5 Interval Timing and Optimal Foraging

Melissa Bateson

CONTENTS

5.1 INTRODUCTION

The natural world is full of temporal regularities, and it makes sense that animals should have evolved clocks that permit them to maximize their fitness by exploiting these regularities. However, proving that animals use internal clocks to schedule their behavior is rarely possible in natural environments. To provide proof that an animal is timing, it is necessary to eliminate any external cues to the passage of time that the animal might be using in place of an internal clock (Killeen et al.,

0-8493-1109-8/03/$0.00+$1.50

1997). Because this is usually difficult in natural environments, data showing that animals can time intervals come from the controlled conditions of the laboratory, where the cues available to the subjects and the temporal properties of their experience can easily be manipulated. Such laboratory experiments have traditionally been the domain of operant psychologists who typically restrict their studies to the behavior of laboratory-reared rats and pigeons tested in Skinner boxes on various schedules of reinforcement. Due to the lack of a clear connection between the behavior of rats and pigeons in the artificial environment of the Skinner box and the problems faced by wild animals in their natural environments, research on interval timing has so far failed to attract widespread interest from ethologists and behavioral ecologists who usually focus on understanding behavior patterns initially identified in wild animals. Evidence for interval timing in animals has mostly been published in the psychological literature, and it is psychologists that have been responsible for setting the agenda in research on interval timing.

As a consequence of the domination of the interval timing literature by psychologists, the focus in most timing research has been on describing the psychophysics of interval timing, with the ultimate goal of elucidating the cognitive and neural mechanisms underlying the interval timing clock (e.g., Hinton and Meck, 1997; Matell and Meck, 2000; Paule et al., 1999). For example, psychologists have paid particular attention to inaccuracy and imprecision in interval timing on the grounds that imperfections in the system can be particularly important for revealing the underlying mechanisms (e.g., Gibbon and Church, 1984). It is taken for granted that the ability to time intervals is useful to animals, and the interval timing clock has consequently been conceived of as a general-purpose stopwatch-like timepiece. A result of this approach has been that questions regarding the evolution and current function of interval timing have remained largely unasked.

In contrast to the psychological approach, ethologists believe that no account of behavior is complete until both its proximate mechanisms and its ultimate evolutionary functions are understood (Tinbergen, 1963). Skinner (1989) stated that “by looking at how a clock is built, we can explain [how] it keeps good time, but not why keeping time is important.” Although Skinner was correct in recognizing that questions about mechanism and function are logically distinct, the strength of Tinbergen’s ethological approach lies in the belief that finding the answer to one type of question will often provide valuable insights into the answer to the other. Thus, although logically it is not necessary to know why the clock has evolved in order to understand how it works, there is reason to believe that considering both questions simultaneously could bring considerable benefits of understanding.

There are a number of reasons why it might be beneficial to study function and mechanism simultaneously. First, it is generally much easier to understand how a mechanism works if you know exactly what it is designed to do. By analogy, deciphering someone else’s computer code is always facilitated if you know precisely what the function of the code is. Although we may feel we understand what clocks are for, it remains the case that a clock that functions to measure the intervals between successive prey captures in a foraging starling may have different design requirements than a clock that functions to measure the rate of the mating display in a male sage grouse. Second, interval timing clocks are likely to have evolved in response to specific types of temporal regularities that confront animals in their natural environments. As a result, the clock may behave very differently when it is probed with the natural stimuli that it has evolved to respond to than when it is probed with unnatural stimuli that are outside the range of variation it encountered during its evolutionary history. In other words, modeling the biological clock as a flexible stopwatch may be misleading; the biological clock may behave differently depending on what it is asked to do. Third, psychologists have described several ways in which the performance of the interval timing clock departs from perfect accuracy and precision. In order to understand the significance of these imperfections for the adaptive behavior of animals, we need to know whether the imperfections of the clock are true constraints of its mechanism or whether perhaps they are artifacts of studying the clock under conditions for which it was not designed. By analogy, if you rev a car on a surface such as ice or sand, the pair of wheels that drive the car will spin. This gives you some clear information about how the car works: you now know whether it has front- or rear-wheel drive. However, the behavior of the car on the sand is very different from what you would observe if you were to rev the same car on a road. Thus, studying timing in the lab may provide very useful insights into the mechanisms underlying the clock; however, it may be misleading if your aim is to understand the role of timing in the generation of adaptive behavior. The evolutionary process has resulted in a complex relationship between function and mechanism in biological systems, and it is only by considering both questions simultaneously that we can ever hope to understand animal behavior fully.

My aim for this chapter is to take an ethological approach to interval timing and attempt to integrate mechanistic and evolutionary approaches. I begin by asking the functional question of why animals need to be able to time intervals, and I present evidence suggesting that interval timing is likely to have a central role in foraging behavior. Next I introduce optimal foraging theory and show exactly how temporal information is important in making optimal foraging decisions. An important component of all optimal foraging models is the constraints that are assumed, and I go on to consider timing as a potential constraint. The most sophisticated attempts to integrate models of timing and foraging have come from studies that focus on scalar timing theory. I give a brief description of the basic scalar timing model and review how it has been adapted to model a range of foraging problems. Throughout, my aim is to highlight the benefits of understanding that have resulted from the combination of scalar timing with optimal foraging.

5.2 INTERVAL TIMING AND FORAGING

So why do animals need the ability to time intervals? Although timing seems a generally useful ability, are there specific behavioral problems faced by animals that are particularly likely to have selected for interval timing abilities? In trying to answer these questions, I am going to argue that several lines of evidence point to the likelihood that interval timing is of major importance in foraging behavior (see Hills, this volume).

Interval timing has been identified in the majority of vertebrate species in which the ability has been investigated (Lejeune and Wearden, 1991), and it is probably safe to assume that the ability is universal in the vertebrates (Bateson, 2001). This implies that interval timing is an evolutionarily ancient ability, and in searching for behavioral problems that might have specifically selected for timing, we need to identify very general problems that are faced by all vertebrate animals and are likely to have also been faced by their common ancestors. All animals need to find food, making foraging behavior a potential candidate for a general problem that might have selected for the ability to time.

Warm-blooded vertebrates such as small mammals and birds have relatively high metabolic rates, the maintenance of which requires large quantities of food on a regular basis. This need is assumed to have imposed a strong selective pressure to produce efficient foraging behavior. Birds in particular have been at the forefront of studies of optimal foraging because the high energetic demands of flight have meant that birds need to ingest particularly large quantities of food of high nutritive value. A small bird may need to eat its body weight in food per day. A blue tit (*Parus caeruleus*) weighing 11 g needs 11 kcal per day in winter, which is equivalent to around 300 small insects, and a rufous hummingbird (*Selasphorus rufus*) weighing 3.5 g will visit a nectar feeder approximately every 10 to 15 min from dawn until dusk. These high rates of foraging coupled with the fact that most birds forage only during the hours of daylight make birds attractive subjects for studies of foraging behavior.

Arthropods are the most common food consumed by small mammals and birds, and they bring with them various interesting foraging problems because the distribution of arthropod prey in the environment is usually both spatially and temporally patchy. For example, for much of the year European starlings forage on leatherjackets, the larvae of tipulid flies. Leatherjackets are hidden beneath the surface of the soil and occur in patches. Starlings search for leatherjackets by probing the soil with their bills. Given that it is impossible to visually assess the density of leatherjackets in a particular location, starlings rely on their cognitive abilities to form estimates of the rates of intake they have experienced in different locations. These estimates can then be used as a basis for making future foraging decisions.

It is possible to analyze all foraging behavior in terms of its costs and benefits to the forager. Finding, consuming, and digesting food all have both energetic and time costs associated with them, because time and energy spent foraging are time and energy taken away from other fitness-promoting activities, such as looking out for predators and reproducing. We therefore expect natural selection to have honed foraging decisions so as to optimize the trade-off between costs and benefits, and thus maximize the lifetime survival and reproductive success of the forager (e.g., Stephens and Krebs, 1986). Because the costs associated with foraging involve the length of time taken, it is likely that selection on foraging decisions has involved selection on the ability to measure these costs accurately.

The final piece of evidence linking timing with foraging comes from the observation that the majority of the comparative evidence for interval timing comes from animals performing on fixed-interval (FI) schedules of reinforcement in which food is used as the reinforcer. In a typical free-operant FI schedule, reinforcement, usually the delivery of a small amount of food, is contingent on a response made by the subject after some fixed period of time has elapsed. The interfood interval serves as the only discriminative stimulus, and the interval requirement is reset after each food reinforcement is delivered. The optimal strategy in a subject trying to maximize the frequency with which it receives food while minimizing the number of responses it has to make when faced with such a schedule is to time the fixed interval and make a single response as soon as the interval has elapsed. Although well-trained subjects never achieve this optimal strategy, they do show a postreinforcement pause that averages about two thirds of the FI value, after which they start responding at a high rate until food is delivered (e.g., Schneider, 1969). Thus, we have good evidence that animals are able to time interfood intervals and also that the delivery of food can serve to reset the animal's interval timer (Matell and Meck, 1999).

In summary, therefore, we have established that (1) all vertebrates need to forage, (2) foraging behavior is likely to have been under strong selective pressure to increase efficiency, (3) efficient foraging involves making decisions that involve timing, and (4) animals can time intervals between food deliveries and that food can reset the clock. Taken together, I suggest that the above evidence points strongly toward the possibility that interval timers may play a major role in natural foraging behavior. Of course, none of the above evidence proves that interval timers initially evolved for the purposes of foraging, or that interval timers are used solely for the purposes of foraging. However, the likely fitness consequences of inefficient foraging do suggest that the selection pressures to improve the efficiency of foraging are very likely to have been an important force in the evolution of the interval timing clock.

