Discrimination, discounting and impulsivity: a role for an informational constraint

David W. Stephens

Ecology, Evolution and Behavior, University of Minnesota, St Paul, MN 55108, USA (dws@forager.cbs.umn.edu)

Feeding animals often prefer small, quickly delivered rewards over larger, more delayed rewards. Students of feeding behaviour typically explain this behaviour by saying that animals discount delayed benefits. Temporal discounting implies that delayed benefits are worth less than immediate benefits. This paper presents a new explanation of short-sighted decision-making called the discrimination advantage model that does not rely on discounting. A new model that includes several possible causes of discounting is developed. This model has many interesting features, but it cannot account for two empirical results: the strength of the 'discounting' effect and the fact that the time between choice presentations (the intertrial interval or ITI) has no effect. This leads to the conclusion that although discounting may be important it is probably not a complete explanation of the experimental facts. In the discrimination advantage model the observation that the ITI does not affect choice is seen as a strategy to make a cleaner discrimination between delayed alternatives in a noisy world. A simple example shows that when discrimination is imperfect a short-sighted choice rule can, in some situations, lead to a higher long-term rate than a rule that actually compares long-term rates. This idea is developed and extended in several ways.

Keywords: foraging; feeding; information; impulsivity; choice; discrimination

1. INTRODUCTION

When a starling (*Sturnus vulgaris*) detects a leatherjacket (e.g. *Tipula paludosa*), it may attack or ignore it. If an attack succeeds, the starling obtains some food after a delay involving pursuit and handling. If instead it ignores this item, it obtains its next meal after a longer delay. The starling's choices determine the timing and quantity of the food it obtains, and this, in turn, influences the starling's survival and reproductive success. It is natural to wonder how the timing and size of food gains affect the starling's fitness. This is a basic issue in behavioural ecology. An animal's actions determine the magnitude *and timing* of fitness gains. Students of foraging have assumed that the effects of many small foraging choices combine to influence fitness over a long time-horizon. Results from the psychological laboratory, however, challenge this long-term view, because they indicate that animal feeding decisions are guided by short-term considerations. Behavioural ecologists have come to view these psychological results as evidence that animals devalue, or discount, delayed food benefits, and have offered models to explain how delay might reduce fitness. This paper argues that discounting is not a very plausible explanation of the observed facts, and offers, instead, a model based on information-processing constraints.

This section reviews the required background material. First, it outlines the evolutionary logic of far-sighted decision-making. Second, it reviews the psychological evidence for animal short-sightedness, or impulsivity. Third,

it reviews existing theoretical attempts to account for animal impulsivity using the concept of discounting, and finally it shows how even a comprehensive model of discounting cannot explain the observed facts.

(a) *The logic of far-sightedness*

Reconsider our hypothetical starling and suppose that one alternative, say attacking the leatherjacket, leads to a one leatherjacket gain after a 10 s delay, while the other, say ignoring the leatherjacket and waiting for something better, leads to a much larger food gain after a longer time (say 100 leatherjacket equivalents after 2 min). Clearly, an animal that only 'considered' the next 10 s would choose to attack, but an animal that considers a longer time-interval (the morning, the day, this foraging bout) would choose to wait for the bigger benefit. Which option makes a larger contribution to the starling's fitness? Food contributes to fitness by enhancing survival and providing resources for reproduction, and crudely speaking we would expect that many feeding events will be required before successful reproduction is possible. Most readers will, I think, agree that the long-term choice will, typically, lead to a bigger increase in fitness. With these ideas in mind, foraging theorists have used the long-term rate of food gain as a proxy for fitness in their models. In formulating the long-term rate it is conventional to recognize two components of foraging: search or travel time, the time spent between encounters with food resources (patch or prey items); and exploitation or handling time, the time spent actually exploiting a particular food resource. The long-term rate is then

$$\frac{E(A)}{E(\tau) + E(t)} \quad (1.1)$$

One contribution of 12 to a Theme Issue 'Information and adaptive behaviour'.

Figure 1. Diagram of a trial within a self-control experiment. After waiting τ time units the apparatus presents a choice between a small immediate option (A_1 delivered after t_1) and a larger delayed option (A_2 delivered after t_2).

where $E(A)$ is the expected amount gained per encounter, $E(\tau)$ is the expected time spent travelling (or searching), and $E(t)$ is the expected time spent exploiting. There are many valid criticisms of this formulation (it is insensitive to variance, it ignores predation risk and so on), but for our purposes it is only important to recognize the long-term rate model as the simplest and most extreme manifestation of our conclusions that animals should be far-sighted.

(b) *Psychological evidence of short-sightedness: self-control procedures*

Notwithstanding our evolutionary arguments that animals should make far-sighted choices, psychological studies provide many examples of animal short-sightedness, or impulsivity. Although this evidence comes from several sources, one especially well articulated case for animal short-sightedness comes from the so-called self-control literature. In self-control experiments, an animal repeatedly chooses between a small-immediate and a large-delayed option. Figure 1 shows a typical trial from a self-control experiment. The subject waits for τ seconds and then is presented with two stimuli (say a red and a green pecking key). From previous experience it has learnt that the red key leads to the small-immediate consequence, and the green key leads to the large-delayed consequence. After the animal makes a choice (say, it pecks the green key) it must wait for the programmed delay to expire before receiving the corresponding amount of food (in our example, it waits t_1 seconds to obtain amount A_1 after pecking the red key, but it waits t_2 seconds to obtain amount A_2 after pecking the green key). After food delivery the animal must wait another τ seconds before the next presentation. Typically, the experimenter designates a small proportion of the trials as forced-choice or no-choice trials, in which only one option is available with the goal of ensuring that the subject has some experience with both options. Many manipulations are possible within this paradigm. For example, the experimenter can parametrically vary the delays and amounts to assess their effects on choice. In addition, this paradigm can be complicated in several interesting ways.

(i) *Relevant self-control results*

Self-control studies over the past 20 years show several key results (see Stephens & Anderson (2001) for a review). Three of these results are relevant here. First, manipulations of delay and amount have the expected qualitative effects. Specifically, reducing the delay associated with an option will increase preference for that option. Similarly increasing the amount associated with an option will also increase preference for that option. Second, delay has a surprisingly powerful effect on preference. For example, it is relatively easy to arrange situations where the subject prefers the small-immediate option even when choosing the large-delayed option would yield a higher long-term rate. Third, the ITI has little or no effect on preference. Gallistel & Gibbon (2000) call this the 'subjective sunk time' effect. This contradicts long-term models because in the long term all time commitments matter, and as explained below, long-term models predict a shift toward large-delayed options when the ITI is long.

