

Stochastic Dynamic Programming Illuminates the Link Between Environment, Physiology, and Evolution

Marc Mangel

Received: 24 December 2013 / Accepted: 30 April 2014 / Published online: 18 July 2014
© Society for Mathematical Biology 2014

Abstract I describe how stochastic dynamic programming (SDP), a method for stochastic optimization that evolved from the work of Hamilton and Jacobi on variational problems, allows us to connect the physiological state of organisms, the environment in which they live, and how evolution by natural selection acts on trade-offs that all organisms face. I first derive the two canonical equations of SDP. These are valuable because although they apply to no system in particular, they share commonalities with many systems (as do frictionless springs). After that, I show how we used SDP in insect behavioral ecology. I describe the puzzles that needed to be solved, the SDP equations we used to solve the puzzles, and the experiments that we used to test the predictions of the models. I then briefly describe two other applications of SDP in biology: first, understanding the developmental pathways followed by steelhead trout in California and second skipped spawning by Norwegian cod. In both cases, modeling and empirical work were closely connected. I close with lessons learned and advice for the young mathematical biologists.

Keywords Stochastic dynamic programming · Parasitoids

Dedicated to my mentor Donald Ludwig on the occasion of his 80th birthday.

M. Mangel (✉)
Center for Stock Assessment Research and Department of Applied Mathematics and Statistics,
University of California, Santa Cruz, CA, USA
e-mail: msmangel@ams.ucsc.edu

M. Mangel
Department of Biology, University of Bergen, Bergen, Norway

1 Understanding, Puzzles, and Surprises in Science

Dayton and Sala (2001, pg 206) wrote “We hold that the goal of useful science is to make *interesting accurate generalizations about nature based on as few relevant parameters as necessary*...By accurate, they must be based on Scotus’ reality [that reason and observation of nature are more powerful means of ascertaining the truth than authority or faith]. While trivial generalizations abound, good science, to have value, must produce generalizations interesting to a wider audience. Finally, the relevant parameters are meant in an exclusive sense; good science must weed out the marginally relevant parameters because all of nature is trivially related. The generalizations must be based on those few parameters that can account for most of the uncertainty [variability], following Ockham’s parsimony principle” (italics in the original).

Mathematical methods and models clearly have a powerful role in achieving useful science. As Charles Darwin, perhaps the greatest theoretician of biology, showed, there is no substitute for immersing oneself in the natural world, either through controlled experiment or careful observation. However, such immersion often leads to puzzles—“why is nature like that?”—and mathematical models can help solve those puzzles. Furthermore, the models may lead to unexpected predictions (surprises) that can then be tested by empirical means (again either observations or experiment), and it is through solving puzzles and responding to surprises that we increase scientific understanding.

In this paper, I will explain how one kind of stochastic optimization method—stochastic dynamic programming (SDP)—has been used to solve puzzles in the biology of organisms, particularly those about behavior and development (growth and sexual maturity leading to reproduction) at the level of the individual organism. It is impossible in this short paper to give a full review of SDP, which has about 550,000 hits in Google Scholar (April 2014) and many books written about it. The origin of the phrase “Stochastic Dynamic Programming” is usually associated with the great mathematician Richard Bellman (1920–1984) who wrote a series of papers and books in the early 1950s introducing and explaining the methods (Bellman 1952, 1954, 1956, 1957). However, the method of dynamic programming is deeply rooted in classical variational problems of applied mathematics, going back to the work of William Rowan Hamilton and Carl Gustav Jacobi in the early nineteenth century (summarized in Courant and Hilbert 1962). Indeed, control theorists often refer to the fundamental equation of SDP as the Bellman–Hamilton–Jacobi equation.

Regardless of the kind of biology in which one is interested, simply stepping outside and briefly observing the world shows that organisms live in risky environments. Before an organism can successfully reproduce, it must gain the resources needed for reproduction while avoiding being killed. There are no Darwinian demons—organisms born fully mature, ready to reproduce, and that live forever. Predicting how animals deal with these two risks was once known as the problem of the common currency (e.g., Krebs and Davies 1978). That is, the natural units for foraging gain are $\frac{\text{energy}}{\text{time}}$ (or energy accumulated over an interval), but the natural characterization of mortality is the rate of death (or the probability of death over an interval), which is unitless, and we are taught from an early age that without commensurate units things cannot

be compared. Here is a puzzle to be solved (see [Lima and Dill 1990](#); [Nonacs and Dill 1990](#) for a status report about 15 years ago).

In the early 1980s, Colin Clark and I and Alasdair Houston and John McNamara independently realized that SDP would allow a uniform treatment of both kinds of risk, thus solving the puzzle of the problem of the common currency. Furthermore, we saw that SDP allows us to connect environment and physiology in a consistent Darwinian framework in which natural selection acts on trade-offs between traits and functional abilities of organisms. This recognition allowed us to attack many problems in which one wants to understand the pattern of behavior or development of organisms, and how those are linked to population dynamics. Furthermore, because the models focus on the individual organism, they are well suited for empirical tests, either by experiment or observation, and thus allow confrontation between the models and data (*sensu* [Hilborn and Mangel 1997](#)). Beginning in 1986 ([McNamara and Houston 1986](#); [Mangel and Clark 1986](#)), we jointly ([Houston et al. 1988](#)), pairwise (summarized in [Clark and Mangel 2000](#); [Houston and McNamara 1999](#); [Mangel and Clark 1988](#); [Mangel and Ludwig 1992](#)), and individually popularized the method and its application in behavioral ecology and beyond. When writing for biologists, we generally call this state-dependent life history theory, since it helps more sharply bring focus to what is being done.

In this paper, I will first derive the two canonical equations of SDP in biology (one for activity choice and one for resource allocation). I will then show how these models can be used to understand the behavior of insect parasitoids, which lay their eggs on or inside other insects [hosts] and whose larvae use the resources of the host to support development. I will then describe two critical experiments that changed the way people think about insect parasitoids, thus demonstrating the contribution of this kind of mathematics to biology. Before closing with a discussion, I give two vignettes showing how these methods have been used to understand the developmental patterns of steelhead trout in California and skipped spawning in Norwegian cod.

2 The Canonical Equations of Activity Choice and Allocation

To illustrate the thinking behind SDP models and state-dependent life history theory, I first derive two of the most fundamental equations for using these methods in biology. These equations are not intended to capture the specifics of any particular system, but rather they are intended to have much in common with many systems, which is what has made them so powerful. That is, they are valuable precisely because they apply to no system in particular, but apply to many systems in general and can be tailored to specific biological problems, as I will illustrate in subsequent sections. We study frictionless springs (or even ones in which the friction is proportional to velocity) for the same reasons.

Throughout, I use a framework in which time is treated discretely rather than continuously, for three reasons. First, this avoids the problem of having to decide which kind of stochastic calculus to use (of which there are an infinite number, not just Ito and Stratonovich, see [Krener 1979](#)); furthermore, the stochasticity that we deal with is often non-Gaussian. Second, using a discrete-time framework from the outset

means that the equations we derive are difference rather than differential or partial differential equations. This helps maximize the number of potential users one can reach. Third, these equations are nonlinear, and thus generally require numerical solutions, which can be easily implemented on a desktop computer. The continuous-time framework leads to partial differential equations that generally have to be numerically implemented with small but discrete steps in time, so that one is back to a discrete formulation in any case.