5.3 OPTIMAL FORAGING THEORY

Optimal foraging theory is the branch of behavioral ecology that seeks to understand how natural selection has shaped foraging behavior. The general strategy adopted is to build models of how animals should forage given various assumptions about the system, and then compare the predictions of these models with the behavior of real animals.

Classical optimal foraging models involve three kinds of assumptions: (1) those regarding the foraging decision being analyzed, (2) those regarding the currency the forager is maximizing, and (3) those regarding the constraints operating on the system (Stephens and Krebs, 1986).

In many classical optimal foraging models the currency that foraging animals are assumed to be maximizing is their long-term net rate of energy intake (Stephens and Krebs, 1986), where long-term rate is defined as the net energy intake divided by the total time spent acquiring this energy. Rate is a proximate currency that is assumed to relate closely to Darwinian fitness if it is maximized over the lifetime of the forager, because an animal that maximizes its rate of energy intake will achieve the greatest amount of energy for use in maintenance, growth, and reproduction in the least possible time, and time not spent foraging is time available for other fitness-promoting activities, such as looking out for predators and reproduction. Given that the computation of rate involves forming an estimate of the time spent foraging, interval timing is likely to be involved in many foraging decisions.

In the following sections I shall consider various foraging problems in which the ability to time intervals is potentially crucial to arriving at the optimal solution. In each case I will describe the problem faced by the forager, outline the solution to the problem that maximizes the rate of energy intake, and describe examples in which the behavior of animals has been shown to approximate the optimal foraging solution.

5.3.1 WHEN TO RETURN TO A RENEWING FOOD SOURCE

Some food sources gradually renew following depletion or depression by a forager. If this renewal process is temporally predictable, then such food sources provide a natural equivalent of a fixed-interval schedule of reinforcement. A forager that returns early to the food source will obtain less food than the maximum possible if it had delayed its return until the source had completely renewed. A forager that returns late to the food source will also obtain less than the maximum rate of food intake possible from that food source because it has waited longer than was necessary to obtain the maximum amount of food, and in situations where there is competition from other foragers, there is the added risk that the food might be lost to a competitor if the forager fails to claim it as soon as it is available. Thus, a forager that can learn the temporal predictability of the food source will be at an advantage over one that cannot, because it can schedule its visits to the source to correspond with the times at which maximum food is available and thus maximize its rate of food intake from the source.

There are several possible natural examples of renewing food sources with predictable temporal properties. For example, in some of the flower species used by nectar-feeding hummingbirds the amount of nectar available in the flower increases predictably and monotonically as a function of the time elapsed since the flower was last visited until the flower has fully refilled. In a field study of long-tailed hermit hummingbirds (*Phaethonis superciliosus*), Gill (1988) studied the responses of wild hummingbirds to 10- and 15-min FI schedules on an artificial nectar feeder. Rather than filling the feeder gradually, as would occur in a natural flower, he refilled the feeder either 10 or 15 min after the feeder had been emptied by a bird, as in a conventional FI schedule. He showed that under conditions of nearly exclusive use of a feeder by a single individual when the risk of competitive loss was low, return intervals increased to longer than the FI, thereby maximizing the probability of obtaining nectar on a visit.

Another possible example of a renewable food source with temporally predictable properties is provided by the amphipod crustacean *Corophium volutator*. *Corophium* species are the major food source of redshanks (*Tringa totanus*), wading birds that forage on tidal mudflats. The *Corophium* species live in burrows and retreat down their burrows when a redshank walks over the sediment surface to a depth where they are inaccessible to the redshanks (Goss-Custard, 1970). This is known as prey depression, and measurements of the feeding rates of redshanks have shown that it takes around 10 min for feeding rates to recover following the previous visit by a bird (Yates et al., 2000). This interval is presumed to correspond to the time it takes the *Corophium* species to resurface from their burrows. There is no information on whether redshanks are able to time this interval, but it would clearly be adaptive for them to do so, because if they return to an area within 10 min of previously feeding there, they will experience a reduced rate of prey capture.

5.3.2 WHEN TO LEAVE A DEPLETING PATCH

A decision commonly faced by foraging animals is that of when to leave a patch of food that is depleting and move on to a new patch. Because the optimal response to depleting patches varies according to whether the patch depletes gradually or suddenly, I will consider these two scenarios separately.

5.3.2.1 Sudden Patch Exhaustion

The availability of food in a patch may suddenly drop to zero. Faced with this situation, a forager needs to detect the change in status of the patch and move on to a new patch, because a forager that remains in an empty patch will be wasting time searching for food that is not there. Sometimes there may be visual or auditory cues available to the forager to indicate that a patch is exhausted, but if there are no such cues, a forager might be able to use its internal clock to detect patch exhaustion. Consider the situation in which there is a fixed intercapture interval and constant probability of sudden patch exhaustion after each prey encounter or after arrival in the patch. If the forager can learn the interfood interval, then it can use this knowledge to detect patch exhaustion. The optimal departure rule is to leave the patch as soon as exhaustion is detected, namely, as soon as the period since the last prey capture exceeds the normal intercapture interval. This optimal departure rule is independent of the time taken to travel to a new patch (Kacelnik et al., 1990).

The foraging of spotted flycatchers (*Muscicapa striata*) provides a possible example of sudden patch exhaustion (Davies, 1977). These birds hunt from a fixed perching site, making periodic forays to catch flying insects that have entered their range. Usually following a capture, the bird returns to the same or a nearby perch, but occasionally it will leave the site and travel to a new perch to continue hunting. Davies was interested in establishing what caused a bird to abandon a foraging site and move on to the next. He found that the birds would move sites when they had waited 1.5 times the average intercapture interval without making a foray, and on the basis of this finding, he suggested that the birds might be using a rule that involved estimating the average intercapture interval and moving on to a new perch when no insect had appeared within a given multiple of this time (specifically, 1.5 times). Such a moving-on rule might be optimal in an environment in which patches of prey deplete suddenly and completely, as might be the case, for instance, if a swarm of flies moves away from the immediate vicinity of the bird's perch, and where there are no external cues to indicate that the patch is now empty. Unfortunately, with the available data, it is impossible to prove that the flycatchers are really using such a time-based moving-on rule (Kacelnik et al., 1990). Proof that the birds are using an internal clock to judge their departure from a site would require eliminating the possibility that the birds do not just move on when they can no longer detect any prey items within their range. Additionally, proof that the birds wait for a fixed multiple of the intercapture interval before moving on would require manipulating intercapture intervals and demonstrating the predicted effects on moving-on times.

5.3.2.2 Gradual Depletion

The rate of food intake within a patch will often be a decelerating function of the amount of time the forager has already spent foraging in the patch. In this situation the forager is faced with the decision about when to leave the current patch and pay the cost of traveling to a new patch. The marginal value theorem (MVT) (Charnov, 1976) describes the behavior that maximizes the long-term rate of energy intake in such a situation. Long-term rate is maximized if the forager stays in each patch until its instantaneous rate within the patch falls below the background rate of gain in the environment as a whole. This background rate will be affected by the travel time between patches, and as the travel time between patches increases, the optimal patch residence time also increases. Because rates are equal to amounts per unit of time, computation of rate requires the ability to time intervals. An important prediction of the MVT is that although patch residence time is predicted to be sensitive to the average travel time, it should not be affected by variance in travel time.

A number of studies have tested the prediction that patch residence time should be positively related to the travel time between patches. Perhaps the most well known is the study by Kacelnik (1984) on European starlings (*Sturnus vulgaris*). During the breeding season starlings make regular forays from their nest to collect food for their chicks. As required by the MVT within each foraging bout, the starling suffers a decelerating rate of food acquisition. This occurs due to the starling's method of foraging, whereby it probes the ground to look for invertebrates, such as leatherjackets, hidden beneath the surface of the soil. As the starling's bill fills up with prey, the bird becomes progressively less efficient at probing the ground for further prey and its rate of prey acquisition declines. Kacelnik tested the MVT prediction that travel time should affect patch residence time by setting up feeding stations for starlings at different distances from their nests. He simulated the loading curve by delivering worms to the birds at progressively greater intervals the longer they stayed at the feeding station. He was able to show that the number of worms the starling collected before returning to its nest increased as the travel distance to the feeder increased. The observed behavior was approximately as predicted if the bird were maximizing the rate at which it delivered worms to its chicks.

5.3.3 HOW TO RESPOND TO VARIABILITY

Many natural food sources are variable either in the exact amount of food they provide or in the time associated with finding or extracting the food. For example, consider a forager faced with one feeding option that it knows will yield food after 5 min of searching vs. another feeding option that it knows will yield the same amount of food after the same average searching time, but the actual time taken to find the food could vary between 1 and 9 min. A rate-maximizing forager should be indifferent to such variability because the computation of the long-term rate of energy intake involves averaging the amounts and times associated with each food source, with the result that both food sources are perceived to be of equal value. However, there are circumstances where the long-term rate of energy intake may not be the currency that correlates best with fitness. Consider a small bird in winter faced with one more foraging decision before the rapidly approaching night. It is vitally important for this bird to achieve a threshold level of energy reserves before dusk in order to survive the long cold night. In this situation, it can be optimal for the forager to pay attention to the variance, or risk, as it is called in the foraging literature, in its food sources. If the bird has no chance of meeting the required threshold for survival in time by choosing the fixed option (i.e., it is on a negative energy budget), then its only chance of survival is to choose the risky option in the hope that it will be lucky and find food quickly. Conversely, if the fixed option will easily take the bird above threshold (i.e., it is on a positive energy budget), then it would be foolish to choose the variable option and risk not getting its final prey item before nightfall. These arguments are summarized in the daily energy budget rule that states that a forager on a positive energy budget should be risk averse, while one on a negative budget should be risk prone.