(ii) *Comparing self-control and patch theory*

This paper offers a new model to explain the surprising strength of animal preferences for immediacy, and the subjective sunk time effect. To understand the need for this new model, one needs to understand the predictions of long-term theory for the self-control situation. Many readers will be familiar with the graphical solution of elementary patch exploitation theory (figure 2). When the time spent travelling from one patch to the next is long, a long-term rate maximizer should spend more time in, and extract more food from, each patch. Moreover, a great many patch studies support this qualitative prediction: experimental increases in travel time increase patch residence time (Stephens & Krebs 1986). A similar situation arises when, as in self-control experiments, a forager must choose one of two options. When, for example, a honeybee (*Apis mellifera*) encounters a pair of flowers, the honeybee's problem is essentially a patch problem: should it spend a long time to extract a larger amount from the blue flower, or choose the more quickly obtainable but less rewarding yellow flower? One expects, by analogy with long-term patch theory, that the larger, more time-consuming flower should be chosen when environmental intake rates are low, but the smaller, less time-consuming flower should be chosen in poor environments (see the experimental study by Stephens *et al.* (1986)). We expect that travel time should influence choice, because the travel time is a key determinant of the environmental intake rate.

(iii) *The disagreement: summary*

The self-control data disagree with traditional foraging models in two ways. First, the effect of delay is too strong, because animals will prefer a sufficiently immediate small food reward even when they could do better in the long term by choosing the delayed alternative. Second, the ITI has no effect in the self-control preparation even though the analogous travel time is theoretically and empirically important in studies of patch exploitation.

(c) *Discounting*

From an economic point of view, it is not surprising that animals prefer immediate rewards. We are all familiar with the idea that delay reduces value—\$10 received one

Figure 2. Comparison of classical patch model (a) and the self-control situation (b). (a) Graphical solution of the optimal patch residence time. The curve shows the relationship between the patch residence time ('patch time') and gains, denoted by $g(t)$. The slope of the line connecting $(-\tau, 0)$ and $(t, g(t))$ is the rate of intake which the model maximizes. If the travel time is short (τ_1), a small patch residence time (t_1) maximizes the intake rate; when the travel time is long (τ_2), a larger patch residence time (t_2) maximizes the intake rate. (b) The self-control patch problem. The 'small, immediate' and 'large, delayed' options can be thought of as two-point gain functions, and one would expect a preference for the small, immediate option for short travel times (when $\tau < \tau^*$), and a preference for the large, delayed option for longer travel times (when $\tau > \tau^*$).

year from today is worth less than \$10 received today. Behavioural ecologists have appealed to this phenomenon of temporal discounting to explain self-control results (Kagel *et al.* 1986; McNamara & Houston 1987; Benson & Stephens 1996; Bateson & Kacelnik 1997).

(i) Causes of discounting

Although delay can reduce value in several ways, we can divide explanations of discounting into two types: collection-risk discounting and opportunity-cost discounting. Each of these can be further divided into two subcategories. *Collection-risk discounting* occurs when delay reduces the likelihood that a benefit will actually be realized. Delay can increase the risk of loss via two related mechanisms: termination risk and interruptions risk. In models of termination-risk discounting, we suppose that the animal collects or searches for benefits in a sequence—collecting one and then the next—and that there is a probability that some event will occur that terminates this sequence of gains. For example, if an animal is collecting benefits of undiscounted value V , and w is the constant probability that the process will continue from the i th to the $i + 1$ th collection, then the expected value of the sequence is

$$V + wV + w^2V + w^3V + \dots \quad (1.2)$$

In termination-risk discounting the second item in the sequence is worth less than the first, simply because there is a chance that the process will end before the second item can be collected. Interruptions-risk discounting has received the lion's share of attention from those trying to account for self-control results. In interruptions-risk discounting some process prevents the collection of a given benefit without terminating the entire sequence of gains. If one focuses on a single collection event, and ignores the properties of the sequence of gains (as psychologists typically do) then there is no need to distinguish termination risk from interruptions risk, as their effect is the same for a single benefit. If, however, we do consider the sequence (as I argue we should) then they are quite different. Let v be the probability that an encountered gain can actually be collected; then the expected value of a sequence of such gains is

$$vV + vV + vV + \dots \quad (1.3)$$

Equation (1.3) is quite different from the expression representing the value of a sequence under termination risk, equation (1.2).

The second class of explanations for discounting is *opportunity-cost discounting*. Animals waiting to obtain a benefit may lose the opportunity to enhance their fitness in other ways. Again, there are two distinct ways this can happen. In the first form of opportunity-cost discounting, waiting for item a delays the collection of future gains. This is simply a rate effect, and one can reasonably call it the rate cost of delay. The second opportunity cost of delay is the lost opportunity to invest (or, perhaps less anthropomorphically, the lost opportunity to use). In waiting, the animal foregoes the opportunity of using a given benefit. The example of a human investor makes the point most clearly. Suppose that an investor must choose between X now, and Y delayed by n periods, and that our investor can expect a rate of return of r per period. This means that when the delayed benefit becomes available, the immediate benefit will be worth $X(1 + r)^n$, so a rational investor should choose the immediate benefit whenever $X(1 + r)^n > Y$, or equivalently if

$$X > \frac{Y}{(1 + r)^n} \quad (1.4)$$

where one says that the delayed benefit is discounted by the factor $1/[(1 + r)^n]$. While this form of discounting is fundamental in human economics, it has received less attention from modellers of non-human discounting. Presumably this is because it is unclear what investment 'returns' might be for non-human animals. However, the more general idea that immediate benefits are more valuable because they can be put to use now is valid even without a specific analogue of 'rate of return' in mind.

To summarize, I have identified four reasons to devalue delayed benefits: termination risk, interruptions risk, the rate cost of delay and the investment cost of delay. I remark that there is a logical connection between termination-risk discounting and the rate cost of delay. A key reason for animals to value high rates is that they have only a finite time to collect resources.