The fundamental metric of evolution by natural selection is the change in the frequency of genotypes in a species or population over time. For most problems, instead of genotypes, we focus on allele frequencies, but even then the effects of changes of allele frequencies on phenotypic traits are usually too hard to compute or predict, since we still lack fundamental understanding on the genetic architecture of most polygenic traits (Giske et al. 2013). So, we use a proxy, refer to it as fitness, and think of it as the expected reproduction or grand offspring of a focal individual. The methods described here can be extended, with some measure of complication, for computing longer term representation of offspring (Mangel et al. 1994; Houston and McNamara 1999, McNamara 2000).

The two canonical equations of SDP share the following features. We consider a season of length T , with t denoting time within the season and the first period of the season $t = 1$. The organism of interest is characterized by a single physiological state $X(t)$ at time t (the dynamic variable) with specific value x .

2.1 The Canonical Equation of Activity Choice

To begin, we assume that the individual does not reproduce between $t = 1$ and $t = T$, so that this is a non-breeding interval, at the end of which the organism has future fitness $\Phi(X(T))$ depending upon its state at that time. This situation clearly applies to juveniles during their developmental period, or migrating birds going from breeding grounds to feeding grounds and then back.

At each time period $t = 1, 2, \dots$ before T , the organism chooses an activity from a collection of activities $i = 1, 2, \dots, I$. In the original formulation, we (Mangel and Clark 1986, 1988) called this the patch selection problem and thought of a forager that each morning chooses to either remain in its burrow (safe from predation but losing energy) or forage in one of the two patches, one more profitable but riskier than the other.

Each activity is characterized by (1) its cost α_i , (2) its risk m_i in the sense that the probability of surviving a single period of that activity is e^{-m_i} , (3) the probability of increasing state if this activity is chosen λ_i , and (4) the increment in state Y_i if that occurs. Thus, if activity i is chosen when $X(t) = x$, the organism survives to time $t + 1$ with probability e^{-m_i} ; with probability λ_i

$$X(t + 1) = x - \alpha_i + Y_i \quad (1)$$

and with probability $1 - \lambda_i$

$$X(t + 1) = x - \alpha_i. \quad (2)$$

We now introduce the fitness function

$$F(x, t) = \max E[\Phi(X(T)|X(t) = x)], \tag{3}$$

where “max” denotes that the maximum is taken over behavioral choices (here $i = 1, \dots, I$ done at each t and x) and E denotes mathematical expectation, taken over the stochastic processes of survival and incrementing the state. Thus, $F(x, t)$ is the maximum expected value of the future fitness at time T given that the current value of the state is x .

To understand the power of SDP, imagine that we wanted to compute $F(x, t)$ forward in time. We would specify $X(1) = x$, then make an activity choice, from which we would simulate whether the organism survives or not and finds food or not. In this case, when individuals survive, the new state ($X(2)$) will be either $= x - \alpha_i + Y_i$ or $x - \alpha_i$ depending upon whether food is found or not. From this new state, we once again choose from one of the I activities and compute $X(3)$ for survivors. We have to repeat this process over and over again until $t = T$. At that point, we have one value of the terminal fitness, depending upon the starting point and the activity choices. But if we want to maximize $\Phi(X(T))$ given the initial state x , we have to investigate, for every starting state and all subsequent ones, every possible activity choice! Sitting with pen and paper and drawing out the tree of possibilities leads one to conclude that there must be a better way finding the pattern of state-dependent activity that maximizes $\Phi(X(T))$ given x at a previous time t . That is, we compute $F(x, t)$ backwards in time.

In light of its definition (Eq. 3), $F(x, t)$ satisfies the end condition

$$F(x, T) = \Phi(x) \tag{4}$$

Consequently, we already know its value at the end of the time interval of interest: $F(x, t)$ is a dynamical entity, but we must focus on its evolution backwards in time. Individuals familiar with Kolmogorov’s approach to probability, in which there are both forward and backward equations for the dynamics of probabilistic quantities (e.g., [Karlin and Taylor 1981](#); [Mangel 2006](#)), will recognize commonalities but not everyone is familiar with this approach to probability.

We already know $F(x, T)$, so imagine stepping backward one time unit to $T - 1$ and that when $X(T - 1) = x$, activity i is chosen. Then, the probability of surviving from $T - 1$ to t and incrementing state is $e^{-m_i} \lambda_i$. Conditioned on those events, the future expected fitness is $F(x - \alpha_i + Y_i, T - 1 + 1) = F(x - \alpha_i + Y_i, T)$ so that the future fitness when activity i is chosen, state is incremented, and the organism survives is $e^{-m_i} \lambda_i F(x - \alpha_i + Y_i, T)$. Similar reasoning applies to the case in which state is not incremented. Thus, future expected fitness if activity i is chosen is $e^{-m_i} [\lambda_i F(x - \alpha_i + Y_i, T) + (1 - \lambda_i) F(x - \alpha_i, T)]$. Since $F(x, T - 1)$ is understood to be maximum fitness taken over activity choice, we conclude $F(x, T - 1) = \max_i e^{-m_i} [\lambda_i F(x - \alpha_i + Y_i, T) + (1 - \lambda_i) F(x - \alpha_i, T)]$. Following this procedure or every value of x allows us to compute $F(x, T - 1)$.

Now that we know that $F(x, T - 1)$ for every value of x , we could step back in time to $t = T - 2$. Iterating this argument, we conclude that the end condition Eq. 4 is complemented by the equation of SDP

$$F(x, t) = \max_i e^{-m_i} [\lambda_i F(x - \alpha_i + Y_i, t + 1) + (1 - \lambda_i) F(x - \alpha_i, t + 1)], \quad (5)$$

which is the canonical equation for activity choice.

2.2 The Canonical Equation of Allocation Processes

Let us now turn to a different kind of problem. Instead of building up resources for future fitness assessed at time T , consider a case in which the organism accumulates fitness throughout the season and then dies at time T . This is true for many insects that live for only one season (or even for less than a season—sometimes just a few days, as trout fishers know) and lay their eggs during that season, with the offspring emerging the following one. Multiple seasons, such as that occur for long-lived birds, can be treated by a modification of the model here, in which we couple the end of one season with the start of the next (Mangel and Clark 1988; Clark and Mangel 2000).

Suppose that there is a payoff to the individual such that if c resources are used in a period, then the increment in fitness is $f(c)$ where $f(0) = 0$ and $f(c)$ is a concave function, at least far away from 0. That is, there are diminishing returns from additional use of resources as c increases. We also allow a stochastic increment in resources, R , which occurs with probability λ in each period. Here—and indeed above—these could be time dependent, but for simplicity, they are taken as constant. As above, we also assume a constant per-period mortality.

In this case, we let $F(x, t)$ denote the maximum expected accumulated fitness between t and T given that $X(t) = x$. Assuming that no fitness can be accumulated at time T or beyond, we have the end condition $F(x, T) = 0$. For previous times, we reason as follows. Assume that allocation is made before the stochastic events of incrementing resources or not and surviving or not. Then, if $X(t) = x$ and an allocation of c is used, then the state at the start of period $t + 1$ will be $x - c + R$ if resources are incremented and $x - c$ otherwise; these events occur with probability λ and $1 - \lambda$, respectively. Thus, in analogy to the logic that lead to Eq. 5, we conclude

$$F(x, t) = \max_c (f(c) + e^{-m} [\lambda F(x - c + R, t + 1) + (1 - \lambda) F(x - c, t + 1)]). \quad (6)$$

Equations 5 and 6 summarize all of stochastic dynamic programming in biology; all that are left are to fill-in the details—which can take a lifetime.