There is a large literature showing that animals are sensitive to risk in both the amount of food and in the delay associated with obtaining food (Kacelnik and Bateson, 1996; Bateson and Kacelnik, 1998). Unfortunately, though, there is little good empirical support for functional explanations for risk sensitivity such as the energy budget rule (but for a beautiful demonstration of the energy budget rule in yellow-eyed juncos, see Caraco et al., 1990). The overall pattern in the literature is that animals tend to be risk averse when there is variability in amount of reward, but risk prone when variability is in delay to reward. This pattern is not readily explained by any of the optimal foraging models, and there has been an ongoing discussion in the literature about whether risk sensitivity is an adaptive response to environmental variability or instead is an artifact of the cognitive mechanisms animals use to assess quantities such as amount, time, and rate (e.g., Bateson and Kacelnik, 1995b; Reboreda and Kacelnik, 1991).

5.3.4 WHEN TO SAMPLE A CHANGING PATCH

The abundance of food in the environment is seldom stable: the availability or quality of particular feeding sites or prey types may change unpredictably over the course of time, perhaps caused by depletion or depression by other foragers or by changes in the weather. As a consequence, one of the problems faced by foraging animals is keeping track of the status of their food resources. Sometimes there may be environmental cues to prey availability, but in instances where there are no such cues, a forager will be forced to sample its environment periodically. Stephens (1987) modeled the problem of how a rate-maximizing forager should track a changing environment. He considered the situation where there are two patch types available to a forager, one that fluctuates between good and bad states and one that remains stable at a value between the two states of the fluctuating patch. The fluctuating patch has a constant probability of changing state, and Stephens assumed that the forager can determine the state of the fluctuating patch as the result of a single sample. Given these assumptions, Stephens asked how a rate-maximizing forager should sample the fluctuating environment. The theoretical results of his model show that an optimal forager should sample the fluctuating patch at regular intervals. The optimal sampling frequency depends on the ratio of two kinds of cost: the lost opportunity experienced in the stable patch when the forager samples the fluctuating patch, and missing the opportunity to forage in the fluctuating site when it is in its good state.

5.3.5 CONSTRAINTS IN OPTIMAL FORAGING MODELS

Traditionally, the kinds of constraints assumed in optimal foraging models have been simple physical constraints, such as the impossibility of simultaneously handling one prey item while searching for another, or the reduction in the rate of prey acquisition as the forager's bill fills up with prey. In general, potential psychological constraints imposed by limitations of an animal's information-processing capacities have received much less attention from optimal foraging theorists. There are several reasons for this neglect. First, theoreticians usually seek to keep models as simple and as general as possible. Second, foraging theorists justify neglecting psychological constraints on the grounds that natural selection should have provided animals with near-perfect psychological mechanisms if these mechanisms have a measurable impact on optimal performance. Finally, due to the lack of integration of the psychological and ecological literatures, behavioral ecologists are often simply unaware of the wealth of data that is available regarding animals' information-processing capacities.

The ignorance of behavioral ecologists of the psychological literature has resulted in some misguided early attempts by them to introduce psychological constraints, such as imperfect timing, into optimal foraging models. For example, Yoccoz et al. (1993) presented a general theoretical framework for understanding the effects of perceptual error on optimal foraging decisions. In order to do this, they assumed explicit relationships between the real energy content and handling time of each food item and what the forager estimates these quantities to be. They incorporated these assumptions into a model of optimal diet choice (Engen and Stenseth, 1984), which is based on the assumption that animals maximize their long-term rate of energy intake, and investigated the effects of adding perceptual error on the optimal solution to the model. In their model, Yoccoz et al. (1993) represent the actual energetic gains and times as random variables G and T, and the animal's perception of these random variables as X and Y, respectively. They construct X and Y from G and T by adding unbiased normally distributed errors such that $X = G + E_G$ and $Y = T + E_T$, where E_G and E_T are the errors in gain and time, respectively. This general strategy is completely sound; however, Yoccoz et al. (1993) assume that the variances of the errors E_G and E_T are independent of the actual gains and times, G and T. This latter assumption flies in the face of over a century's worth of psychophysics showing that the error in the perception of the magnitude of a stimulus is not independent of its magnitude, but rather increases with stimulus magnitude (Bateson and Kacelnik, 1995a), a relationship known as Weber's law (see Figure 5.1).

The value of Yoccoz et al.'s (1993) model is that it demonstrates that introducing perceptual error can result in constrained rate-maximizing solutions that differ substantially from naïve, unconstrained treatments of the same problem. However, if such models are to be of use in understanding the behavior of real animals, then they need to include more realistic assumptions regarding the nature of perceptual constraints.

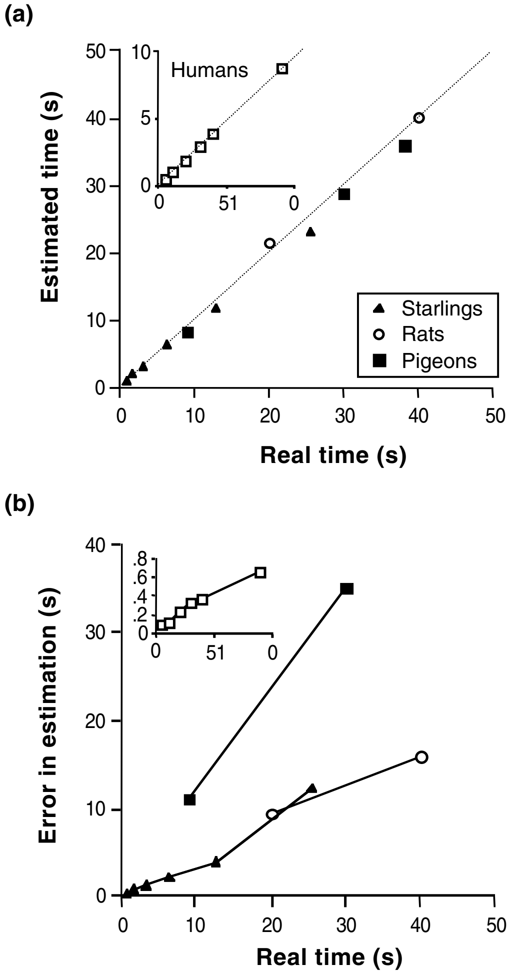

FIGURE 5.1 Examples of (a) estimated time, and (b) error in estimation of time, both vs. actual time from a variety of species. The data are taken from experiments in which subjects were required to reproduce time intervals. (Reproduced from Bateson, M. and Kacelnik, A., *Anim. Behav.*, 50, 431–443, 1995a. Copyright © 1995 by the Association for the Study of Animal Behavior. With permission.)

In the next section I introduce scalar timing theory, a sophisticated information-processing model of interval timing that has formed the basis of some of the most successful attempts to integrate realistic psychological constraints into optimal foraging models.

5.4 SCALAR TIMING THEORY

Scalar timing theory (see Church, this volume; Gibbon et al., 1984) is an information-processing account of interval timing that grew out of scalar expectancy theory

(SET) (Gibbon, 1977). All models of interval timing require three basic functions: a clock function that measures elapsed time by converting it to some physical representation, a memory function in which a recorded time interval can be represented and stored, and a decision function that uses output from the clock and the memory components to control behavior (Church, 1997). In scalar timing theory these different functions are embodied in discrete components described as follows.

The clock subsystem consists of a pacemaker, a switch and an accumulator. Time measurement is achieved by collecting pulses from the pacemaker in the accumulator when the switch is closed. The pacemaker is assumed to continuously emit pulses at a rate Λ. When the switch is closed in response to an external signal to begin timing, pulses are transmitted to the accumulator, where they accumulate until the switch is opened again. The value in the accumulator, m_t, thus represents the amount of time that has elapsed since the switch was closed and timing began. This estimate of real time grows linearly with real time, t, such that

$$m_t = \Lambda(t - T_0)$$

where T_0 is the mean latency between the external start signal and the beginning of time accumulation.

The current estimate of elapsed time, m_t, can be transferred either to working memory for immediate use or to reference memory, where it can be stored for future reference. In a fixed-interval schedule the usual cue for transferring a value to reference memory is the delivery of food. When food is delivered (i.e., t = FI, the value of the fixed interval), the scalar timing theory assumes that the value of m_t, which we will now refer to as m_{FI}, is transferred to reference memory as

$$m_{FI}^{*} = k * m_{FI}$$

where k^* is a translation constant that is assumed to vary between trials. An important assumption of scalar timing theory is that m_{FI} is assumed to be represented in reference memory as a distribution of the various values of m_{FI}^{*} transferred to it on different trials. The form of this distribution is determined by the way in which k^* varies between trials. For the purposes of modeling timing performance on a fixed-interval procedure, the mean of k^* is usually assumed to be unity such that the mean estimate in memory is equal to the mean estimate of the current time at which reinforcement occurs. Once an animal is fully trained on a fixed-interval schedule, it is assumed to have built up a reference memory representation of the interval that has a mean of m_{FI} and a standard deviation proportional to this mean. If there is variance in the real time between reinforcements, as is the case, for instance, on a variable-interval schedule, then the memory representation formed will be equivalent to the sum of the distributions that would be formed for each of the component intervals in the variable mixture. Due to the fact that the standard deviation of the representation of an interval grows with the value of the interval, the memory distributions of variable intervals will be asymmetrical and skewed to the right. The different memory representations resulting from a range of different mixtures of intervals are illustrated in Figure 5.2.

