

Time discounting and time preference in animals: A critical review

Benjamin Y. Hayden¹

© Psychonomic Society, Inc. 2015

Abstract Animals are an important model for studies of impulsivity and self-control. Many studies have made use of the *intertemporal choice task*, which pits small rewards available sooner against larger rewards available later (typically several seconds), repeated over many trials. Preference for the sooner reward is often taken to indicate impulsivity and/or a failure of self-control. This review shows that very little evidence supports this assumption; on the contrary, ostensible discounting behavior may reflect a boundedly rational but not necessarily impulsive reward-maximizing strategy. Specifically, animals may discount weakly, or even adopt a long-term rate-maximizing strategy, but fail to fully incorporate postreward delays into their choices. This failure may reflect learning biases. Consequently, tasks that measure animal discounting may greatly overestimate the true discounting and may be confounded by processes unrelated to time preferences. If so, animals may be much more patient than is widely believed; human and animal intertemporal choices may reflect unrelated mental operations; and the shared hyperbolic shape of the human and animal discount curves, which is used to justify cross-species comparisons, may be coincidental. The discussion concludes with a consideration of alternative ways to measure self-control in animals.

Keywords Discounting · Intertemporal choice · Animal cognition · Impulsivity · Self-control

✉ Benjamin Y. Hayden
benhayden@gmail.com

¹ Department of Brain and Cognitive Sciences and Center for Visual Science, University of Rochester, Rochester, NY 14620, USA

Impulsivity is a basic character trait that affects nearly every aspect of our lives (Bechara, 2005; Dalley, Everitt, & Robbins, 2011; Evenden, 1999; Tangney, Baumeister, & Boone, 2004). Its conjugate, self-control, predicts both lifetime success measures and susceptibility to diseases like addiction, obsessive-compulsive disorder, Tourette syndrome, and even depression and obesity (Bickel & Marsch, 2001; Mischel, Shoda, & Rodriguez, 1989). In 2002, Frederick and colleagues published a critical review of the psychology of time preferences in humans (Frederick, Loewenstein, & O'Donoghue, 2002). Much of their criticism was directed at the human intertemporal choice task. They argued that many of the findings from this task fail to generalize to other contexts, fail to make reliable predictions, and produce inconsistent data from study to study (for other critical reviews, see Scholten & Read, 2010; Van den Bos & McClure, 2013). Frederick et al. suggested that humans' choices in the intertemporal choice task do not arise solely from time attitudes, but are strongly biased by other extraneous factors beyond temporal preferences, thus throwing into doubt the meaning of the measures of discounting.

The goal of the present study is to critically review the analogous task for animals, the animal intertemporal choice task. Like Frederick et al. (2002), the discussion focuses on evidence that factors besides time preferences strongly bias measures of discounting in animals, although for different reasons. This evidence challenges the idea that intertemporal preferences provide a good measure of impulsivity or self-control in animals.

Methods for measuring animal impulsivity and self-control have become increasingly important for several reasons. First, invasive measures of brain function, such as single-unit recordings, are much more readily used in animals than in humans. Thus, animals are used as a model for the neural basis

of human cognitive processes (e.g., Hwang, Kim, & Lee, 2009; Louie & Glimcher, 2010). Second, animals can be more readily subjected to techniques that manipulate brain activity, such as optogenetics or psychopharmacology (e.g., Cardinal, Winstanley, Robbins, & Everitt, 2004; Heilbronner & Meck, 2014). Third, techniques devised for animal studies are often used in preverbal humans as a way of studying cognitive development. Fourth, animals are often used as a model for human diseases of self-control (Dalley et al., 2011; Monterosso & Ainslie, 1999). Finally, animals provide an essential counterpart to humans in studies of comparative cognition. Such studies can potentially provide insights into the evolution of patience (Stevens, Rosati, Heilbronner, & Mühlhoff, 2011; Stevens & Stephens, 2009).