2.3 Numerical Solution

Even if Eqs. 5 and 6 were written in continuous time, they rarely have analytical solutions (for a counter-example see Mangel 1992). Often, some kinds of preliminary analysis can be conducted on the model. However, numerical solution of the SDP equation can provide exceptional insight—both qualitative patterns and detailed numerical predictions. That is, by conducting sensitivity analyses, we can develop the same kinds of intuition that mathematical analysis often provides. Indeed, as noted in

Mangel and Clark (1988) and Clark and Mangel (2000), very often the intuition from a numerical model can be so powerful that one no longer needs the model to understand the phenomenon. New and deep understanding of the biology of the problem of interest often emerges from building, running, and interpreting the model. This is one of the reasons that Colin Clark and I declined offers to develop software packages for SDP—because the biologist needs to do the modeling in order to achieve maximum intellectual gain from the model. Depending upon the problem that one investigates, there may be a variety of other numerical issues that have to be addressed—such as non-integer values of state—but these can be dealt with (Mangel and Clark 1988; Clark and Mangel 2000).

Now let us get to some of the specific examples.

3 Application to Insect Parasitoids

3.1 The Classification of Parasitoid Life Styles

Insect parasitoids deposit their eggs on or in the eggs, larvae, or adults of other insects. After the offspring hatch, they consume the host body to fuel their development. Insect parasitoids are commonly used for biological control of insect pests, so understanding them with a view of their application firmly puts this work in Pasteur's Quadrant (Stokes 1997) in which one is motivated by an applied problem but seeks fundamental understanding.

Parasitic wasps, in particular, are incredibly rich and diverse, which is one of the things that makes them so interesting. Parasitoid life histories can be classified using population, behavioral, and physiological criteria (Fig. 1a). First, parasitoids may have one generation (univoltine) or more than one generation (multivoltine) per calendar year. Second, females may lay one egg (solitary) or more than one egg (gregarious) in hosts. Third, females may be born with essentially all of their eggs (pro-ovigenic) or may mature eggs (synovigenic) throughout their lives. Each dichotomous choice leads to a different kind of life history.

3.2 The Pro-ovigenic, Gregarious Parasitoid

3.2.1 The Puzzle

A useful starting point for our discussion is the pro-ovigenic and gregarious parasitoid that lives for one season, of length T (after which the parasitoid dies, for example because the first freeze has come), has all of her eggs upon emergence (or at least a very large number), but may lay more than one egg in a host. For example, in *Aphytis lingnanensis*, the average size of a daughter emerging from a clutch of size c is the decreasing function of clutch size $S(c) = \max[0.2673 - 0.0223c, 0]$, and the number of eggs a female harbors depends upon her size according to $X_{\max}(S) = \max[181.8S - 26.7, 0]$ (Rosenheim and Rosen 1992). Combining these, we conclude that if a female lays a clutch of size c in a host, the grand offspring produced from this clutch is $cX_{\max}(S(c))$, which is a concave function of the clutch size (Fig. 1b).

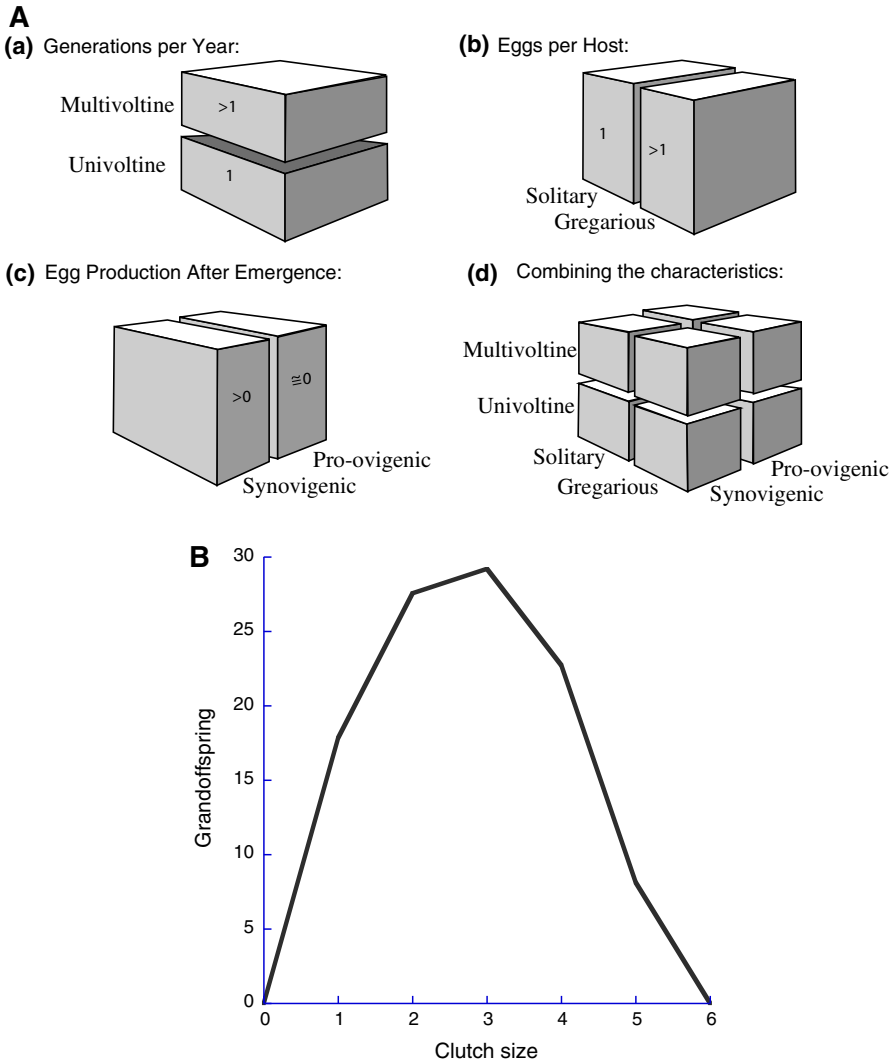


Fig. 1 **a** A classification scheme for parasitoid life histories (from Mangel 2006). **b** The fitness increment, approximated by the number of grand offspring, for an *Aphytis lingnanensis* laying eggs in a single host (from Mangel 2006)

Prior to the introduction of state-dependent life history, predicting the number of eggs a female lays in hosts proceeded along three lines. First, one might note that in a single host, a clutch of size 3 produces the largest number of potential grandchildren, and predict that females will lay this single host maximum, called the Lack clutch size after the great evolutionary biologist David Lack whose work on birds was foundational (Charnov and Skinner 1984, 1985; Mangel and Clark 1988; Clark and Mangel 2000). Second, because the curve in Fig. 1b is concave, the highest per egg increment in fitness comes from laying one egg, so that we could predict that females only lay

one egg per host. Third, one could apply the classical marginal value theorem of foraging theory (Charnov and Skinner 1984, 1985; Mangel 2006, pp. 5–8) to this case by interpreting clutch size in terms of oviposition time, extending the x -axis of Fig. 1b to the left, noting the travel time between hosts and drawing the line tangent between the point on the x -axis corresponding to travel time and the curve, thereby reading off the optimal oviposition time (and thus clutch). All of these methods predict that there is a single optimal clutch, independent of egg complement, rate of mortality, or time in the season.