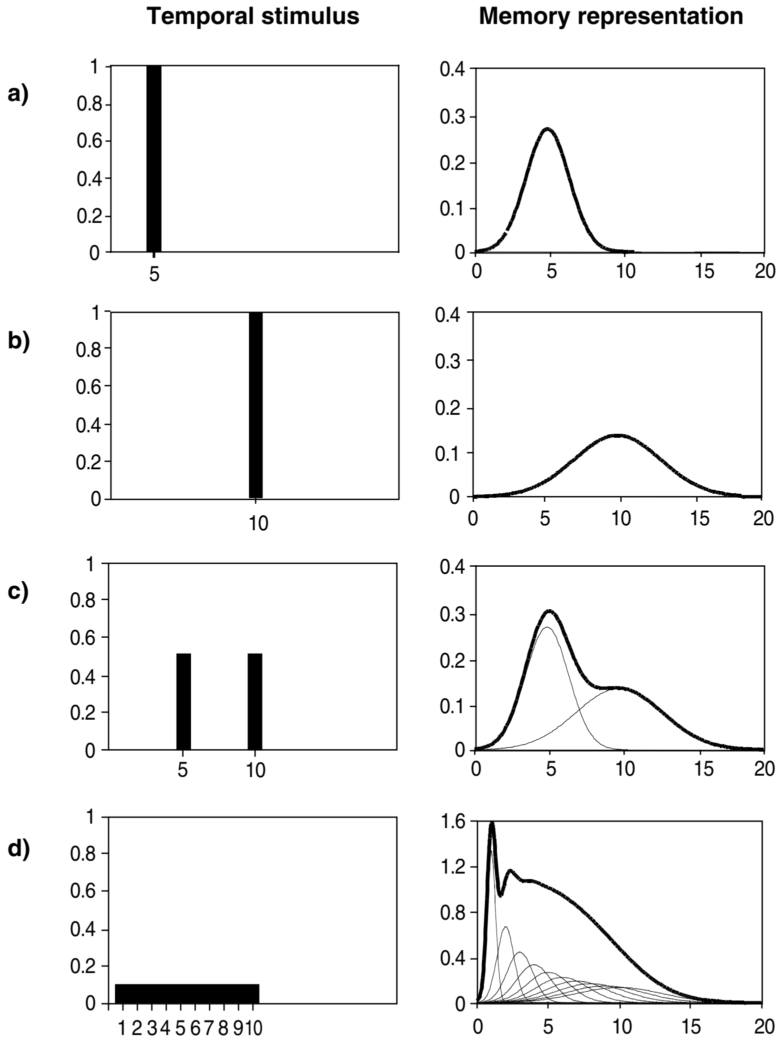

FIGURE 5.2 Examples of how a range of different temporal stimuli are represented in memory according to scalar timing theory. Temporal stimuli are shown in the left-hand panels and their associated memory representations in the right-hand panels. In all panels, probability is shown on the y-axis and the duration of the interval on the x-axis. In generating the memory representations, a coefficient of variation of 30% in the perception of the interval was assumed. (a) A fixed interval of 5 sec. (b) A fixed interval of 10 sec. (c) A variable interval that has an equal probability of being 5 or 10 sec. (d) A variable interval that has an equal probability of being 1, 2, 3, 4, 5, 6, 7, 8, 9, or 10 sec. In panels (c) and (d) the overall memory representation of the mixture is indicated by the bold line, and the representations of the individual intervals of which it is the sum are indicated by the thinner lines. Note that due to the constant coefficient of variation, the longer an interval, the less precisely it is represented; i.e., the scalar property applies. As a result of the scalar property, the memory representations for variable intervals are asymmetrical and skewed to the right.

In scalar timing theory the decision subsystem is assumed to receive input from both working and reference memories and to compare these inputs, usually using a ratio rule, to produce a behavioral output. For the purposes of modeling fixed-interval timing performance, at the beginning of each trial the decision subsystem is assumed to receive a single input from reference memory and continuously compare this with the current value of elapsed time, m_t, in working memory. Importantly, the value of m^*_{FI} used in this comparison process is assumed to be a single random sample from the distribution of m^*_{FI} stored in reference memory. The rule used to compare the values of m_t and m^*_{FI} is the discrimination ratio

$$\frac{m_t - m^*_{FI}}{m^*_{FI}}$$

When this ratio becomes greater than or equal to a threshold value, *b*, a decision is made to start responding. The value of *b* can also vary between trials, but for fixed-interval schedules it is usually assumed to have a mean value of approximately zero. Variation between trials in either *k** or *b* will induce the commonly observed relationship between the mean and variance of timing data known as the scalar property, whereby the coefficient of variation (i.e., the standard deviation or mean) is constant. Thus, the above scalar timing model is capable of producing the basic features of fixed-interval timing performance. The model is summarized in Figure 5.3.

One of the attractions of the scalar timing model is that the three components — clock, memory, and decision — are clearly modular, and the representation of information in memory is clearly separated from the way in which information is used in decisions. This modularity makes the scalar timing model very flexible and, as a result, particularly tractable for modeling a range of different foraging problems. The details of both the memory and decision components of the scalar timing model are altered depending on the specific task being modeled, as I shall demonstrate in the following section.

5.5 APPLICATIONS OF SCALAR TIMING TO FORAGING PROBLEMS

My aim in this section is to review examples where the scalar timing model has successfully been applied to some of the optimal foraging problems previously described. The general strategy has been to incorporate the perceptual error produced by the scalar timing model as an explicit psychological constraint in conventional optimal foraging models, and to explore whether the modified models do a better job of explaining observed behavior than their unconstrained counterparts. I shall describe how the basic scalar timing model is modified to cope with different foraging problems and highlight the benefits in understanding that have resulted. Although the following studies all tackle foraging problems that have been identified in animals in their natural habitats, the experiments described all involve analogs of these foraging problems translated to the operant laboratory. In attempting to replicate the natural foraging problems, the schedules developed for study in the lab are

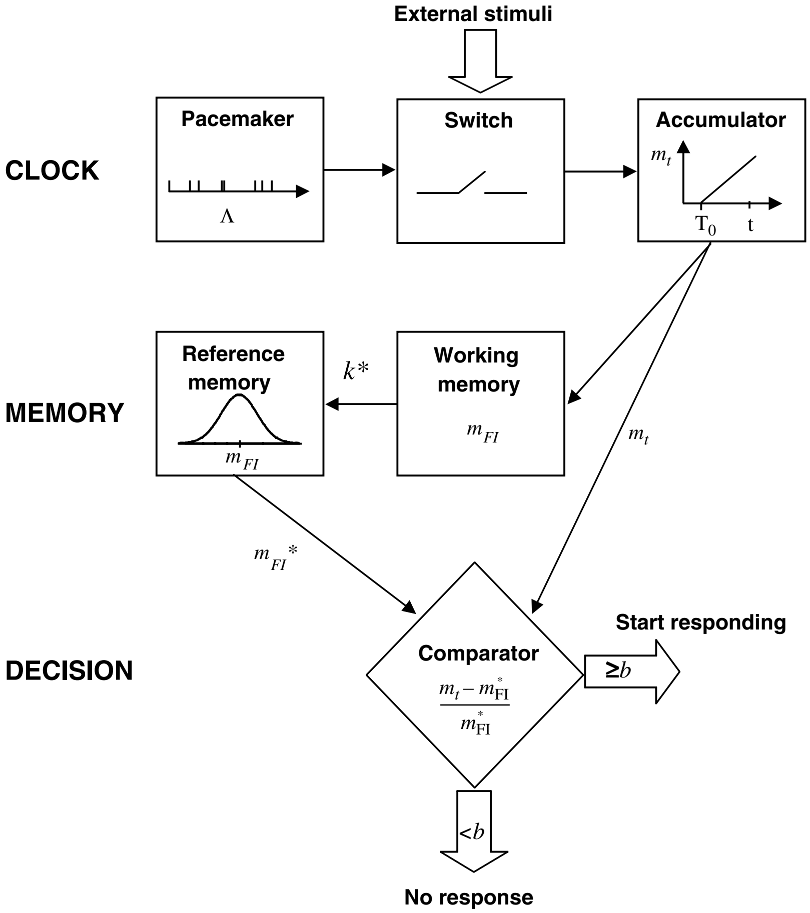

FIGURE 5.3 Diagrammatic representation of the basic scalar timing model. This version of the model predicts when a subject will start responding on a standard FI schedule. In a free-operant FI schedule the external stimulus that resets the clock and begins timing is the occurrence of food. Alternatively, in a discrete trial version of an FI schedule the external stimulus would be the stimulus that indicates the beginning of a trial, which is typically a tone or light. See the text for details of the model.

often more complex than the types of reinforcement schedules traditionally studied by operant psychologists.

5.5.1 SCALAR TIMING AND PATCH DEPARTURE: SUDDEN PATCH EXHAUSTION

Brunner and Kacelnik (Brunner et al., 1992, 1996; Kacelnik et al., 1990; Kacelnik and Brunner, 2002) set out to analyze the timing problem suggested by Davies' spotted flycatchers using European starlings foraging on an operant analog of the patch departure scenario in the lab. Their design simulated an environment in which food was distributed in patches. Each patch contained a random number of prey items (N = 0 to 4) that were delivered according to a fixed-interval schedule until the patch ended with sudden depletion. The time elapsed since the last prey item was the only cue the bird had to detect patch depletion. Once the patch had depleted, the bird could leave the patch and travel to a new patch by flying between two perches. As described above, the optimal patch departure rule, given perfect timing, is to abandon the patch as soon as exhaustion is detected, namely, when a prey is not encountered after waiting for the programmed fixed interval.

