutton, R. S., and A. G. Barto. 2018. Reinforcement learning: an introduction. MIT press, Cambridge, Massachusetts.
- Werner, E. E. and S. D. Peacor. 2003. A review of trait-mediated indirect interactions in ecological communities. *Ecology* **84**:1083-1100.