3.2.2 Theory

Let us now consider a SDP model, building off of the canonical equation for allocation. Because the insect is pro-ovigenic, there is no increment in resources so that $R = 0$. Suppose that there are I different kinds of hosts, each with (a) an associated fitness increment $f_i(c)$ (determined for example, by the size of the host), (b) encounter probability per unit time λ_i , (c) rate of mortality m_i while ovipositing, and (d) that laying c eggs requires time $\tau(c)$, with the understanding that $\tau(0) = 1$.

Imagine that the parasitoid with egg complement x at time t has encountered host type i , which occurs with probability λ_i . If she lays c eggs in this host, which takes time $\tau(c)$, then she immediately receives an increment in lifetime fitness equal to $f_i(c)$, survives to time $t + \tau(c)$ with probability $e^{-m_i\tau(c)}$, and starts looking for hosts again at that time with $x - c$ eggs. Thus, given that a host of type i has been encountered and she lays c eggs, the combination of current and future fitness is $f_i(c) + e^{-m_i\tau(c)} F(x - c, t + \tau(c))$. For this host type, time, and egg complement, the maximum fitness is then found by picking the value of c that gives the largest combination of current and future fitness. However, at the start of period t , the parasitoid has not yet encountered a host, so fitness is computed by averaging overall possible host encounters, with the understanding that $\lambda_0 = 1 - \sum_{i=1}^I \lambda_i$ is the probability of not encountering any host and $f_0(c) = 0$ for all c . Thus, the analog of Eq. 6 is

$$F(x, t) = \sum_{i=0}^I \lambda_i \max_c [f_i(c) + e^{-m_i\tau(c)} F(x - c, t + \tau(c))]. \tag{7}$$

The numerical solution of Eq. 7 is straightforward, and consideration of the equation suggests that we could use experimental methods to investigate how clutch size depends upon egg complement, encounter rate with hosts, and natural mortality (or cues of natural mortality). Figure 2a shows the prediction of the first period clutch of *Aphytis* for two values of the rate of natural mortality for the case of a single host type (see Clark and Mangel 2000, Ch 4 for further details).

A strong prediction of this theory, which was published in its first form in Mangel (1987), is that clutch size will increase with increasing egg complement and rate of mortality and decrease as t moves further away from T . Although these may seem intuitive, they were not when first published and they seem intuitive now because of the way the mathematical argument has been constructed. Since previous theories lacked consideration of state and mortality, they could not capture the predictions based on SDP.

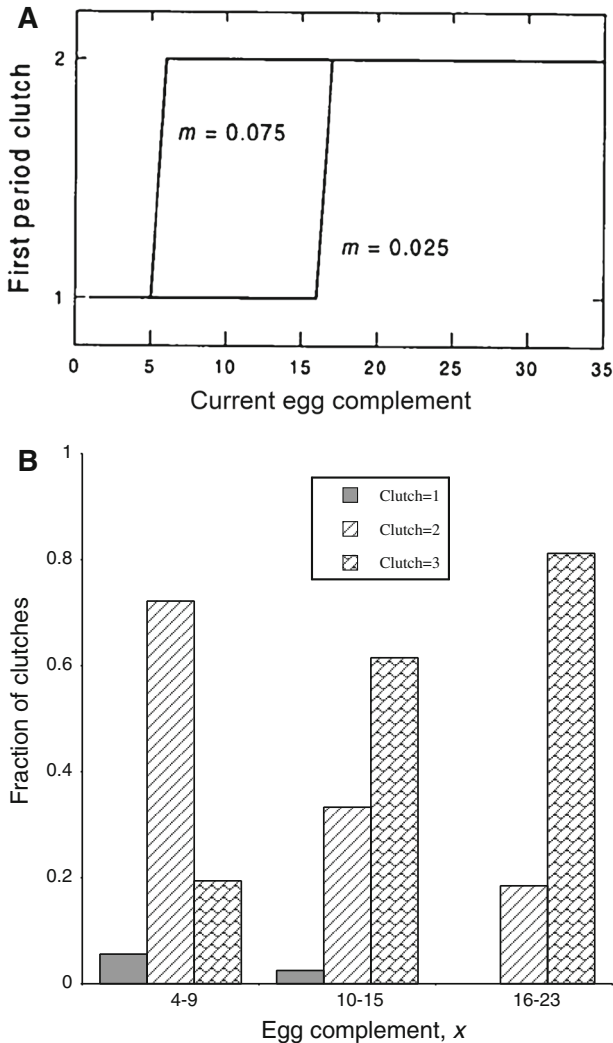


Fig. 2 The simplest state variable model for insect oviposition predicts that clutch size will respond to egg complement, mortality rate, and encounter rate. **a** Predictions of first period clutch for the fitness increment of *Aphytis* for a single host with $\lambda = 0.5$, maximum egg complement $x_{max} = 35$ and $T = 60$. **b** Results of the clutch size experiment conducted by Rosenheim and Rosen. The raw data can be found in Hilborn and Mangel (1997, chapter 6) and other versions of the aggregated data in Rosenheim and Rosen (1991)

3.2.3 Experiment

In his post-doctoral work with David Rosen, Jay Rosenheim tested these predictions (Rosenheim and Rosen 1991, 1992). One of their results is shown in Fig. 2b.

They used *Aphytis lingnanensis*, which is actually synovigenic; however, they chose protocols to ensure that parasitoids had matured a large complement of eggs before behavior was assayed. The hosts used in the experiments were large virgin third instar

females of the California red scale *Aonidiella aurantii*, maintained as virgins because the females develop a hard protective shell after mating. Rosenheim and Rosen provided parasitoids having different egg complements with uniform experience of host encounter. They manipulated egg load by using parasitoids of different sizes or raising parasitoids at low temperature, which slows egg production. They confined parasitoids with a single host, and observed oviposition behavior. After that, the parasitoid was given a second host, and only parasitoids ovipositing in both hosts were used in the data collection; after the second oviposition, they dissected parasitoids to determine the remaining egg complement. The egg complement at the start of the experiment was the sum of the eggs laid plus those counted during dissection. Virtually, all of the clutches were one, two, or three eggs and as egg complement increased, the frequency of larger clutches increased (Fig. 2b). Indeed, no clutches of size 1 were observed for an individual that had more than 13 eggs.

In summary, the experiments of Rosenheim and Rosen support the conclusion that oviposition behavior is fundamentally dynamic and responds to changes in physiological state. Now, to be sure, dissecting individual parasitoids after having ensured that they had controlled encounters with hosts is a lot of work. But this effort was rewarding in showing that parasitoids did indeed respond to egg complement when making oviposition decisions.

3.3 Superparasitism and Environmental Conditions

I now turn to an example of using the canonical equation of activity choice, in this case for a solitary parasitoid, *Leptopilina heterotoma*, which attacks the eggs of *Drosophila*. Solitary parasitoids differ from gregarious ones in that usually only one individual emerges from a host, even if more than one egg is laid. When a female lays an egg in a host that already has an egg, she is said to be superparasitize. If she is laying the second egg close to when the first egg was laid, then the chance that her egg emerges is generally greater than if the eggs are laid at two very different times, because in the latter situation, the first offspring usually outcompetes and often kills the second one.

In this case, the puzzle is to predict when a female is predicted to oviposit in a previously parasitized host.