(ii) *Modelling discounting**Exponential model*

The simplest and most basic model of discounting is the exponential model, which supposes that value is lost at a constant rate per unit time k :

$$\frac{dV}{dt} = -kV \text{ leads to } V(t) = V(0)e^{-kt}. \quad (1.5)$$

This formulation focuses on the value of single delayed benefit, and not a sequence of benefits; as such this model cannot account for any rate cost of delay. The parameter k in this model may be interpreted as measuring any of the other three discounting effects—termination risk, interruption risk, investment costs—or some combination of them (recall, however, that termination risk and interruption risk are not meaningfully different for a single benefit).

The exponential model is mathematically appealing and convenient, but fatally flawed. The key evidence against it is the phenomenon of preference reversal. In preference reversal, experimentalists arrange a situation in which a small, immediate food reward is preferred to a larger delayed food reward. Once this preference is established the experimenter increases the delay of both options, and this shifts preference from the small more immediate reward to the larger more delayed reward. The exponential model cannot predict this, because if

$$V_1 e^{-kt_1} > V_2 e^{-kt_2} \quad (1.6)$$

then it must also be that $V_1 e^{-k(t_1 + \Delta)} > V_2 e^{-k(t_2 + \Delta)}$ because the laws of exponents make the additional delay behave like a common factor that can be cancelled away.

The single-trial hyperbolic model

To account for preference reversal, animal psychologists have offered the hyperbolic model of discounting,

$$V(t) = \frac{V(0)}{1 + kt}, \quad (1.7)$$

which does predict preference reversal. The single trial hyperbolic model fits the data, but explanations of why animals should discount in this way over a single trial have, typically, been unconvincing, requiring, for example, that interruption rates decline with time in an *ad hoc* way (Green & Myerson 1996). Recently, however, Souza (1998) has offered a stronger rationale for hyperbolic discounting. As explained above (equation (1.5)), if the interruption rate is constant, we have exponential discounting. Souza imagines, instead, that the interruption rate itself varies stochastically. He shows that if the interruption rate is drawn from an exponential distribution, then one expects hyperbolic discounting and that the hyperbolic parameter k should equal the mean interruption rate. Souza's model is a great improvement, but it focuses rather narrowly on how interruptions influence choice over a single trial. Indeed, Souza rejects the rate costs of delay that one would expect when animals make repeated choices. This is rather odd, because all animal discounting experiments have repeated choices, but none have interruptions.

A basic problem with both the exponential and hyperbolic models is that they require extremely high k values to account for observed animal preferences. For example,

although these data are not reported directly, we can calculate hyperbolic k values from a well-cited study by Mazur (1987) averaging 2.8 s^{-1} ; this implies that 75% of a food reward's value is lost in the first second of delay. Similar calculations (using the same study) indicate exponential k values ranging from 0.07 to 0.1 s^{-1} , and implying a loss of 7–10% of an item's value in the first second of delay. The problem here is that the effects of interruptions risk and the investment costs of delay are bound to be small effects, even in combination. For example, even my most modest estimate of 0.07 s^{-1} (in the exponential model) is equivalent to 4.2 interruptions per minute!

(d) *Discounting from first principles: the EDS model*

We have a problem. Although many behavioural ecologists believe that discounting causes animal impulsiveness, we do not have a plausible, first-principles model of discounting. To fill this gap, I have developed a simple model that, in effect, combines the exponential and rate models. In its simplest incarnation, this model supposes that an animal collects amount A every $\tau + t$ seconds, so that the sum of exponentially discounted gains would be

$$e^{-\lambda t}(A + e^{-\lambda(\tau+t)}A + e^{-2\lambda(\tau+t)}A + e^{-3\lambda(\tau+t)}A \dots), \quad (1.8)$$

where λ is the discount rate, combining termination and lost opportunity effects. (Note that the entire sequence is discounted by the factor $e^{-\lambda t}$ because we are interested in the value at the point where the animal makes a choice. In the self-control preparation, the choice point is after the ITI and before the delay.) I call this the EDS model, because it simply sums up the exponentially discounted benefits of a sequence of food gains. The term in parentheses is the sum of a geometric series and so we can write

$$\frac{e^{-\lambda t}A}{1 - e^{-\lambda(\tau+t)}} \quad (1.9)$$

as a more compact (closed) form of our model. In the slightly more complex problem, where a forager chooses option 1 (amount A_1 , delay t_1) with probability p , and option 2 (A_2 , t_2) with probability $1 - p$, we have

$$\frac{pe^{-\lambda t_1}A_1 + (1-p)e^{-\lambda t_2}A_2}{1 - e^{-\lambda\tau}[pe^{-\lambda t_1} + (1-p)e^{-\lambda t_2}]}, \quad (1.10)$$

which gives the discounted benefits in a self-control-like situation when the forager cannot (or simply does not) express an absolute preference.

(i) *Advantages of the EDS model*

The EDS model combines the logic of the single-trial exponential and multiple-trial long-term rate models to create a model of discounted value that has many advantages as follows.

- (i) For small discounting rates (λ) it is approximately hyperbolic. To show this, we rearrange equation (1.9) to find

$$\frac{e^{-\lambda t}A}{1 - e^{-\lambda(\tau+t)}} = \frac{A}{e^{\lambda t} - e^{-\lambda\tau}}$$

Now, if λt is small, we can approximate $e^{\lambda t} \approx 1 + \lambda t$, substituting this into the expression above and rearranging we find

$$\frac{A}{1 + \lambda t - e^{-\lambda t}} = \frac{\frac{A}{1 - e^{-\lambda \tau}}}{1 + \frac{\lambda}{1 - e^{-\lambda \tau}} t},$$

where the expression on the right is in hyperbolic form (the numerator represents the discounted value of A when there is no delay—one A obtained every τ time units—and the hyperbolic rate parameter is $k = \lambda/(1 - e^{-\lambda \tau})$).

- (ii) For zero discounting rates ($\lambda = 0$) it becomes the long-term rate model (modified, as appropriate, to incorporate a term measuring the likelihood of collection).
- (iii) It is qualitatively consistent with the observed phenomenon of preference reversal.
- (iv) It is qualitatively consistent with the asymmetric effect of delay even when 'cycle' time is held constant (cf. Green *et al.* 1981)
- (v) It agrees with empirical results (Kacelnik & Bateson 1996) in predicting risk-prone choice over variable delays.