Frederick et al. (2002) offered the helpful distinction between *time discounting*—devaluation of rewards for any reason, including maximizing long-term rewards by reducing opportunity costs—and *time preference*, which refers to the more specific situation in which the discounting is driven solely by a preference for sooner rewards (i.e., impulsivity). Thus, in their parlance, time discounting is consistent with perfect patience and does not necessarily imply anything about time preference, such as impulsivity. In the context of the intertemporal choice task, for example, selection of a smaller sooner (SS) option may reflect mental devaluation of the larger later (LL) reward because the delay cheapens it (time preference), or it may be a shrewd response to an environment in which choice of the SS option provides a sooner next trial, and thus better prospects in the long run (i.e., time discounting arising from a neutral time preference). A theoretical long-term rate-maximizing decision-maker with no impulsivity (i.e., an optimal forager) will prefer an SS to an LL option if the opportunity costs imposed by the LL choice outweigh its benefits. However, this discounter may, to a naive observer, appear to be impulsive (Bateson & Kacelnik, 1996; Blanchard, Pearson, & Hayden, 2013; Kacelnik, 2003; Namboodiri, Mihalas, Marton, & Hussain Shuler, 2014; Namboodiri, Mihalas, & Shuler, 2014; Pavlic & Passino, 2010; Stephens, 2002; Stephens & Anderson, 2001).

Here it is argued that the field has underappreciated this second motivation, rate maximization, for choice of the SS option in animal intertemporal choice tasks. Evolution may have driven animal choices to be (1) roughly long-term-rate maximizing and (2) computationally efficient. These two goals conflict, leading to systematic biases in the treatment of postreward delays. These biases have weak effects in natural contexts, but lead to steep and hyperbolic discounting in the intertemporal choice task (Bateson & Kacelnik, 1996). This discounting seems to imply strong time preference and poor self-control, but it is an artifact of a failure to fully understand the task structure. The conclusions here suggest that after factoring out these biases, animals' true *time preferences* are closer to neutral than is commonly believed (but likely still somewhat impulsive).

The intertemporal choice task

On each trial of the animal intertemporal choice task, animals choose between a large reward offered after a long delay and a small reward immediately or after a short delay (Fig. 1). Note that this task is sometimes called the *delay-discounting task*, *self-control task*, or *intertemporal trade-off task*. Here, the term *intertemporal choice task* is used to refer uniquely to a task with the structure shown in Fig. 1.

A typical finding in this task is that animals are indifferent between a small reward offered immediately and a reward twice as large offered several seconds later (Stevens & Stephens, 2009). Even the most patient nonhuman species, chimpanzees, will only wait over the order of minutes (Beran & Evans, 2009; Dufour, Pelé, Sterck, & Thierry, 2007; Rosati, Stevens, Hare, & Hauser, 2007). By contrast, when humans perform the task, equivalence is typically found when the delay gets out to several months—a difference of several orders of magnitude (Frederick et al., 2002; Kable & Glimcher, 2007; Kirby & Maraković, 1996; McClure, Laibson, Loewenstein, & Cohen, 2004). These results are consistent with the commonplace view of animals as much more impatient than humans (Kacelnik, 2003; Stevens & Stephens, 2009).

Psychologists have found high levels of discounting (suggesting impulsivity) in rats, several species of monkey, lemurs, nonhuman apes, jays, chickens, bumblebees, and other animals (e.g., Abeyesinghe, Nicol, Hartnell, & Wathes, 2005; Ainslie, 1974; Bateson & Kacelnik, 1996; Cheng, Peña, Porter, & Irwin, 2002; Hayden & Platt, 2007; McDiarmid &




Fig. 1 (A) Schematic of the structure of the standard intertemporal choice task. On each trial, the subject chooses between a smaller sooner (SS) and a larger later (LL) option. (Reward size here is indicated by the size of a water reward.) Choice of the SS is often taken to indicate impulsivity. (B) In most implementations, a postreward buffer is added to make sure that the trial lengths for both options are equivalent; otherwise, the subject can advance to the next trial sooner. This postreward buffer is identified here as a potentially problematic element of the design of the task

Rilling, 1965; Rachlin & Green, 1972; Rosati et al., 2007; Siegel & Rachlin, 1995; Snyderman, 1983; Stephens & Anderson, 2001; Stevens & Mühlhoff, 2012). Interestingly, these findings extend to normally patient humans in some contexts. For example, when discounting food, drink, or the opportunity to view photographs of other people, humans show high discounting (Hayden, Parikh, Deaner, & Platt, 2007; McClure, Ericson, Laibson, Loewenstein, & Cohen, 2007; Smith et al., 2010). These data suggest that the well-known patience in humans may not be all that robust (Frederick, Loewenstein, & O'Donoghue, 2002). Future studies will be needed to fully delineate the contexts that promote steeper discounting in humans.