3.3.1 Theory

We modify the canonical equation for activity choice as follows. Consider that there are only two kinds of hosts. Unparasitized hosts are encountered with per-period probability λ_u and provide an increment in lifetime fitness f_u , still measured in terms of potential grandchildren. Previously, parasitized hosts are encountered with per-period probability λ_p and provide an increment in lifetime fitness f_p . We also assume that the rate of mortality while searching for hosts is m_s and while ovipositing is m_o and that the time required for oviposition is τ_o . The constraints on these parameters are that they are all positive, that $\lambda_u + \lambda_p \leq 1$, that $f_u > f_p$, and that $m_o \geq m_s$ (*Leptopilina* is almost frozen on the host as she oviposits). Since parameters are constant, it is easy

to see that if an unparasitized host is encountered, then we predict that the female will always oviposit in it.

Suppose we let $V_{reject}(x, t)$ denote the fitness value of rejecting a previously parasitized host given that $X(t) = x$ and a previously parasitized host is encountered. Then

$$V_{reject}(x, t) = e^{-m_s} F(x, t + 1). \quad (8)$$

If $V_{accept}(x, t)$ is the fitness value of accepting a previously parasitized host under the same conditions, then it is determined by the combination of current and future fitness

$$V_{accept}(x, t) = f_p + e^{-m_o} F(x - 1, t + \tau_o). \quad (9)$$

The canonical equation for activity choice becomes

$$F(x, t) = (1 - \lambda_u - \lambda_p) e^{-m_s} F(x, t + 1) + \lambda_u [f_u + e^{-m_o} F(x - 1, t + \tau_o)] + \lambda_p \max[V_{reject}(x, t), V_{accept}(x, t)]. \quad (10)$$

When this equation is solved, we generate fitness $F(x, t)$ as a function of state and time, and the boundary in egg complement–time space at which the inferior host is predicted to be accepted (Fig. 3a). In addition, we can generate a similar boundary in the plane determined by the value of m_s and the time to start superparasitizing (Fig. 3b); this is actually a three-dimensional plot because of the dependence on egg complement, so the figure is a slice through the three-dimensional surface.

3.3.2 Experiment: Growth Chambers and Thunder Chambers

Figure 3b suggests two experiments. First, imagine that the rate of mortality is fixed but we give the parasitoid a cue that it is later in the season than it really is. The consequence of this is that a point in the m_s /time to start superparasitizing plane moves upward and crosses the boundary (Fig. 4a). Alternatively, if time was fixed, but we gave the parasitoid a cue that mortality is about to greatly increase, then we can cross the boundary horizontally (Fig. 4b). Bernie Roitberg and I did these experiments (Roitberg et al. 1992, 1993).

We used a Dutch (i.e., temperate zone) strain of *L. heterotoma* obtained from the University of Wageningen in the Netherlands. The strain was reared on *D. simulans* under summer conditions (20 °C, 16 h light: 8 h dark) for approximately 20 generations before experiments began. We used 4- or 5-day-old wasps for all experiments and maintained wasps on dilute honey and water. We tested each wasp once and then placed in the colony cage to help maintain the culture.

Leptopilina heterotoma is an appropriate animal for testing predictions from our theory for three reasons. First, it has a relatively long life span (up to 30 days as an adult in the laboratory). This means that a 4- or 5-day-old individual can expect to live a further 10 days or so during the summer in the field. Second, the Dutch strain regularly

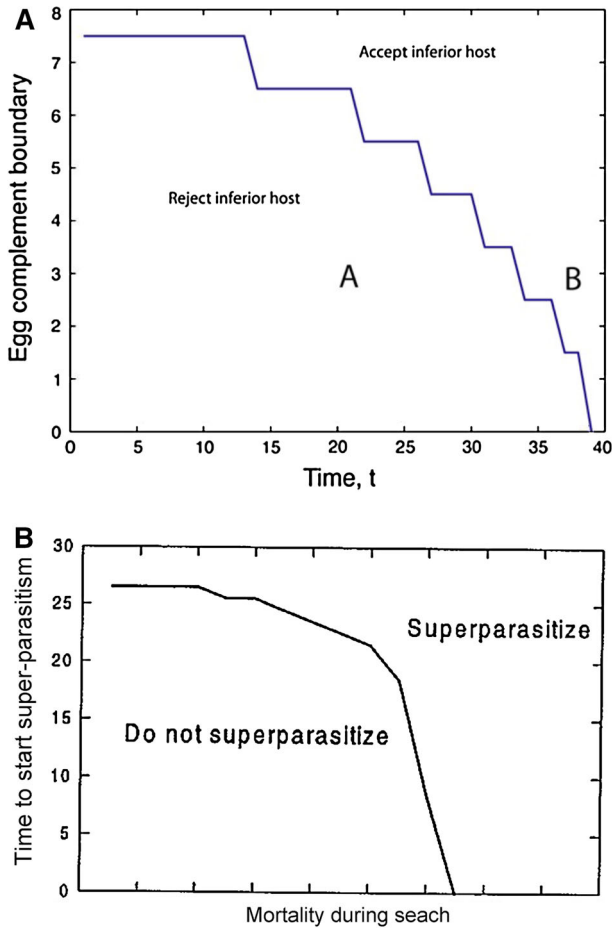
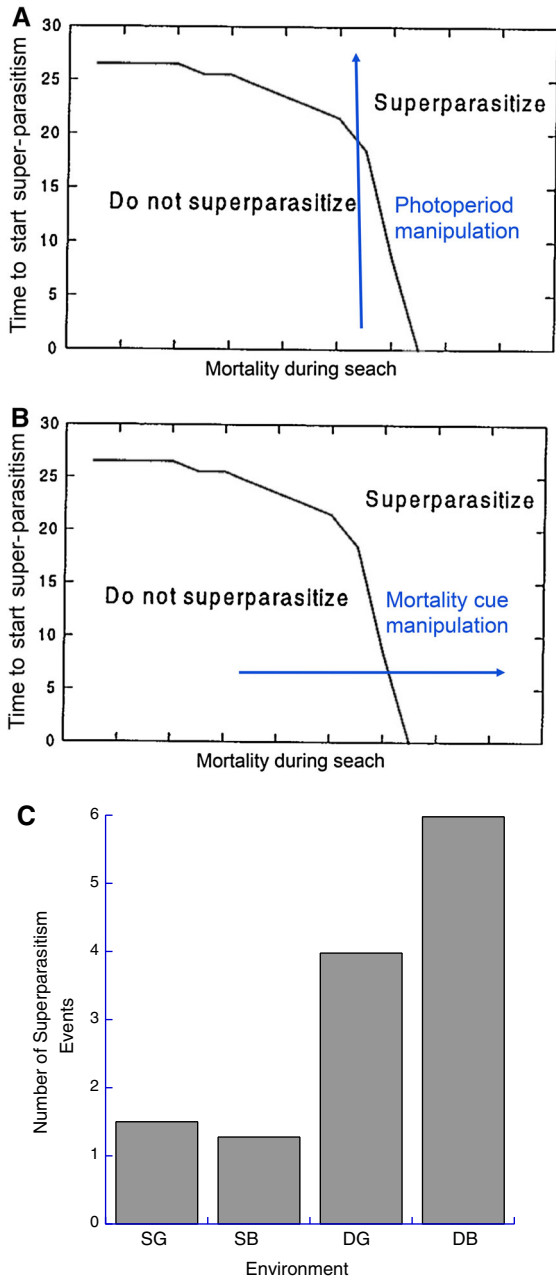


Fig. 3 **a** The solution of Eq. 10 generates a boundary in the time–egg complement plane separating regions in which we predict that a previously parasitized host will be rejected (Region A) and one in which it will be accepted (Region B). **b** The theory also generates predictions about the relationship between the rate of mortality during search and the time at which the parasitoid will start superparasitizing for fixed egg complement

experiences cool conditions during the fall, meaning that it would be unable to forage effectively even before the weather became so cold that it died. Third, *Drosophila* hosts are frequently very patchily distributed.