(ii) Failure of the EDS model

Although the EDS model is a considerable improvement over existing explanatory models of discounting (single-trial exponential and long-term rate), it falls short of being a complete explanation of self-control results. There are two problems. First, self-control studies typically show (see, for example, Bateson & Kacelnik (1996); Mazur *et al.* (1985)) that τ (the ITI) has little or no effect on preference even though the EDS model predicts that τ should affect preference. This is the 'subjective sunk time' effect of Gallistel & Gibbon (2000). In Gallistel and Gibbon's analogy animals erroneously treat the ITI as a sunk cost because it precedes choice. This is an error because although previously experienced ITIs are economic sunk costs, so are previously experienced delays and amounts. Previously experienced delays, amounts and ITIs are important because they are reliable guides to the magnitudes of future delays, amounts and ITIs, which are, most certainly, not sunk costs.

Second, there is a significant quantitative disagreement between the EDS theory and the self-control data. Specifically, discount rate parameters large enough to account for the data lead to a model that is closer to the repudiated single-trial exponential model than to the empirically supported hyperbolic model. Consider the claim of Mazur (1987) that the hyperbolic k parameter is typically close to one (i.e. 1 s^{-1} as the experimental units of delay are seconds). This implies (as mentioned in § 1c) that the first second of delay causes a 50% loss in value. Assuming a typical ITI of 60 s, we need to find a discount rate (λ) such that

$$\frac{1}{2} \frac{A}{1 - e^{-\lambda 60}} = \frac{e^{-\lambda 1} A}{1 - e^{-\lambda 61}} \quad (1.11)$$

to account for 50% loss in 1 s. For even modest values of λ , $e^{-\lambda 60} \approx 0.0$ so equation (1.11) is approximately

$\frac{1}{2} = e^{-\lambda}$ or $\lambda = 0.693$. The problem here is that the EDS model is approximately exponential for these parameter values, and so we lose many of the appealing hybrid properties of the EDS model. Notice also that from the point of view of explanatory ideas about discounting, a discount rate of 0.7 (per second!) is ridiculous. If, for example, we interpret λ as a measure of termination risk, $\lambda = 0.7$ implies an average feeding bout duration of less than 2 s, and we would find similarly absurd values if we considered lost investment costs or interruptions accounts of discounting. So while the EDS model is appealing because it incorporates the economic logic of discounting and gives many of the correct qualitative effects, it does not account for the strong effects of delay or the absence of an ITI effect in the self-control situation.

(e) Background summary

This introduction has explained the conflict between self-control results and evolutionary models of feeding behaviour, and outlined a comprehensive model of discounting. This discounting model has many advantages but it cannot account for the self-control results. In the remainder of this paper, I argue that we must incorporate constraints on discrimination to understand the economically puzzling features of the self-control results.

2. DISCRIMINATION ADVANTAGE HYPOTHESIS

In my introductory animal behaviour lectures I give the following demonstration. I place a small weight (say a heavy pen) in one of two opaque plastic shopping bags. I hand both bags to several students asking them to tell me which is heavier. The students easily identify the heavier bag. Then I add a heavy textbook to each bag, so one bag contains a textbook and a small weight, while the other contains only a textbook. Now the students have great difficulty telling me which bag is heavier, even though the *difference* in weights is the same as before. Most readers will recognize this as a demonstration of Weber's Law which states that the smallest detectable change in stimulus intensity (ΔI , called the just noticeable difference) is proportional to the initial stimulus intensity (I); in symbols $\Delta I = cI$ where c is a constant of proportionality (see Gescheider 1985). The additional weight of the pen (ΔI) is easy to detect when the bags are otherwise empty, as $\Delta I > cI_{\text{small}}$ but more difficult to detect with the additional weight of the textbooks, as $\Delta I < cI_{\text{large}}$. This paper argues that it makes sense to ignore the ITI in some choice situations (even though the ITI is economically important), because the ITI makes it more difficult to discriminate small differences in delay just as the weight of the textbooks makes it more difficult to discriminate a small difference in weight.

(a) A motivating example

This section develops an example that shows the importance of informational constraints. In this example, a forager that decides on the basis of a short-term comparison achieves a *higher long-term rate* of gain than a forager that makes a direct comparison of long-term rates. This surprising result occurs because discrimination is imperfect, and short-term comparisons provide a stronger, more easily discriminated signal.

Consider a self-control situation as shown in figure 1. To simplify the presentation, I assume that the discount rate λ is small so that the EDS model of fitness value is well approximated by the long-term rate. Specifically, I assume that

$$\frac{pA_1 + (1-p)A_2}{\tau + pt_1 + (1-p)t_2} \quad (2.1)$$

gives the approximate fitness benefits associated with choosing option 1 with probability p and option 2 with probability $1-p$.

(i) Scores and evaluations

I assume that our hypothetical animal evaluates its options by comparing 'scores' that summarize its experience with each option. For example, an animal following a long-term rate rule would compare

$$\ell_1 = \frac{A_1}{\tau + t_1} \text{ with } \ell_2 = \frac{A_2}{\tau + t_2} \quad (2.2)$$

to decide which alternative is best (assuming, of course, that experience has provided it with reasonable estimates of the experimental parameters A_i , t_i and τ).

Recall that the EDS model predicts that the ITI τ should affect preference, even though empirical results show that τ has little influence. Our example focuses on this disagreement by asking whether a forager should attend to τ in choosing between alternatives. An animal using the long-term scores (equation (2.2)) is clearly attending to τ but an animal comparing

$$s_1 = \frac{A_1}{t_1} \text{ with } s_2 = \frac{A_2}{t_2} \quad (2.3)$$

is obviously not. These short-term scores consider only the delays to food and the associated amounts. An evaluation of this form agrees with many data. It has been suggested by several authors (Baum & Rachlin 1969; Bateson & Kacelnik 1996) and it is very similar to the hyperbolic model of Mazur (1987) when the t_i terms are large (technically, when kt_i is large, as this is dimensionless number with a meaningful magnitude). I call the s_1 versus s_2 comparison a short-term comparison, because it ignores food gains beyond the current trial. The difference between the ℓ_1 versus ℓ_2 comparison and the τ -less s_1 versus s_2 comparison provides a convenient contrast between long- and short-term rules.