Brunner and Kacelnik tested how the patch departure of starlings was affected by the value of the fixed-interval schedule in the patch by examining the behavior of the birds tested with six different values of the fixed interval, ranging between 0.8 and 25.6 sec. As predicted by the unconstrained optimality model, the giving-in time at which a bird stopped attempting to obtain food from a patch in each trial (defined as the last peck in the patch) increased linearly as the fixed interval increased. However, in contrast to the predictions of the unconstrained model, the line relating the value of the fixed interval to the giving-in time had a slope of 1.49 rather than 1.00. Thus the starlings waited for a fixed proportion of the fixed interval before abandoning the patch. This result is particularly interesting in the light of data showing that the starlings knew accurately when food should have been expected. An analysis of the patterns of pecking in the final, unreinforced interval of each patch revealed that the birds showed a peak of pecks centered accurately on the value of the fixed interval: the line relating the value of the fixed interval to the peak time had a slope that was not significantly different from 1.00. Therefore, despite apparently knowing accurately when food should have been delivered, the birds still chose to wait 1.49 times the usual interfood interval before giving in; i.e., a multiplicative bias was introduced by the decision mechanism that controls the bird's decision to give up searching for more food in the patch.

How can we explain the multiplicative bias in the giving-in rule? It is possible to answer this question either by considering the functional implications of adopting different giving-in rules or by considering the timing processes responsible for the decision. We will start with the functional approach. If the birds knew accurately when food should have been delivered in the patch, why didn't they give in and leave the patch as soon as it did not appear? The answer to this question is to be found in the behavioral variability of the birds. Although the peak times of the birds show perfect accuracy, the standard deviations of the pecking functions are linearly related to the fixed interval: in other words, the timing functions of the birds display the scalar property evident in all timing data. The giving-in times also display the scalar property because their interquartile range also increased linearly with the fixed interval. In order to understand why it is optimal to introduce a multiplicative bias in the giving-in rule, we need to consider the potential costs of giving in early (before the patch is exhausted) vs. the costs of giving in late, as was observed. It is easy to show that from the perspective of rate of energy intake, it is much more costly to give in too early than too late. This asymmetry occurs because an animal that gives in early fails to get the last food item available in the patch, and this has a much bigger impact on the rate of energy intake than the relatively small time cost imposed by giving in late. Thus we can predict that if imprecision in timing is a constraint,

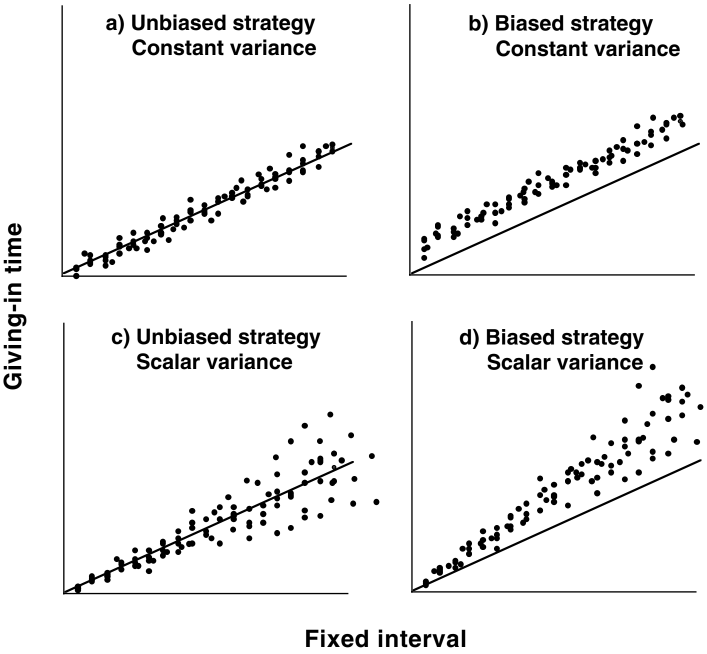

FIGURE 5.4 Variance in timing behavior and optimal giving-in time strategies. In each panel the solid line has a slope of 1.00 and indicates the optimal giving-in time predicted if timing is both perfectly accurate and precise. The top two panels (a and b) show constant variability in giving-in time, and the lower two panels (c and d) proportional or scalar variability in giving-in time. When timing is imprecise, the optimal strategy is biased in order to minimize the probability of leaving the patch before it is exhausted (panels b and d). When variability is constant, the optimal strategy is: giving-in time = fixed interval + bias (panel b); however, when variability is scalar, the optimal strategy is: giving-in time = fixed interval*bias (panel d). Brunner et al.'s (1992) starlings displayed a multiplicative bias as in panel (d) consistent with their also displaying the scalar property in their estimates of when food should have occurred. (Redrawn from Brunner, D., Kacelnik, A., and Gibbon, J., *Anim. Behav.*, 44, 597–613, 1992. Copyright © 1992 by the Association for the Study of Animal Behavior. With permission.)

then giving-in times should be biased to be longer than the fixed interval to guard against giving in too early. If, as the data revealed, the precision of timing is linearly related to the fixed interval (i.e., the scalar property applies), then the bias needs to be multiplicative, as opposed to a constant, in order to reduce the probability of giving in too early. These arguments are presented graphically in Figure 5.4.

It is also possible to analyze the behavior of the starlings from the mechanistic perspective of scalar timing theory. Scalar timing theory can be used to explain both the pecking patterns of the bird in the final unreinforced interval of the patch and the giving-in times. Both of these behavioral measures can be analyzed using the version of the scalar timing model proposed to deal with fixed-interval performance described above. The final unrewarded interval in the patch is exactly analogous to a probe trial from the peak-interval procedure. In order to use the scalar timing theory to model this performance, only a small modification to the decision mechanism is needed. Because birds start pecking at a high rate sometime prior to the FI and then stop pecking at this high rate sometime after the FI, the decision rule needs to produce both the start and stop of the high rate responding. This can be achieved by modifying the comparison made to

$$\left| \frac{m_t - m_{FI}^*}{m_{FI}^*} \right|$$

The starling responds at the high rate when this absolute ratio is greater than or equal to *b*. When applied over several trials, this model will result in pecking functions that are centered on the value of the fixed interval with a standard deviation proportional to the fixed interval, exactly as observed in the starlings.

The giving-in times can be modeled with another small modification of the basic model. In this case, we assume that the birds are using the same reference memory of reinforcement times used above to generate the pecking patterns, but a different decision rule. The bird leaves the patch at time *g* such that

$$\frac{m_g - m_{FI}^*}{m_{FI}^*} \geq b_g$$

where m_g is the perception of the current time, *g*, and b_g is a new threshold that is larger than *b*. Before Brunner et al.'s application of scalar timing theory to an optimal foraging problem, it had previously been assumed that thresholds were fixed (e.g., Gibbon 1977); however, one of the most important outcomes of Brunner et al.'s experiment is the suggestion that in fact a threshold such as b_g may be optimized by natural selection to maximize the rate of energy intake of the birds given the constraint of their imprecise timing mechanisms. The predictions of such an optimality approach are that the threshold, b_g, and consequently the giving-in times should be bigger when the cost associated with travel or the energy gain associated with the reinforcement are increased.

This study provides a beautiful example of the benefits of the ethological approach of integrating mechanistic and evolutionary analyses of the same problem. Without the scalar timing model, we would not have understood why the starlings' giving-in times are a fixed proportion of the interval between food items. Without the evolutionary approach, we would not have understood that the biases assumed in scalar timing are not arbitrarily chosen, but may actually be the outcome of an evolutionary optimization of the costs of leaving the patch too early and too late.

5.5.2 SCALAR TIMING AND PATCH DEPARTURE: THE MARGINAL VALUE THEOREM

Kacelnik and Todd (1992; Todd and Kacelnik, 1993) set up an operant analog of the MVT scenario in order to pursue the details of how foraging pigeons respond to travel time. A red flashing light signaled to the pigeon that food was available in the patch. As soon as the pigeon pecked the light, it changed to a steady red light, and a white light was also turned on. At this point the pigeon could choose to peck either the red light to obtain food in the patch or the white light to leave the patch and initiate the travel component of the schedule (that was actually a waiting time in this task). The food in the patch was delivered according to a progressive-interval schedule, such that the delay between food items increased with each successive prey delivered (as required for the MVT to apply). Using this schedule, Kacelnik and Todd (1992) could change various features of the travel component of the schedule and measure the effect that their manipulations had on the dependent variable, which was the number of prey per patch visit (PPV) taken by the bird before it pecked the white key and initiated the next travel component of the schedule.

In their first study, Kacelnik and Todd (1992) chose to test the MVT prediction that PPV should be sensitive only to mean travel time and should not be affected by variance in travel time. They tested pigeons in three treatments, all of which had mean travel times of 95 sec, but which differed in the variance in travel time. Treatment 10t consisted of a random order of ten different travel times with a coefficient of variation of 60.5%, treatment 2t consisted of a random order of two travel times with a coefficient of variation of 95%, and treatment 1t consisted of a single travel time. Contrary to the predictions of the MVT, the distribution of travel times had an effect on the mean PPV, with the birds having the highest PPV in the 1t treatment, an intermediate PPV in the 10t treatment, and the lowest PPV in the 2t treatment. Thus, PPV decreased as the coefficient of variation in travel time increased.