To manipulate perceived life expectancy (i.e., the difference between current time t and T), we reared *L. heterotoma* on larvae of *D. simulans* under two photoperiods: summer (16 h light: 8 h dark) and fall (12 h light: 12 h dark) and held temperature constant at 22 °C. We assumed that the fall photoperiod treatment would provide recently emerged wasps with information that t is close to T , whereas the summer photoperiod would indicate that $t \ll T$. We manipulated the perceived quality of patches by altering the quality of hosts wasps experienced within those patches. We did

Fig. 4 a Given a cue that t is closer to T than it actually is, we predict that the parasitoid will switch from rejecting previously parasitized hosts to accepting them. **b** Similarly, given a cue that the rate of mortality is much higher than it actually is, we also predict that the parasitoid will switch from rejecting previously parasitized hosts to accepting them. **c** We conducted an experiment in which we measured the number of superparasitism events for parasitoids that had environments with Steady (S) or Dropping (D) barometric pressure—the latter indicating a rain storm is on its way—and few previously parasitized hosts (Good environments) or many previously parasitized hosts (Bad environment)



this by setting circular 3-cm-diameter (1-mm height) viscous suspensions of baker's yeast onto the center of 0.5-cm-deep agar bases set within 10-cm-diameter glass petri dishes. Thirty minutes after preparation, each yeast patch received 30 48-h-old *D. simulans* larvae. One hour later, we collected two 4- to 7-day-old *L. heterotoma* females

from our stock colony and released them onto half of the patches. If a female did not actively search for larvae, then we replaced her. We allowed the wasps to search their patches and parasitize hosts until both females left the patch and moved to the lid of the petri dish. Preliminary observations showed that wasps generally did not abandon patches until most, if not all, of the hosts had been parasitized. Thus, this treatment produced two types of patches: all larvae unparasitized and essentially all larvae parasitized.

We divided summer-treatment wasps into two prior experience groups: Good World (G) and Bad World (B); fall-treatment wasps were only treated with Good World conditions. In the Good World treatment, we released wasps on yeast patches that contained unparasitized hosts. We released wasps in Bad World treatments on patches that contained the already-parasitized hosts. We allowed each female to search until she abandoned the patch for the lid of the petri dish, at which the point we returned her to a 50-ml glass vial containing dilute honey. We repeated the treatment the next day. On day 3, we coded all wasps by treatment and individually placed them on patches containing 30 *D. simulans* larvae that had already been parasitized by other *L. heterotoma*. We observed each wasp under a stereomicroscope and counted superparasitisms (ovipositions are easily counted in this wasp because of a very clear behavior when she lays an egg).

For simplicity, here I discuss only the Summer-Good World and Fall-Good World treatments (all of the results are found in [Roitberg et al. \(1992\)](#), Table 2). Wasps experiencing the fall photoperiod spent twice as much time looking for hosts (about 29 vs 13 min, different at the $p < 0.01$ statistical level); this is something not in the theory described above and I encourage a reader to think about how one might make this modification. They also superparasitized more, both in absolute number (5.4 on average vs 1.8 on average) and as a percentage of encounters (about 39 vs 28 %); both of these results also significantly different at the 0.005 level. Thus, part of the puzzle is solved: we confirmed that parasitoids become less choosy later in the season, as T approaches, presumably because the balance of current and future reproduction shifts toward current reproduction.

To test the prediction about a cue of impending mortality, we once again raised wasps with Good (G) or Bad (B) conditions with summer photoperiod, but then gave them either Steady (S) barometric pressure or Dropping (D) barometric pressure. We did this by enclosing the petri dishes in sealed containers that were connected to pumps which allowed us to maintain barometric press (S) or slowly drop it (D). Steady barometric pressure is intended to correspond to a fine summer day, while dropping barometric pressure is indicative of a large rainstorm on its way. The rest of the protocol was as above, and I show the results in [Fig. 4c](#). Parasitoids, given a cue that their expectation of life was small, greatly decreased choosiness. In addition, wasps that were placed in chambers with dropping barometric pressure and at sites with poor availability of hosts (DB) laid by far the most eggs in low-quality hosts when they encountered them, whereas those under steady barometric pressure and at a good site (SG) were by far the most choosy and laid the fewest eggs into low-quality hosts; the statistical significance here was at $p < 0.001$. As Bernie said in his 2008 Gold Medal address to the Entomological Society of Canada “We can conclude that wasps know the theory and apply it appropriately” ([Roitberg 2008](#), p. 164).

4 Two Vignettes: Steelhead Trout and Skipped Spawning in Cod

I now briefly describe another set of studies done by my group and then turn to work done by others. A search of the literature will lead to many other such papers, but the intention of this article is to be brief not comprehensive.

4.1 Steelhead Trout

My group conducted a coordinated research program involving models, field work, and laboratory experiments to explore and refine the application of state-dependent life history models to anadromy in *Oncorhynchus mykiss* in California, USA. *O. mykiss* is facultatively anadromous, meaning that individuals can complete their entire lifecycle within freshwater (rainbow trout), or they can follow an anadromous lifecycle (steelhead). The puzzle here is to explain when a juvenile fish will become a steelhead trout (and thus migrate to the ocean) or remain as a resident rainbow trout.

Our work began with modifying the model of [Thorpe et al. \(1998\)](#) to suit aspects of Northern California steelhead ([Satterthwaite et al. 2009](#)), which are distinct from the typical Atlantic salmon case and tested key model predictions using lab experiments and field observations of a well-studied population on a small coastal stream in Central California, Scott Creek. We then extended the modeling framework to explain differences in the life histories observed on small coastal streams such as Scott Creek, compared to larger rivers in the Central Valley of California (the Mokelumne River and the lower American River; see [Satterthwaite et al. 2010](#)), to predict impacts of environmental change and flow manipulation, and to explore the relaxation of some assumptions about timing to explore hypothesized explanations of some unexpected results observed in the field.

For Scott Creek, our model identified a size threshold of approximately 100–110-mm fork length at the end of December for a fish to initiate smolt development and emigrate the following spring. When coupled with observed sizes and growth rates, we predicted no Scott Creek fish initiating smolting as young of the year and thus no age-1 emigrants, along with substantially more fish reaching the threshold by age-1 with many age-2 emigrants, and the slowest growers reaching the threshold by age-2 and emigrating at age-3. We predicted a higher smolting size threshold for the American and Mokelumne Rivers, but most Mokelumne River fish were able to reach this higher threshold in their first year of life (due to more rapid growth) and essentially all American River fish surpassed it. Thus, the model predicted mostly young emigrants from the Mokelumne (although those fish not smolting were predicted to mature rather than emigrate at older ages) and all young emigrants from the American River. Empirical observations were largely, but not entirely, consistent with model predictions ([Sogard et al. 2012](#)).