(ii) Comparison and imperfect discrimination

The scores themselves are only a partial specification of a choice mechanism. We must also consider how the comparison is made and how it is translated into action. Psychophysical results support a ratio evaluation (Gescheider 1985; Stevens & Stevens 1986; Gallistel & Gibbon 2000). This means, for example, that an animal using the short-term scores s_1 and s_2 compares their relative magnitudes via the ratio s_1/s_2 ; tending to choose option 1 when $(s_1/s_2) > 1$ and tending to choose option 2 when $(s_1/s_2) < 1$.

I say 'tending to choose' because discrimination is necessarily a noisy process. Noise can be introduced at several levels: stimuli vary, sensory systems are imperfect, neural representations may vary and so on. This noise

Figure 3. Two hypothetical dose-response or psychometric functions. The preference for alternative 1 increases with the ratio s_1/s_2 ; we expect indifference when $s_1/s_2 = 1$. The shallow curve shows a situation where discrimination is difficult, while the steeper, more step-function-like curve shows an easier discrimination problem.

means that we must think of discrimination as a graded and probabilistic process. When two stimuli (or 'scores' in our situations) are very different, the animal will almost always detect the difference. If, however, they are very similar, differences will seldom be detected. Psychophysicists use a sigmoid (or ogival) psychometric function to describe imperfect discrimination. Figure 3 shows how one can apply this idea to our choice situation. If s_1/s_2 is much smaller than one, p will be near zero; if s_1/s_2 is much larger than one, p will be near one; but, if $s_1/s_2 = 1$, the two options are equivalent and p should be near $\frac{1}{2}$. Now, the slope of the central portion of this psychometric function tells us something about the noisiness of our discrimination problem. If the slope is very steep, the psychometric function approaches a step function indicating a relatively easy, noiseless discrimination problem (figure 3). A shallow slope indicates a noisier and more difficult discrimination problem (figure 3). (I remark that signal-detection theory provides a more sophisticated way of incorporating noise in 'economic' problems, but this psychometric function approach is adequate for the purposes of our example).

(iii) The fixed amount case

Now consider a self-control situation where both options offer the same amount of food, say A . This is an especially simple situation because under virtually any model of value the animal should choose the option with the smallest delay, and so it faces a relatively simple temporal discrimination problem. In our scenario, however, the animal cannot set $p = 1$ or $p = 0$, because discrimination is noisy and imperfect.

Consider how our two types of scores do in this situation. An animal using the long-term scores forms the ratio

$$R_\ell = \frac{\frac{A}{\tau + t_1}}{\frac{A}{\tau + t_2}} = \frac{\tau + t_2}{\tau + t_1} \quad (2.4)$$

and an animal using the short-term scores forms the analogous ratio

$$R_s = \frac{\frac{A}{t_1}}{\frac{A}{t_2}} = \frac{t_2}{t_1}. \quad (2.5)$$

These two ratios always agree about the direction of preference, because R_ℓ will be greater than one whenever R_s is greater than one. They will, however, be quite different in magnitude. Consider what happens to R_ℓ as τ increases. Obviously, when $\tau = 0$ then $R_\ell = R_s$, but as τ increases it begins to dominate the quotient $\tau + t_2/\tau + t_1$, and the quotient approaches one, effectively ‘damping out’ any difference in delays (t_1 and t_2). We conclude that the long-term ratio, R_ℓ , will always be closer to one than the short-term (delay only) ratio, R_s .

Now consider a sigmoid psychometric function $f(R)$ where, as required, $f(1) = \frac{1}{2}$, and $f(R) \rightarrow 1$ for R large, and $f(R) \rightarrow 0$ for R small (figure 3). It follows that an animal that evaluates its options via the long-term ratio, R_ℓ , will be closer to indifference than an animal using the short-term ratio, R_s , and conversely that animal using the short-term ratio will show a stronger preference for the option with the shorter delay. *In other words, the animal using the short-term rule comes closer to the long-term optimum than the animal using the long-term rule!*

This is precisely analogous to the weight detection demonstration described above: the ‘ t_1 versus t_2 ’ discrimination is easier than the ‘ $\tau + t_1$ versus $\tau + t_2$ ’ discrimination in the same way that the ‘empty bag versus bag with small weight’ discrimination is easier than the ‘heavy bag versus heavy bag plus small weight’ discrimination. This observation indicates a simple explanation of the ‘subjective sunk time’ effect: even if including τ accurately reflects the evolutionary economics of choice, including τ in a comparative rule can make it *more difficult to detect differences in delay*. The reader will recall that self-control results disagreed with evolutionary models in two ways: the effect of τ on preference, and the surprisingly strong effect of delay on preference. The discrimination advantage hypothesis can, in principle, explain both. Delay will have a much more powerful effect for an animal using a rule based on the short-term rate (A/t), than for an animal using the economically reasonable long-term rate ($A/(\tau + t)$) (see figure 4).

Notice that the interpretation of our example depends on our assumption that both options yield the same amount of food, because this fixed amount makes the choice problem a temporal discrimination problem. Consider a different arrangement in which amount and delay are proportional, say $A_i = ct_i$; in this case, a short-term rule could be disastrous because both options lead to $A/t = c$ even when one option provides a much larger long-term rate of intake. The discrimination advantage gained by adopting a short-term rule depends on the properties of the feasible set of alternatives that the forager must discriminate. § 2b considers how this feasible set interacts with discounting rates to influence choice rules.

(b) Generalizing the model

Consider a slightly generalized rule in which the forager compares

Figure 4. The effect of τ on a rate model of discounting, $V = A/(\tau + t)$, where V is the discounted value, A is the amount, τ is the ITI (or travel time), and t is the delay. As the figure shows when τ is small, the delay (t) has a larger effect on the value. So a choice model that ignores τ will also have stronger discounting level.

$$s_1 = \frac{A_1}{\beta\tau + t_1} \text{ and } s_2 = \frac{A_2}{\beta\tau + t_2}, \text{ where } 0 \leq \beta \leq 1. \quad (2.6)$$

If $\beta = 0$, we have the short-term comparison (as in equation (2.3)), and if $\beta = 1$, we have the long-term comparison (as in equation (2.2)). Our problem is now to find the β value that gives the best-discounted sequence of benefits; to find this we must first specify a dose-response relationship (or psychometric function) that translates the comparison of s_1 and s_2 into a choice probability. For the purposes of this preliminary development I will use the matching law

$$p = \frac{s_1}{s_1 + s_2}, \quad (2.7)$$

which provides a simple, tractable dose-response relationship. Finally, I use the EDS model

$$\frac{pe^{-\lambda t_1}A_1 + (1-p)e^{-\lambda t_2}A_2}{1 - e^{-\lambda\tau}[pe^{-\lambda t_1} + (1-p)e^{-\lambda t_2}]} \quad (2.8)$$

to represent the discounted benefits gained by a forager facing a given choice (i.e. a particular combination of (A_1, t_1) , (A_2, t_2) and τ).