In order to explain the observed inverse relationship between PPV and variance in travel time, Kacelnik and Todd (1992) considered a modification of the MVT in which the representation of travel times in memory was not a perfect mean, as assumed in the unconstrained MVT, but was instead a distribution, as assumed in the scalar timing model. As shown in Figure 5.2, the travel time in the 1t treatment will be represented by a symmetrical memory representation, but the 10t and 2t distributions become progressively asymmetrical and skewed to the right. Kacelnik and Todd (1992) combined this memory representation of travel times with the MVT by assuming that in order to form an estimate of the background rate of energy intake available in the environment, the bird draws a single random sample from its reference memory holding the representation of experienced travel times. This value is then used to calculate the optimal PPV in the current patch. Due to the skew in the representations of the variable travel times, random samples drawn from the 10t representation will have an average shorter than 95 sec, and samples drawn from the 2t representation will have an average that is shorter still. Because shorter travel times will lead to a smaller optimal number of PPV, the modified MVT model is capable of explaining the observed effects of variance in travel time. In addition to explaining the effects of variance in travel time on PPV, the modified MVT model also explains another feature of the data not accommodated by the unconstrained MVT. The unconstrained MVT predicts that in the 1t treatment there should be no variance in the PPV taken by a bird; however, in reality the pigeons did have variance in their PPV in the 1t treatment. The scalar timing modification of the MVT explains this variance, because even a single travel time is represented in memory as a distribution from which random samples are taken for the purposes of decision making.

Despite the successes of the modified MVT, molecular analysis of the data from the 2t and 10t treatments revealed an important result that was not predicted by either unconstrained or modified MVT. When Kacelnik and Todd (1992) examined the PPV taken by a bird in relation to the previous travel time the bird had experienced, they found a significant positive effect, with the pigeons taking more PPV if the previous travel time had been long. Thus the pigeons appear to be not only responding to the mean and variance of the mixture of travel times, but also weighting the most recently experienced travel time more highly in their decision making. Todd and Kacelnik (1993) confirmed this finding in a subsequent study using the same paradigm that was designed to explicitly address the relative roles of the mixture of travel times and the most recently experienced travel time on PPV. They tested pigeons in two treatments that differed in mean travel time but had similar coefficients of variation (60.7 and 67.9% for the short and long mean treatments, respectively). Importantly, two of the travel times (1 and 13 sec) were contained in both mixtures. Todd and Kacelnik (1993) studied the effects of travel time by comparing the PPV after various travel times within each treatment and by comparing the PPV after equal travel times (1 and 13 sec) between treatments. The within-treatment analysis showed that in the long mean travel time treatment, PPV was correlated with the previous travel time, as they had found in the previous study. When the same travel times were compared between treatments, the pigeons were found to take higher PPV in the long mean travel time treatment than in the short mean travel time treatment. Because this effect cannot be accounted for by the within-treatment effect of the previous travel time, it implies that there is a different and independent effect on PPV of the mixture of travel times.

Because scalar timing theory is a steady-state model that makes no statements about how the reference memory representations are built up, it currently does not address short-term changes in behavior due to very recent experience. However, from a functional point of view it makes sense that the adaptive length of a memory window could vary depending on the stability of the environment. In very stable environments, it would make sense to base decisions on all of the available experience; however, in more changeable environments, it might be adaptive to weight recently acquired information more heavily than information acquired longer ago. In order to remedy this problem, Todd and Kacelnik (1993) developed a dynamic version of their previous model that involved adding an explicit learning algorithm to scalar timing theory and combining this with the MVT. Their aim was to develop a model that could handle the parallel effect of both recent and longer-term memory on foraging decisions that was suggested by their empirical data. Scalar timing theory provided a natural way to approach this problem because its parallel structure, whereby samples may be read at the same time from working and reference memory, allows a way to separate the effects of current percepts of time from remembered experience.

Todd and Kacelnik's new model contains two innovations that allow recent experience to have greater impact on decision making. First, the reference memory representation is built up and continually modified according to recent experience.

The reference memory is defined as a probability density function with bins corresponding to travel times each assigned a probability, and the total area under the function always equal to 1. Following each travel, the reference memory is updated in two steps. First, a fraction (α) of the area under the probability function is subtracted by devaluing each bin in proportion to its probability value at the time, such that the sum of the devaluations equals α. Second, an area the size of α is added back to the probability in the bin corresponding to the current travel time in working memory. Thus, following updating, the total area under the probability density function remains at unity, but the shape of the distribution is shifted toward the most recently experienced travel time, with the size of this shift controlled by the value of the parameter α. Low values of α correspond to little weight being given to recent experience, as would be predicted in a stable environment, whereas high values of α correspond to a high weight being given to recent experience, as would be predicted in changeable environments. The second innovation involves the decision subsystem. Rather than just using a random sample of travel time from reference memory in order to choose the PPV for the patch, Todd and Kacelnik (1993) assumed that value of travel time used was a weighted average of the value currently in working memory representing the most recently experienced travel time and a random sample taken from reference memory. They introduced a second parameter, β, that controlled the relative weight given to the values from working and reference memory. Just as for α, low values of β correspond to little weight being given to recent experience, whereas high values correspond to a high weighting of recent experience. As in their previous model, the weighted average of the two travel times was used as the input to the MVT to produce the optimal PPV.

Simulations of the above model run for the two different treatments experienced by the pigeons produced results that mimicked both of the main empirical results: the model produced both the observed treatment difference, with higher PPV in the high mean travel time patch than in the low mean travel time patch, and the within-treatment molecular effect of higher PPV directly following longer travel times.

Again, this study provides a clear example of the benefits of the ethological approach. Without the scalar timing model, we did not have an explanation for why pigeons should respond differently to fixed and variable travel times of the same mean. However, the scalar timing model shows that these effects occur because travel times are stored in reference memory as a distribution rather than as a mean, and as we have seen previously, the distribution representing a fixed travel time is symmetrical, whereas the distribution representing a variable travel time is asymmetrical (see Figure 5.2). The contribution of the evolutionary approach is the realization that it may not always be adaptive to weight all previous experience equally, and therefore the scalar timing model needs mechanisms that control the weighting given recent and past experience.

5.5.3 SCALAR TIMING AND RESPONSE TO RISK

There have been several experimental investigations of how foraging animals respond to variability in both delay to reward and amount of reward. For example, Bateson and Kacelnik (1995b) used an operant paradigm to study risk sensitivity in starlings. The starlings' preferences were tested in two treatments, one with variability in amount of reward and one with variability in delay to reward, and the energy budgets of the birds were maintained unchanged throughout the experiment. In the variable-amount treatment, the starlings were initially trained that a flashing red light indicated the availability of a fixed interval of 20 sec that culminated with a reward of 5 units of food, and that a flashing green light signaled the availability of a fixed interval of 20 sec culminating in a reward that was either 3 or 7 units of food with equal probability. In the variable-delay treatment, the fixed option was identical to the variable-amount treatment, but the other color signaled the availability of a variable interval of either 2.5 or 60.5 sec with equal probability, culminating in 5 units of reward. The assignment of colors to options was balanced across birds. The birds initially learned about the two options available in a treatment in forced trials in which only one option was presented. Following this training, the birds' choices were tested in trials in which both options were presented simultaneously, and the bird was required to commit to one of the options by choosing to peck one of the two flashing lights.

Bateson and Kacelnik (1995b) recorded two measures of preference in both treatments: the latency to peck the flashing lights in the training trials, and the option chosen by the birds in the choice trials. The choice data revealed that the birds were indifferent to risk in the variable-amount treatment, but strongly risk prone in the variable-delay treatment. The latency results showed that the birds had a shorter latency to begin a fixed-amount trial in the variable-amount treatment and a shorter latency to begin a variable-delay trial in the variable-delay treatment. Thus the results provide support for the general finding that animals tend to be risk averse to variability in amount, but strongly risk prone to variability in delay (Kacelnik and Bateson, 1996), because the birds preferred the variable-delay option even though the fixed delay of 20 sec was well below the mean delay of 31.5 sec in the variable-delay option.

The above results are not predicted by any purely evolutionary model; however, Reboreda and Kacelnik (1991) proposed that this pattern of preference could be explained by a modification of a version of scalar timing theory first proposed by Gibbon et al. (1988) to explain the preferences for delayed rewards seen in concurrent schedules. Gibbon et al. proposed that when animals are subjected to two alternative options between which they are required to choose, they build up a separate reference memory representation for each option. Thus, in Bateson and Kacelnik's experiment, the starlings in the variable-delay treatment would build up one reference memory for the fixed option and one for the variable option. As explained previously, the memory representation for the fixed interval would be a symmetrical distribution centered on the experienced interval, whereas the memory representation for the variable interval would be an asymmetrical distribution resulting from the sum of a symmetrical low variance distribution with a mean of 2.5 and a symmetrical high variance distribution with a mean of 60.5. The starling is assumed to choose between the two options by taking one random sample from each reference memory distribution, comparing these two samples, and choosing the option that yielded the lower sample (see Figure 5.5). Due to the skew in the representation of the variable delay (see Figure 5.2), this distribution will on average yield smaller samples that the