The models performed well both qualitatively (in identifying the predominant life histories for watersheds with very different growth environments, as well as relative smolt threshold sizes) and quantitatively (identifying the most common age of smolting). Of course, no model can capture every detail observed in nature, nor should we expect it to do so. Nevertheless, mismatches between model predictions and observa-

tions suggest important new directions for theoretical and empirical exploration (see the cited papers for details). Taken on balance, these results suggest that the approach used in the state-dependent model of exploring life history variation in the context of growth/fecundity opportunities trading off with survival captures an important driver of life history variation in California steelhead.

4.2 Skipped Spawning in Arctic Cod

Here is a case in which a surprise from a model led to an empirical program. During his PhD studies at the University of Bergen, Norway, Christian Jørgensen developed life history models of Arctic cod *Gadus morhua* (Jørgensen and Fiksen 2006). The models were based on state-dependent life history theory, and Jørgensen used SDP to find the optimal state-dependent strategy for energy allocation between growth and reproduction. One of the unexpected predictions of these models was that under certain circumstances, the fish would forgo reproduction in a year (Jørgensen et al. 2006). This phenomenon of skipped spawning was predicted to occur early in life. Researchers at the Institute of Marine Research (IMR), Bergen, Norway, who had worked with the stock throughout their careers, said that skipped spawning did not occur. However, Russian colleagues had data suggesting that skipped spawning did indeed occur (Yaragina 2010). These data and the predictions from the SDP models motivated scientists at IMR to conduct a comprehensive field study with detailed histological examination of the fish. The results of the field study matched the predictions from the model with very high fidelity to the overall level and the age pattern of skipped spawning (Skjæraasen et al. 2012).

Skipped spawning had already been documented in other species of fish (Rideout et al. 2005) and is found throughout the animal kingdom (see Cayuela et al. 2014 for a brief review) but nobody knew why it occurred. For example, although half of the potentially second-time spawners of Norwegian herring *Clupea harangues* are missing from the spawning grounds based on scale readings, nobody knows where they are or why they are not on the spawning grounds (Engelhard and Heino 2005). The work of Jørgensen and colleagues motivated studies of skipped spawning in herring, and at this point, the question remains unresolved (e.g., Kennedy et al. 2011). There is anecdotal evidence that Antarctic toothfish (*Dissostichus mawsoni*) may also exhibit skipped spawning (I. Everson, personal communication), but in larger individuals, this has important implications for fishery management.

5 Lessons Learned: What has SDP Done for Biology?

In this brief paper, I have shown how stochastic dynamic programming, a classical method of applied mathematics, can help us solve puzzles in biology, and lead to unexpected predictions. As I have used it, SDP is a technique for modeling individual behavior; there are other methods for modeling individuals and a good introduction to them is Railsback and Grimm (2012). However, SDP remains the fundamental method for understanding how natural selection shapes trade-offs in behavior and development.

Although I focussed on insect behavior with brief discussions of fish, there are three general lessons to be learned from the development of state-dependent life history theory, as implemented by SDP.

First, details matter. The power of mathematical methods is that they allow us to see the commonalities of apparently disparate phenomena. However, empiricists must be deeply involved with the details of the particular systems that they study or they would never make progress. Hence, the mathematical biologist who wants to help empiricists must also care about the details. In my opinion, we advance the field most rapidly by focussing on specific biological questions and then developing general methods to answer them. The developments of general equations for the dynamics of populations by Lotka and Volterra or for the firing of nerves by Hodgkin and Huxley are examples of this. In this regard, it is helpful to remember that if we are to really affect biology, then the methods we develop and use must be transferrable and accessible; telling a biologist that he or she really needs an undergraduate degree in mathematics (usually meant to be pure mathematics, rather than applied mathematics) before being able to use mathematical methods to solve problems in biology is a guarantee to drive people away.

Second, be patient. It takes a while for new ideas to be understood and absorbed. When I published the first paper on the state variable approaches to insect behavioral ecology in 1987, a colleague wrote to me that this was clearly the best way to approach the issue but that he was too lazy to make the requisite measurements. As I described above, three years later, another colleague, Jay Rosenheim, had done so. At the Entomophagous Insect Workshop in 1994, everyone was talking about the importance of the state of the insect, not just the rate of gain of fitness.

Third, seek to inspire experiments and/or field observations. Richard Feynman began his lectures on physics: “The principle of science, the definition, almost, is the following: *The test of all knowledge is experiment*. Experiment is the *sole judge* of scientific ‘truth’. But what is the source of knowledge?” (Feynman et al. 1963, p. 1, paragraph 5). If we are patient and care about the details, then our mathematical models can be the source of the knowledge that motivates empiricists. If it is done right, then mathematical biology can be the foremost way of guiding empirical work. Even better—mathematical biologists can themselves learn (as Feynman did) how to do experiments. Give yourself a rotation in a biology laboratory of your choice.

Finally, we should always keep in mind Samuel Karlin’s comment when he gave the 11th RA Fisher Memorial Lecture (20 April 1983): “The purpose of models is not to fit data but to sharpen questions.”

6 Conclusion: Here is to You, Nicolas Rashevsky

This journal was founded (under a somewhat different name) by Nicolas Rashevsky. As an undergraduate at the University of Illinois, in fall 1969, I discovered his books on mathematical biophysics, which had been republished by Dover Press (Rashevsky 1969). These books had a profound effect on my life. I was a physics major; in the upper division course on electricity and magnetism, we did complicated calculations in the morning and tested them in the laboratory in the afternoon. This was breathtaking, and I decided that I wanted a career blending theory and experiment, but in biology

(of which I also had a great love, particularly plant biology and marine biology). Rashevsky's books were the first indication that such a career is possible.

I still own these books, the bindings retaped many times. Looking at them again for this paper, I am struck particularly by the confrontation between models and data in them (*sensu* Hilborn and Mangel 1997). Mathematical biology has an illustrious past and a very bright future.

Acknowledgments My work on insect behavioral ecology was financially supported over the years by the NSF and the USDA and intellectually supported by the wonderful collaborations with Bernie Roitberg and Jay Rosenheim. The work on steelhead (which was supported by the Cal-Fed Science Program) was preceded by a decade of collaboration with Felicity Hungtingford, Neil Metcalfe, and John Thorpe on Atlantic salmon and supported by a NATO travel grant, California Sea Grant, and the US National Marine Fisheries Service. Preparation of this paper was partially supported by NSF Grant EF-0924195. I thank Suzanne Alonzo, Charles Fisher, Holly Kindsvater, Susan Mangel, Bernard Roitberg, Jay Rosenheim, and Katriona Shea for comments on a previous version of the manuscript, Simone Vincenzi for a careful and thoughtful review of the penultimate version, and two anonymous referees for their comments on the submission. To a very real extent, my career was shaped by Morton Brussels and Donald Ginsberg, who taught that course in electricity and magnetism in academic year 1969–1970 and by my first scientific mentor, William W. Hay, who guided me through my first paper in which models were used to illuminate a scientific question (Mangel 1971).