With these three pieces in hand, we can express the expected benefits associated with the rule parameter β , denoted by $W(\beta)$, by substituting (2.6) into (2.7), and substituting the result into (2.8). We can then study the effect of β on benefits by simple differentiation; although this process is rather involved, the results can be stated simply. First, β does not affect the sign of $W'(\beta)$, so that (for a given choice) the optimal β (β_{opt}) must be either 0 (the short-term rule) or 1 (the long-term rule). Second, we can show that the long-term rule is best if

$$t_2 > t_1 \text{ and } \frac{A_2 e^{-\lambda t_2}}{1 - e^{-\lambda(\tau + t_2)}} > \frac{A_1 e^{-\lambda t_1}}{1 - e^{-\lambda(\tau + t_1)}} \quad (2.9)$$

while the short-term rule is best if

$$t_2 > t_1 \text{ and } \frac{A_1 e^{-\lambda t_1}}{1 - e^{-\lambda(\tau + t_1)}} > \frac{A_2 e^{-\lambda t_2}}{1 - e^{-\lambda(\tau + t_2)}}. \quad (2.10)$$

In other words, the long-term rule is best when the more delayed alternative yields the larger benefit, while the short-term rule is best when the more immediate option is better. This generalizes the intuition we developed in the constant amount case: when choosing the smallest delay the short-term rule provides a cleaner discrimination.

This result gives us a simple way to study the combined effects of ‘choice set’ and discounting rate. Consider a fixed relationship between delay and amount, say $A_i = f(t_i)$. Now we know that any set of delays such that the benefit

$$\frac{f(t)e^{-\lambda t}}{1 - e^{-\lambda(\tau + t)}} \quad (2.11)$$

always increases with t represents a range of choices for which the long-term rule is best, and conversely that whenever benefits decrease with t a short-term rule is best.

(i) *An example*

The power law $A_i = bt_i^c$ admits several interesting possibilities; for example, the constant amount case ($c = 0$) where we expect the short-term rule to do well, the simple linear case ($c = 1$) where we expect the long-term rule to do best, and relationships that either ‘bend up’ ($c > 1$) or bend down ($c < 1$). In this case, we can show that there exists a critical delay, $\hat{t}$, such that the long-term rule is better if t_1 and t_2 are both less than $\hat{t}$, and the short-term rule is better if both delays are greater than $\hat{t}$. This critical time value is given by

$$c - \hat{t}\lambda - ce^{-\lambda(\tau + \hat{t})} = 0. \quad (2.12)$$

Unfortunately, this exact equation cannot be solved analytically. We can, however, show that

$$\hat{t} \approx \frac{c}{\lambda} \left(1 - \frac{\exp\left(-\lambda\left(\tau + \frac{c}{\lambda}\right)\right)}{1 - c \exp\left(-\lambda\left(\tau + \frac{c}{\lambda}\right)\right)} \right). \quad (2.13)$$

The leading term c/λ provides insight into the interaction between discounting rate and properties of the choice set (expressed here as the curvature parameter c). For example, if $c = 0$ we have the constant amount case, and as outlined above the short-term rule is always best (as $\hat{t} = 0$). When $c > 0$, the critical delay $\hat{t}$ increases with c and decreases with discount rate λ , giving the intuitively reasonable result that conditions favouring the short-term rule are broader when the discount rate is high.

(ii) *Numerical explorations*

These calculations show that the long-term rule is best when all possible delays are less than the critical value, $\hat{t}$, and that the short-term rule is best when all possible delays exceed $\hat{t}$. Of course, this ‘all above or all below’ distinction does not tell us what to expect when some delays are less than $\hat{t}$, and others are greater than $\hat{t}$. I performed numerical calculations to address this question. These calculations assume that delays (t_1 and t_2) are drawn from a uniform distribution with lower bound $\bar{d} - (r/2)$ and upper bound $\bar{d} + (r/2)$, so $\bar{d}$ is the mean, and r is the range, so that the expected benefit is the double integral

Figure 5. A typical result from numerical studies of the effect of $\bar{d}$ on the optimal value of the rule parameter β . As $\bar{d}$ increases there is an abrupt transition from $\beta_{\text{opt}} = 1$ (the long-term rule) to $\beta_{\text{opt}} = 0$ (the short-term rule).

$$\overline{W(\beta)} = \frac{1}{r^2} \iint_{t_1, t_2 \in [d \pm \frac{r}{2}]} W(\beta, t_1, t_2) dt_1 dt_2. \quad (2.14)$$

I calculated this integral numerically for a range of parameter values and found the best β value for each set of parameters. The results closely parallel the ideas presented in the previous paragraph. When the average delay ($\bar{d}$) is small the long-term rule ($\beta = 1$) provides the larger expected benefit, and when the average delay is large the short-term rule ($\beta = 0$) is better. Figure 5 shows a typical plot of $\bar{d}$ versus β_{opt} . The figure indicates that we can think of the mean delay that gives an β_{opt} value of one half as a ‘critical mean delay’ (denoted by t_{crit}), such that long-term rules are generally better when $\bar{d} < t_{\text{crit}}$ and short-term rules are better otherwise. I calculated the t_{crit} using numerical root-finding procedures for a wide range of discount rates (λ), curvature parameters (c), travel times (τ) and ranges (r). Figure 6 shows the result. Each plot shows four curves: a dashed line shows the root calculated from equation (2.12) for the short travel time case ($\tau = 30$), while a dashed line marked with circles shows the computed t_{crit} values for the same short travel time case; a solid line shows the root calculated for the long travel time case ($\tau = 120$), while the solid curve marked with squares shows the corresponding computed values. The figure is arranged as a trellis of nine plots with each row showing the results for one of three ranges ($r \in [5, 30, 60]$), and each column showing the result for a different value of the curvature parameter ($c = 0.5$, giving a negatively accelerated delay–amount relationship, $c = 1$ a simple straight line, and $c = 1.5$ a positively accelerated delay–amount relationship). The figure shows several key points as follows.