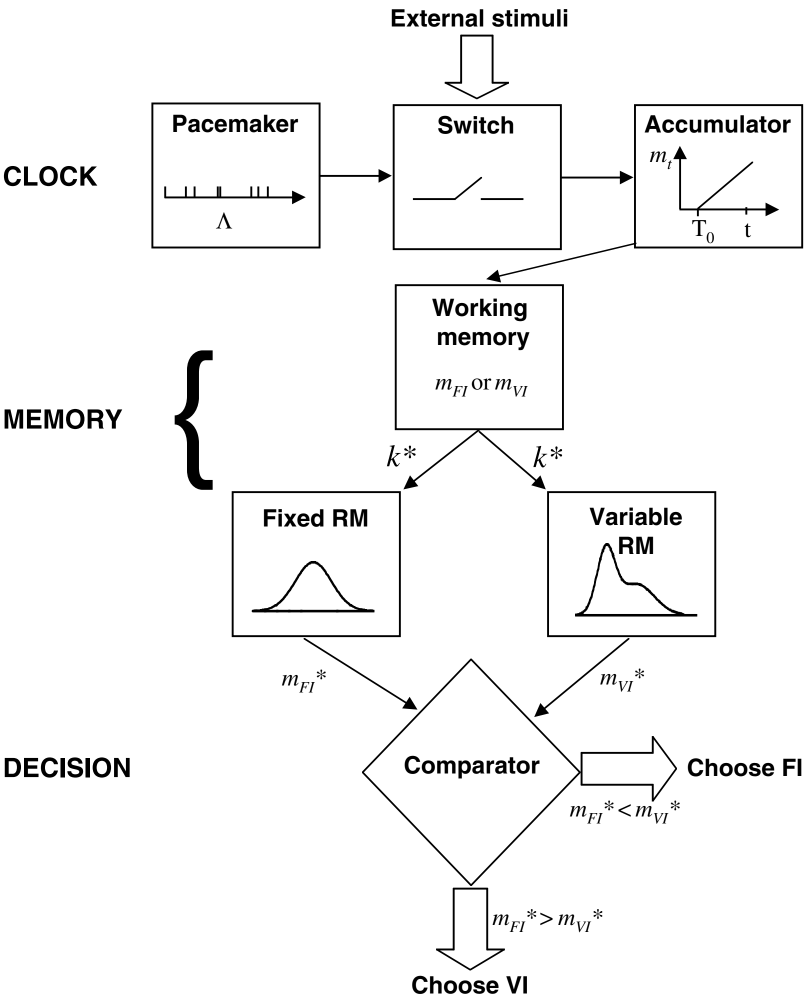

FIGURE 5.5 Diagrammatic representation of the scalar timing model applied to a choice scenario. In this case, the model is applied to the choice between a fixed interval and a variable interval (composed of two different intervals), as is found in many risk-sensitive foraging experiments.

symmetrical distribution representing the fixed option. This model can therefore explain why animals might be risk prone to variance in delay to reward. Reboreda and Kacelnik's (1991) innovation was to realize that the same argument could explain risk aversion to variance in amount. They suggested that if the birds form memory representations similar to those amounts assumed for delays, then the variable-amount option will be represented by an asymmetrical distribution formed from the sum of a symmetrical low variance distribution with a mean of 3 and a symmetrical higher variance distribution with a mean of 7. This is not an unreasonable assumption because in most experiments manipulating the amount of reward, the size of a reward is positively correlated with the time taken to consume it; thus consumption time could be used as a mechanism for measuring amount. A random sample drawn from the asymmetric distribution representing the variable-amount option will, as for the variable-delay option, on average be smaller than a random sample drawn from a symmetrical distribution with a mean of 5. However, in the case of amounts, the bird prefers the option yielding the larger sample, because although short delays to food are preferable to long delays, large amounts of food are preferable to small amounts.

One of the attractive features of the scalar timing theory account of choice is that the model not only makes predictions about the direction of preference, but also makes quantitative predictions about the exact magnitude of preference. Bateson and Kacelnik (1995b) proved that if a fixed-delay option is compared with a variable-delay option in which the delay is either short or long with equal probability, scalar timing theory predicts that the two options will become subjectively equivalent when the delay in the fixed option is equal to the geometric mean (i.e., $\sqrt{\text{short delay} \times \text{long delay}}$) of the two delays in the variable option. Bateson and Kacelnik (1996) tested this prediction in a subsequent starling experiment. They used a titration procedure to find the fixed delay equal to the variable mixture of 2.5- and 60.5-sec delays used in their previous experiment, and showed that indifference occurred when the fixed delay was equal to 5.61 sec, which is significantly below the geometric mean of 2.5 and 60.6 that equals 12.30 sec. Therefore, although scalar timing theory predicts the qualitative features of choice for variable delays, in this instance it failed to predict the quantitative detail of the results.

In the case of risk-sensitive foraging, the main outcome of considering the mechanistic account provided by scalar timing theory is that behavioral phenomena previously interpreted by behavioral ecologists as adaptive (e.g., Caraco et al., 1990) potentially now emerge as nonadaptive artifacts of the mechanisms by which animals maximize their rate of energy intake.

5.5.4 SCALAR TIMING AND SAMPLING A CHANGING ENVIRONMENT

Shettleworth et al. (1988) tested Stephens' (1987) sampling model, described previously, by setting up an operant analog of the sampling problem in the laboratory. Their experiment was carried out with pigeons in a shuttlebox in which the two feeding sites were represented by feeders and keys at either end of the box. On the stable side of the box, a random ratio schedule delivered rewards with some probability that remained unchanged throughout an experimental treatment. On the fluctuating side of the box, the schedule varied between no reward and a random ratio schedule equal to or lower (i.e., delivering more frequent rewards) than the stable side. Two different colored lights provided information about the state of the fluctuating side as soon as the bird pecked there. Following each reinforcement, the pigeon was required to return to the center of the box to initiate a new trial. Changes in the state of the fluctuating side occurred with a probability of .002 per trial. Shettleworth et al. (1988) investigated how the pigeons allocated their choices between the stable and fluctuating sides of the box in a range of treatments in which both the probability of reinforcement in the stable patch and the probability of reinforcement in the good state of the fluctuating patch were manipulated.

As predicted by Stephens' optimality model, the results showed that when the fluctuating patch was bad, pigeons chose the stable patch most of the time, only infrequently sampling the fluctuating patch. However, if on a sampling visit the fluctuating patch was in its good state, the pigeon would stay there until it reverted to its bad state. Furthermore, as predicted by the optimality model, the pigeons increased their rate of sampling the fluctuating patch as the probability of reinforcement in the constant patch was decreased. However, the behavior of the pigeons deviated in three important ways from the predictions of the unconstrained optimality model. First, sampling did not occur at the regular intervals predicted by Stephens' model, but instead occurred at random intervals. Second, sampling frequency was not sensitive to the probability of reinforcement in the good state of the fluctuating patch, and sampling still occurred when the probability of reinforcement in the good state of the fluctuating patch was equal to that in the stable patch. Finally, when the fluctuating patch was in its good state, the birds occasionally visited the constant patch, thus reducing their rate of food intake.

Because scalar timing theory was developed to model experiments in which rewards occur on the basis of time, but in this experiment Shettleworth et al. (1988) used ratio schedules in which rewards were delivered on the basis of responses, it was necessary to make the assumption that the pigeons responded at a constant rate in order to use scalar timing theory to model the data. Having made this assumption, the problem faced by the pigeon converts into a simple choice between two different distributions of delays to reward. The pigeon can simply be viewed as asking, "What is the delay to food if I stay with the current stable side vs. what is the delay to food if I sample?" According to scalar timing theory, the memory representation of the stable side will be an exponential distribution. The memory representation of the fluctuating side is more complex, and Shettleworth et al. (1988) assume that the time to food can be simplified to the sum of the time in the bad state until the good state begins and the time in the good state until food, which are both assumed to be represented by exponential distributions. As in the risky choice version of scalar timing described in the previous section, the pigeon is assumed to take a random sample from each of its two memory representations and use these to decide which side offers the shortest delay to food. Once the good state on the fluctuating side has been entered, a new reference memory for the fluctuating side applies in which the delay to reward is represented by an exponential distribution of delays as on the stable side. Thus the choice is now between two exponentially distributed delays to reward.

The scalar timing model can provide an explanation for some of the ways in which the data deviated from the unconstrained optimality model. First, sampling is not predicted to occur at regular intervals because decisions are based on random samples from memory distributions. For the same reason, both sampling of the fluctuating patch when the probability of reward is the same as the probability of reward in the stable patch, and reverse sampling of the stable patch when the fluctuating patch is currently in its good state, are predicted to occur, because occasionally the memory sample drawn from the memory for the suboptimal side will suggest a shorter delay to food than the sample drawn from the memory of the optimal side.

5.6 CONCLUSIONS

My intention has been to give some insight into how, by combining the evolutionary approach of optimal foraging theory with the hypotheses about the psychological mechanisms underlying interval timing, we can move toward a fuller understanding of both foraging behavior and timing. In all of the foraging examples described above, the basic strategy has been to take the constraints embodied in the scalar timing model, apply these to different foraging problems, and investigate how they affect optimal foraging decisions.

In all cases, the addition of scalar timing explains details of the birds' behavior that could not be accounted for with an unconstrained evolutionary approach. Three key assumptions of the scalar timing model emerge as being particularly important in producing these results. The first important assumption is the scalar property, whereby the duration of an interval is proportional to the precision with which it is represented in memory. The second assumption is that time intervals are assumed to be represented in memory as distributions of the intervals experienced rather than as a single statistic. Finally, the third crucial assumption is that reference memory is accessed via random sampling. There is extensive evidence for the scalar property; however, the latter two assumptions are harder to test empirically, and given the failure of the basic scalar timing model to account for some quantitative features of behavioral data (e.g., Bateson, 1993; Bateson and Kacelnik, 1996; Brunner et al., 1997), it is probably wise to maintain an open mind about the exact details of how memory is represented and accessed at this stage (e.g., Kacelnik and Brito-E-Abreu, 1998).

The examples described also highlight how our understanding of interval timing can benefit by thinking in detail about what animals use timing for. Specifically, the applications above have given insights into how the functional demands of different uses of temporal information may have shaped how this information is both represented and used in decision making. The basic scalar timing model makes no assumptions about the learning process and assumes that animals are equipped with memories that represent all possible experiences. However, the study of Todd and Kacelnik (1993) makes it clear that scalar timing models need to take account of learning, and that it may make sense for the memory to be weighted in favor of recent experience. In terms of decision making, the study of Brunner et al. (1992) shows that the biases assumed in the decision stage of scalar timing should be thought of as reflecting the optimal trade-off between responding early and late, and that the position of this trade-off will differ depending on the decision being made.