References

- Bellman R (1952) On the theory of dynamic programming. *Proc Natl Acad Sci USA* 38:716–719
- Bellman R (1954) The theory of dynamic programming. *Bull Am Math Soc* 60:503–515
- Bellman R (1956) Dynamic programming and Lagrange multipliers. *Proc Am Acad Sci USA* 42:767–769
- Bellman R (1957) *Dynamic programming*. Princeton University Press, Princeton
- Cayuela H, Besnard A, Bonnaire E, Perrt H, Rivoalen J, Miaud C, Joly P (2014) To breed or not to breed: past reproductive status and environmental cues drive current breeding decisions in a long-lived amphibian. *Oecologia* (in press)
- Charnov EL, Skinner SW (1984) Evolution of host selection and clutch size in parasitoid wasps. *Fla Entomol* 67:5–21
- Charnov EL, Skinner SW (1985) Complementary approaches to understanding parasitoid oviposition decisions. *Environ Entomol* 14:383–391
- Clark CW, Mangel M (2000) *Dynamic state variable models in ecology. Methods and applications*. Oxford University Press, New York
- Courant R, Hilbert D (1962) *Methods of mathematical physics*. Wiley, New York
- Dayton PK, Sala E (2001) Natural history: the sense of wonder, creativity and progress in ecology. *Sci Marine* 65(Suppl 2):199–206
- Engelhard GH, Heino M (2005) Scale analysis suggests frequent skipping of the second reproductive season in Atlantic herring. *Biol Lett* 1:172–175
- Feynman RP, Leighton RB, Sands M (1963) *The Feynman lectures on physics. Mainly mechanics, radiation, and heat*. Addison-Wesley Publishing Company, Reading
- Giske J, Eliassen S, Fiksen O, Jakobsen PJ, Aksnes DL, Jørgensen C, Mangel M (2013) Effects of the emotion system on adaptive behavior. *Am Nat* 182:689–703
- Hilborn R, Mangel M (1997) *The ecological detective. Confronting models with data*. Princeton University Press, Princeton
- Houston AI, McNamara JM (1999) *Models for adaptive behavior. An approach based on state*. Cambridge University Press, Cambridge
- Houston A, Clark CW, McNamara JM, Mangel M (1988) Dynamic models in behavioural and evolutionary ecology. *Nature* 332:29–34
- Jørgensen C, Fiksen Ø (2006) State-dependent energy allocation in cod (*Gadus morhua*). *Can J Fish Aquat Sci* 63:186–199

- Jørgensen C, Ernande B, Fiksen Ø, Dieckmann U (2006) The logic of skipped spawning in fish. *Can J Fish Aquat Sci* 63:200–211
- Karlin S, Taylor HM (1981) A second course in stochastic processes. Academic Press, New York
- Kennedy J, Skjæraasen JE, Nash RDM, Slotte A, Geffen AJ, Kjesbu OS (2011) Evaluation of the frequency of skipped spawning in Norwegian spring-spawning herring. *J Sea Res* 65:327–332
- Krebs JR, Davies NB (1978) Behavioural ecology. An evolutionary approach. Blackwell Scientific Publications, Oxford
- Krener AJ (1979) A formal approach to stochastic integration and differential equations. *Stochastics* 3:105–125
- Lima SL, Dill LM (1990) Behavioral decisions made under the risk of predation: a review and prospectus. *Can J Zool* 68:619–640
- Mangel M (1971) A treatment of complex ions in sea water. *Marine Geol* 11:M24–26
- Mangel M (1987) Oviposition site selection and clutch size in insects. *J Math Biol* 25:1–22
- Mangel M (1992) Rate maximizing and state variable theories of diet selection. *Bull Math Biol* 54:413–422
- Mangel M (2006) The theoretical biologist's toolbox. Quantitative methods for ecology and evolutionary biology. Cambridge University Press, Cambridge
- Mangel M, Clark CW (1986) Towards a unified foraging theory. *Ecology* 67:1127–1138
- Mangel M, Clark CW (1988) Dynamic modeling in behavioral ecology. Princeton University Press, Princeton
- Mangel M, Ludwig D (1992) Definition and evaluation of behavioral and developmental programs. *Ann Rev Ecol Syst* 23:507–536
- Mangel M, Rosenheim JA, Adler FR (1994) Clutch size, offspring performance, and intergenerational fitness. *Behav Ecol* 5:412–417
- McNamara JM (2000) A classification of dynamic optimization problems in fluctuating environments. *Evol Ecol Res* 2:457–471
- McNamara JM, Houston AI (1986) The common currency for behavioral decisions. *Am Nat* 127:358–378
- Nonacs P, Dill LM (1990) Mortality risk vs. food quality trade-offs in a common currency: ant patch preferences. *Ecology* 71:1886–1892
- Railsback SF, Grimm V (2012) Agent-based and individual-based modeling. A practical introduction. Princeton University Press, Princeton
- Rashevsky N (1969) Mathematical biophysics. Physico-mathematics foundations of biology, vol 1 and 2. Dover Press, New York
- Rideout RM, Rose GA, Burton MPM (2005) Skipped spawning in female iteroparous fishes. *Fish Fish* 6:50–72
- Roitberg BD (2008) Gold Medal Address. *Bull Entomol Soc Can* 40:162–166
- Roitberg BD, Mangel M, Lalonde R, Roitberg CA, van Alphen JJM, Vet L (1992) Seasonal dynamic shifts in patch exploitation by parasitic wasps. *Behav Ecol* 3:156–165
- Roitberg BD, Sircom J, Roitberg CA, van Alphen JJM, Mangel M (1993) Life expectancy and reproduction. *Nature* 364:351
- Rosenheim JA, Rosen D (1991) Foraging and oviposition decisions in the parasitoid *Aphytis lingnanensis*: distinguishing the influences of egg load and experience. *J Anim Ecol* 60:873–893
- Rosenheim JA, Rosen D (1992) Influence of egg load and host size on host-feeding behaviour of the parasitoid *Aphytis lingnanensis*. *Ecol Entomol* 17:263–272
- Satterthwaite WH, Beakes MP, Collins E, Swank DR, Merz JE, Titus RG, Sogard SM, Mangel M (2009) Steelhead life history on California's Central Coast: insights from a state dependent model. *Trans Am Fish Soc* 132:532–548
- Satterthwaite WH, Beakes MP, Collins E, Swank DR, Merz JE, Titus RG, Sogard SM, Mangel M (2010) State-dependent life history models in a changing (and regulated) environment: steelhead in the California Central Valley. *Evol Appl* 3:221–243
- Skjæraasen JE, Nash RDM, Korsbrette K, Fonn M, Nilsen T, Kennedy J, Nedreaas KH, Thorsen A, Withames PR, Geffen AJ, Høiea H, Sigurd Kjesbu O (2012) Frequent skipped spawning in the worlds largest cod population. *Proc Natl Acad Sci USA* 109:8995–8999
- Sogard SM, Merz JE, Satterthwaite WH, Beakes MP, Swank DR, Collins EM, Titus RG, Mangel M (2012) Contrasts in habitat characteristics and life history patterns of *Onchorhynchus mykiss* in California's Central Coast and Central Valley. *Trans Am Fish Soc* 141:747–760
- Stokes DE (1997) Pasteur's quadrant: basic science and technological innovation. Brookings Institution, Washington

-
- Thorpe JE, Mangel M, Metcalfe NB, Huntingford FA (1998) Modelling the proximate basis of life history variation, with application to Atlantic salmon, *Salmo salar* L. *Evol Ecol* 121:581–600
- Yaragina NA (2010) Biological parameters of immature, ripening, and non-reproductive mature northeast Arctic cod in 1984–2006. *ICES J Marine Sci* 67:2033–2041