- (i) The root of equation (2.12), which was derived for the simpler case where all delays are above or below the critical value, provides an excellent approximation for this more general case in which we focus on the position of the mean delay. In most cases the

Figure 6. Results of a systematic numerical study of how the discount rate (λ), curvature parameter (c) and range (r) affect the value of t_{crit} . (a–c) Range = 5: (a) $c = 0.5$, (b) $c = 1$, (c) $c = 1.5$; (d–f) range = 30: (d) $c = 0.5$, (e) $c = 1$, (f) $c = 1.5$; (g–i) range = 60: (g) $c = 0.5$, (h) $c = 1.5$, (i) $c = 1$. Each panel shows four plots of t_{crit} versus λ : a dashed line shows the approximation based on the root of equation (2.13) for the short travel time case ($\tau = 30$); a dashed line marked with circles shows the exact, computed values of t_{crit} for the short travel time case; similarly a solid line shows the approximate solution for the long travel time case ($\tau = 120$); and a solid line marked with squares shows exact computations for the long travel time case. In several cases the approximate and exact solutions are so close that only one curve can be seen.

approximation is so good that we cannot distinguish separate lines in the figure. Not surprisingly, this approximation does worse when the range of feasible delays is large (compare the range equals 5 and the range equals 60 cases).

- (ii) The critical delay decreases with the discount rate. When the discount rate is high, the short-term rule is favoured even when the average delay is small.
- (iii) The critical delay increases with travel time. This may seem surprising because as travel time increases a long-term comparison does a poorer job of discriminating delays. But, when travel times are long, it is less important to discriminate delays, because they comprise a smaller portion of the total foraging cycle. The value of the critical delay, therefore,

seems to reflect a trade-off between these two effects.

- (iv) The critical delay increases with the curvature parameter c ; in general, short-term rules are more likely to provide a discrimination advantage when the amount changes little with delay ($c < 1$).

While these results represent only a beginning in our understanding of how discounting and discrimination may interact to influence animal decision-making, they provide a concrete example of this interaction. These results show how discounting and discrimination can combine to create a situation that favours a short-term decision rule; and this indicates that implausibly large discounting rates are not necessary to explain observed animal impulsivity.

Figure 7. Diagram of a patch-like choice procedure. In this procedure the subject must decide whether to stay and collect amount A_2 or to leave and start another ITI. In this situation the short-term rule and long-term rule always agree about the direction of preference, but the short-term rule allows a cleaner discrimination.

3. DISCUSSION

I have argued that economic discounting cannot explain observed preferences for immediacy, because these data require implausibly large discounting rates. Delay does reduce the value of food resources, but these effects (e.g. collection risk or lost opportunity risk) should be relatively small. In addition, as the EDS model shows, the properties of observed preferences such as the subjective sunk time effect contradict economic models of discounting. I have offered the discrimination advantage hypothesis in an attempt to overcome the shortcomings of models based entirely on discounting. I suspect, however, that discrimination advantage is only part of the solution.

(a) Context

Binary-choice tests, as in the self-control experiment, make short-term rules look worse than they really are. Consider the short-term rule ($\beta = 0$) that considers only the delay to the next food item as discussed in the previous paragraph. In binary-choice tests, an animal applying the short-term rule sometimes performs abysmally, obtaining much less food than a hypothetical animal applying the long-term rule. Consider, however, the patch-like situation shown in figure 7: here the forager travels for time τ , enters a patch and waits t_1 seconds to obtain A_1 , and then it must decide whether to leave and search for new patch, or to wait an additional t_2 seconds for the additional amount A_2 . Here, the short-term rule compares $A_1/(\tau + t_1)$ (because $\tau + t_1$ is the delay to the next food item if the forager leaves) and A_2/t_2 . This short-term comparison agrees precisely with the long-term rule (it is, in fact, algebraically equivalent to the long-term rule). As I have argued elsewhere (Stephens & Anderson 2001), short-term rules can have long-term consequences in many naturalistic sequential choice situations. If we combine this 'context' effect with the discrimination advantage of short-term rules, we have a powerful explanation of animal short-sightedness.

In summary, I propose that three elements favour short-term rules. First, delayed benefits are discounted. Second, short-term rules have long-term consequences in many natural decision contexts, especially when animals make sequential rather than binary decisions. Third, as explained in this paper, short-term rules have a discrimi-

nation advantage over long-term rules. Discounting and its putative causes (e.g. interruptions) have received the lion's share of attention in models of animal impulsiveness, while context and discrimination effects have seldom even been recognized. I speculate that discounting is the least important of these three 'causes of short-sightedness', because I believe that plausible discounting rates are very small. Discounting would be a plausible explanation if the self-control evidence showed that animals preferred food delivered now to food delivered next week, but the self-control data show that delays measured in seconds, even tenths of seconds, have strong effects on preference.

Assessing the plausible level of discounting has important implications in behavioural ecology. If we reject the interpretation outlined above, and interpret the self-control results to mean that discounting rates are very high, then this means that feeding behaviour should, quite generally, be influenced by short-term consequences. For example, under strong discounting, it may be very difficult to establish cooperative behaviour (Stephens *et al.* 1995; Stephens 2000).

(b) More general models?

The ideas presented here could be generalized in several ways. One might, for example, consider a more complex psychometric function. I have used the algebraically simple matching law, but ideas from scalar expectancy theory (Gibbon 1977; Gibbon *et al.* 1988) would connect this approach to a well-supported model of temporal discrimination. As I have shown, the 'environment' can have a profound effect on the rule, but I have explored only a tiny subset of possible environments—consider the possible distributions of τ , t_1 , t_2 , A_1 and A_2 , and the relationships between them. Indeed, there so many possibilities that empirical observations would be invaluable in focusing our attention on plausible cases.