The flexibility of the scalar timing model has proved central in applying the model to a range of different foraging problems. The modularity of the model makes it easy to modify the details of how temporal information is represented in memory and how this information is used. For example, if we look first at how information is represented in the various examples, some of the applications described require a single reference memory, whereas others involving choice assume that there are two (Bateson and Kacelnik, 1995b; Reboreda and Kacelnik, 1991) or more (Shettleworth et al., 1988) memories. In most applications, experience is weighted equally in reference memory, but in some cases, it is necessary to assume that memory is biased toward recent experience (Todd and Kacelnik, 1993). If we look at how the temporal information represented in memory is subsequently used, again all of the applications make different assumptions. In some cases, a single sample from memory is compared with the current value in the accumulator (Brunner et al., 1992); in other cases, a weighted average of a sample from memory and the value in the accumulator is used in decision making (Todd and Kacelnik, 1993); and in yet other applications, the critical comparison is between two samples drawn from different reference memories (Bateson and Kacelnik, 1995b; Reboreda and Kacelnik, 1991; Shettleworth et al., 1988). An interesting point that emerges from Brunner et al.'s (1992) study is that the same reference memory may be accessed by two or more different decision-making mechanisms.

As a final thought, it is interesting to note that the flexible use of the same basic components described above has a close analogy with current thinking about the evolutionary process. The neural mechanisms responsible for producing new adaptive behavior patterns are not created from scratch, but are shaped by natural selection from small modifications of existing mechanisms. Thus, it is plausible that the tinkering necessary to apply scalar timing to a range of foraging problems is an accurate reflection of how interval timing mechanisms have evolved.

ACKNOWLEDGMENTS

I thank the Royal Society and the Natural Environment Research Council for financial support. The ideas presented in this chapter have benefited from the many discussions about timing that I have had with my mentors and colleagues, including John Gibbon, Sue Healy, Alex Kacelnik, and Warren Meck.

REFERENCES

Bateson, M., Currencies for Decision Making: The Foraging Starling as a Model Animal, unpublished doctoral dissertation, Oxford University, U.K., 1993.

Bateson, M., Interval timing, in *Frontiers of Life*, Baltimore, D., Dulbeco, R., Jacob, F., and Levi-Montalcini, R., Eds., Academic Press, New York, 2001, pp. 241–250.

Bateson, M. and Kacelnik, A., Accuracy of memory for amount in the foraging starling (*Sturnus vulgaris*), *Anim. Behav.*, 50, 431–443, 1995a.

Bateson, M. and Kacelnik, A., Preferences for fixed and variable food sources: variability in amount and delay, *J. Exp. Anal. Behav.*, 63, 313–329, 1995b.

Bateson, M. and Kacelnik, A., Rate currencies and the foraging starling: the fallacy of the averages revisited, *Behav. Ecol.*, 7, 341–352, 1996.

Bateson, M. and Kacelnik, A., Risk-sensitive foraging: decision making in variable environments, in *Cognitive Ecology*, Dukas, R., Ed., Chicago University Press, Chicago, 1998, pp. 297–341.

Brunner, D., Fairhurst, S., Stolovitzky, G., and Gibbon, J., Mnemonics for variability: remembering food delay, *J. Exp. Psychol. Anim. Behav. Process*., 23, 68–83, 1997.

Brunner, D., Kacelnik, A., and Gibbon, J., Optimal foraging and timing processes in the starling, Sturnus vulgaris: effect of inter-capture interval, *Anim. Behav.*, 44, 597–613, 1992.

Brunner, D., Kacelnik, A., and Gibbon, J., Memory for inter-reinforcement interval variability and patch departure decisions in the starling, *Sturnus vulgaris*, *Anim. Behav.*, 51, 1025–1045, 1996.

Caraco, T., Blanckenhorn, W.U., Gregory, G.M., Newman, J.A., Recer, G.M., and Zwicker, S.M., Risk-sensitivity: ambient temperature affects foraging choice, *Anim. Behav.*, 39, 338–345, 1990.

Charnov, E.L., Optimal foraging: the marginal value theorem, *Theor. Popul. Biol.*, 9, 129–136, 1976.

Church, R.M., Timing and temporal search, in *Time and Behaviour: Psychological and Neurobiological Analyses*, Bradshaw, C.M. and Szabadi, E., Eds., Elsevier, Amsterdam, 1997, pp. 41–78.

Davies, N.B., Prey selection and the search strategy of the spotted flycatcher (Muscicapa striata), *Anim. Behav.*, 25, 1016–1033, 1977.

Engen, S. and Stenseth, N.C., A general version of optimal foraging theory: the effect of simultaneous encounters, *Theor. Popul. Biol.*, 26, 192–204, 1984.

Gibbon, J., Scalar expectancy theory and Weber's law in animal timing, *Psychol. Rev.*, 84, 279–325, 1977.

Gibbon, J. and Church, R.M., Sources of variance in an information processing theory of timing, in *Animal Cognition*, Roitblat, H.L., Bever, T.G., and Terrace, H.S., Eds., Erlbaum, Hillsdale, NJ, 1984, pp. 465–488.

Gibbon, J., Church, R.M., Fairhurst, S., and Kacelnik, A., Scalar expectancy theory and choice between delayed rewards, *Psychol. Rev.*, 95, 102–114, 1988.

Gibbon, J., Church, R.M., and Meck, W.H., Scalar timing in memory, in *Timing and Time Perception*, Gibbon, J. and Allan, L., Eds., Academic Press, New York, 1984, pp. 52–77.

Gill, F.B., Trapline foraging by hermit hummingbirds: competition for an undefended renewable resource, *Ecology*, 69, 1933–1942, 1988.

Goss-Custard, J.D., Feeding dispersion in some overwintering wading birds, in *Social Behaviour in Birds and Mammals*, Crook, J.H., Ed., Academic Press, London, 1970, pp. 3–35.

Hinton, S.H. and Meck, W.H., How time flies: functional and neural mechanisms of interval timing, in *Time and Behaviour: Psychological and Neurobiological Analyses*, Bradshaw, C.M. and Szabadi, E., Eds., Elsevier, New York, 1997, pp. 409–457.

Kacelnik, A., Central place foraging in starlings (Sturnus vulgaris): I. Patch residence time, *J. Anim. Ecol.*, 53, 283–299, 1984.

Kacelnik, A. and Bateson, M., Risky theories: the effects of variance on foraging decisions, *Am. Zool.*, 36, 402–434, 1996.

Kacelnik, A. and Brito-E-Abreu, F., Risky choice and Weber's law, *J. Theor. Biol.*, 194, 289–298, 1998.

Kacelnik, A. and Brunner, D., Timing and foraging: Gibbon's scalar expectancy theory and optimal patch exploitation, *Learn. Motiv.*, 33, 177–195, 2002.

Kacelnik, A., Brunner, D., and Gibbon, J., Timing mechanisms in optimal foraging: some applications of scalar expectancy theory, in *NATO ASI Series*, Series G, *Ecological Sciences*, Vol. G 20, Hughes, R.N., Ed., Springer-Verlag, Berlin, 1990, pp. 63–81.

Kacelnik, A. and Todd, I.A., Psychological mechanisms and the marginal value theorem: effect of variability in travel time on patch exploitation, *Anim. Behav.*, 43, 313–322, 1992.

Killeen, P.R., Fetterman, J.G., and Bizo, L.A., Time's causes, in *Time and Behaviour: Psychological and Neurobiological Analyses*, Bradshaw, C.M. and Szabadi, E., Eds., Elsevier, Amsterdam, 1997, pp. 79–131.

Lejeune, H. and Wearden, J.H., The comparative psychology of fixed-interval responding: some quantitative analyses, *Learn. Motiv.*, 22, 84–111, 1991.

Matell, M.S. and Meck, W.H., Reinforcement-induced within-trial resetting of an internal clock, *Behav. Process.*, 45, 159–171, 1999.

Matell, M.S. and Meck, W.H., Neuropsychological mechanisms of interval timing behaviour, *Bioessays*, 22, 94–103, 2000.

Paule, M.G., Meck, W.H., McMillan, D.E., Bateson, M., Popke, E.J., Chelonis, J.J., and Hinton, S.C., The use of timing behaviors in animals and humans to detect drug and/or toxicant effects, *Neurotoxicol. Teratol.*, 21, 491–502, 1999.

Reboreda, J.C. and Kacelnik, A., Risk sensitivity in starlings: variability in food amount and food delay, *Behav. Ecol.*, 2, 301–308, 1991.

Schneider, B.A., A two-state analysis of fixed-interval responding in the pigeon, *J. Exp. Anal. Behav.*, 12, 677–687, 1969.

Shettleworth, S.J., Krebs, J.R., Stephens, D.W., and Gibbon, J., Tracking a fluctuating environment: a study of sampling, *Anim. Behav.*, 36, 87–105, 1988.

Skinner, B.F., The origins of cognitive thought, *Am. Psychol.*, 44, 13–18, 1989.

Stephens, D., On economically tracking a variable environment, *Theor. Popul. Biol.*, 32, 15–25, 1987.

Stephens, D.W. and Krebs, J.R., *Foraging Theory*, Princeton University Press, Princeton, MA, 1986.

Tinbergen, N., On aims and methods of ethology, *Z. Tierpsychol.*, 20, 410–433, 1963.

Todd, I.A. and Kacelnik, A., Psychological mechanisms and the marginal value theorem: dynamics of scalar memory for travel time, *Anim. Behav.*, 46, 765–775, 1993.

Yates, M.G., Stillman, R.A., and Goss-Custard, J.D., Contrasting interference functions and foraging dispersion in two species of shorebird (Charadrii), *J. Anim. Ecol.*, 69, 314–322, 2000.

Yoccoz, N.G., Engen, S., and Stenseth, N.C., Optimal foraging: the importance of environmental stochasticity and accuracy in parameter estimation, *Am. Naturalist*, 141, 139–157, 1993.