(c) Experiments

One might test these ideas in several straightforward experiments. A simple study, suggested by the constant amount example, would be to test sensitivity to τ (travel time or ITI) in two treatments when the amount is constant ($A = b$) and when the amount increases with delay (e.g. $A = br^c$, $c \geq 1$). If foragers adjust their rule as they experience new sets of possibilities, then we would predict more sensitivity to ITI when amount increases with delay. Assessing the role of discounting presents a more challenging problem, because it is difficult to control or measure discounting. Perhaps the most feasible solution is to increase discounting levels by creating programmed interruptions.

(d) Why is the discrimination constraint tolerated?

The discrimination advantage proposed here arises because I assume, following established psychophysical principles, that animals make ratio comparisons. It is not obvious why this should be so. Consider, for example, the constant amount case. An animal making a ratio comparison uses

$$\frac{\frac{A}{\tau + t_1}}{\frac{A}{\tau + t_2}} = \frac{\tau + t_2}{\tau + t_1} > 1 \quad (3.1)$$

and, as argued earlier, increasing the quantity τ tends to obscure the difference between t_1 and t_2 . If animals made additive comparisons,

$$\tau + t_2 > \tau + t_1, \quad (3.2)$$

then τ would not interfere with the comparison of delays, because $\tau + t_2 > \tau + t_1$ is algebraically equivalent to $t_1 > t_2$. My model, therefore, assumes that ratio comparisons are a constraint on psychophysical systems. I do not know why such a constraint should arise. The prevalence of ratio comparisons in psychophysics indicates that neural systems are predisposed to ratio comparisons. So, perhaps, ratio comparisons reflect a mechanistic constraint. Alternatively, it may be that ratio comparisons offer some adaptive advantage that is yet to be understood.

(e) Information constraints and impulsivity

This paper may seem an unlikely contribution to a volume on the adaptive significance of information use, especially because the development requires several digressions into complex issues, such as discounting models and psychological procedures, that have little to do with information. The observed impulsivity of animals, however, presents a long-standing challenge to models of feeding behaviour. Here, I develop the novel argument that this observed short-sightedness is a response to an informational constraint. The world is a noisy place, and short-sighted choice rules can lead to better long-term results because they provide a cleaner discrimination of delayed alternatives.

I thank A. Houston and S. Dall for helpful comments on the manuscript. The research described here was partly supported by the National Science Foundation (IBN 9896102) and the National Institute of Mental Health (R01-MH64151-01). I am grateful for the support of these agencies.

REFERENCES

- Bateson, M. & Kacelnik, A. 1996 Rate currencies and the foraging starling: the fallacy of the averages revisited. *Behav. Ecol.* **7**, 341–352.
- Bateson, M. & Kacelnik, A. 1997 Starling's preferences for predictable and unpredictable delays to food: risk-sensitivity or time-discounting? *Anim. Behav.* **53**, 1129–1142.
- Baum, W. & Rachlin, H. 1969 Choice as time allocation. *J. Exp. Analysis Behav.* **12**, 861–874.
- Benson, K. & Stephens, D. W. 1996 Interruptions, tradeoffs and temporal discounting. *Am. Zool.* **36**, 506–517.

- Gallistel, C. R. & Gibbon, J. 2000 Time, rate, and conditioning. *Psychol. Rev.* **107**, 289–344.
- Gescheider, G. A. 1985 *Psychophysics: method, theory, and application*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Gibbon, J. 1977 Scalar expectancy theory and Weber's law in animal timing. *Psychol. Rev.* **84**, 279–325.
- Gibbon, J., Church, R. M., Fairhurst, S. & Kacelnik, A. 1988 Scalar expectancy theory and choice between delayed rewards. *Psychol. Rev.* **95**, 102–114.
- Green, L. & Myerson, J. 1996 Exponential versus hyperbolic discounting of delayed outcomes: risk and waiting time. *Am. Zool.* **36**, 496–505.
- Green, L., Fisher Jr, E. B., Perlow, S. & Sherman, L. 1981 Preference reversal and self-control: choice as a function of reward amount and delay. *Behav. Analysis Lett.* **1**, 43–51.
- Kacelnik, A. & Bateson, M. 1996 Risky theories—the effects of variance on foraging decisions. *Am. Zool.* **36**, 402–434.
- Kagel, J. H., Green, L. & Caraco, T. 1986 When foragers discount the future: constraint or adaptation? *Anim. Behav.* **34**, 271–283.
- McNamara, J. M. & Houston, A. I. 1987 A general framework for understanding the effects of variability and interruptions on foraging behaviour. *Acta Biotheor.* **36**, 3–22.
- Mazur, J. E. 1987 An adjusting procedure for studying delayed reinforcement. In *Quantitative analyses of behavior: the effect of delay and of intervening events on reinforcement value* (ed. M. L. Commons, J. E. Mazur, J. A. Nevin & H. Rachlin), pp. 55–73. Hillsdale, NJ: Lawrence Erlbaum.
- Mazur, J. E., Snyderman, M. & Coe, D. 1985 Influences of delay and rate of reinforcement on discrete-trial choice. *J. Exp. Psychol.: Anim. Behav. Processes* **11**, 565–575.
- Souza, P. D. 1998 On hyperbolic discounting and uncertain hazard rates. *Proc. R. Soc. Lond. B* **265**, 2015–2020. (DOI 10.1098/rspb.1998.0534.)
- Stephens, D. W. 2000 Cumulative benefit games: achieving cooperation when players discount the future. *J. Theor. Biol.* **205**, 1–16.
- Stephens, D. W. & Anderson, D. 2001 The adaptive value of preference for immediacy: when shortsighted rules have far-sighted consequences. *Behav. Ecol.* **12**, 330–339.
- Stephens, D. W. & Krebs, J. R. 1986 *Foraging theory. Monographs in behavior and ecology*. Princeton University Press.
- Stephens, D. W., Lynch, J. F., Sorensen, A. E. & Gordon, C. 1986 Preference and profitability: theory and experiment. *Am. Nat.* **127**, 533–553.
- Stephens, D. W., Nishimura, K. & Toyer, K. B. 1995 Error and discounting in the Iterated Prisoner's Dilemma. *J. Theor. Biol.* **176**, 457–469.
- Stevens, S. S. & Stevens, G. 1986 *Psychophysics: introduction to its perceptual, neural, and social prospects*. New Brunswick, NJ: Transaction Books.

GLOSSARY

- EDS: exponentially discounted sequence
ITI: intertrial interval