Choice in the intertemporal choice task is most commonly fit with the hyperbolic choice function (see, e.g., Hwang et al., 2009; Mazur, 1987).

$$v_{discounted} = \frac{v}{1 + k \cdot D}.$$

The term $v_{discounted}$ indicates the subjective value of a reward of size v when it is delayed by D units of time. It is usually taken to reflect mental discounting due to time preference. However, the parameters derived from the equation do not, by themselves, provide a way to distinguish between time preference and time discounting.

The only free parameter in this equation is the discount factor k . Its units are the reciprocal of time, as measured in whatever unit is used for the variable D . Because the discount function is hyperbolic, the reciprocal of k indicates the half-life of the initial reward value (this is only approximate, due to the +1 in the denominator, which prevents the ratio from going to infinity as D goes to zero). A k value of 0.5, for example, indicates that reward is worth roughly half as much when the animal must wait 2 s to obtain it. The value of k is often taken to measure impulsivity, with larger values indicating greater impulsivity.

Intertemporal choice and self-control

The conceptual link between the intertemporal choice task and self-control crystallized with the publication of an influential review by George Ainslie in 1975. Prior to then, a large number of studies had quantified the effects of delay on animals' learning rate, incentive strength, and proportional preference (e.g., Hull, 1943; Renner, 1964). Much of this work was in the behaviorist tradition, and so avoided ascribing mental operations like discounting or psychological concepts like self-control to choices. Ainslie's advance was to link this literature with the psychology of self-control in humans, including in economic theory and social psychology, especially in Mischel's early research into self-control in children

(e.g., Mischel, Ebbesen, & Zeiss, 1972). Ainslie, trained as a psychiatrist, even brought in Freud. Emphasizing the parallels between these diverse traditions, Ainslie argued that the animal intertemporal choice task captures something fundamentally akin to human self-control, and thus provides a measure of self-control in animals. This view has been influential ever since. Ainslie's review motivated numerous studies on animal intertemporal choice behavior, much of it performed with an eye toward understanding human self-control.

One of the central arguments that Ainslie made to support the link between intertemporal choice and self-control is related to preference reversals in the intertemporal choice task (Ainslie, 1975; see also Rachlin, 2000). Animals, like humans, make more abstemious choices when comparing rewards in the distant future (e.g., a pigeon may prefer a large reward in 31 s to a small reward in 30 s), and more impulsive when considering reward in the immediate future (e.g., it may prefer a small reward immediately to a large reward in 1 s). In other words, their performance looks superficially very much like poor self-control in humans. It resembles the experience of a dieter who promises to give up milkshakes in the coming year on December 31st, but succumbs when confronted with one on January 7th. The intuitive similarity of these situations was a critical ingredient in Ainslie's (1975) argument that animal intertemporal choice tasks capture self-control as we conventionally think about it.

A human seeking to avoid milkshakes may avoid temptation by simply avoiding places where they are served. This behavior is prudent if a decision-maker recognizes that his or her self-control may fail when temptation arises. Thus, Ainslie (1975) argued that precommitment would provide additional evidence that intertemporal choices reflect self-control. In a study published the previous year, he had observed what appeared to be precommitment in pigeons (Ainslie, 1974): These birds willingly chose options that would reduce future choice by committing to an LL option early in long trials (Ainslie, 1974; Rachlin, 2000; Rachlin & Green, 1972). This precommitment is intuitively consistent with the idea that when the reward is close, the pigeon faces a battle against temptation that it knows it may lose (Thaler & Shefrin, 1981). Thus, precommitment, along with preference reversal and the intuitive similarity of the intertemporal choice task and self-controlled behavior in humans, justified Ainslie's (1975) argument that this task measures animal self-control.

Following work by Chung and Herrnstein (1967), Ainslie (1975) suggested that this pattern of behavior could be modeled with hyperbolic discount curves. A *discount curve* indicates the subjective value of a reward as a function of the delay to its receipt (Fig. 2). It starts at 1 (meaning that a reward immediately is not devalued) and declines as time extends into the future. Some discounting is normative; for example, time imposes both a collection risk and an opportunity cost

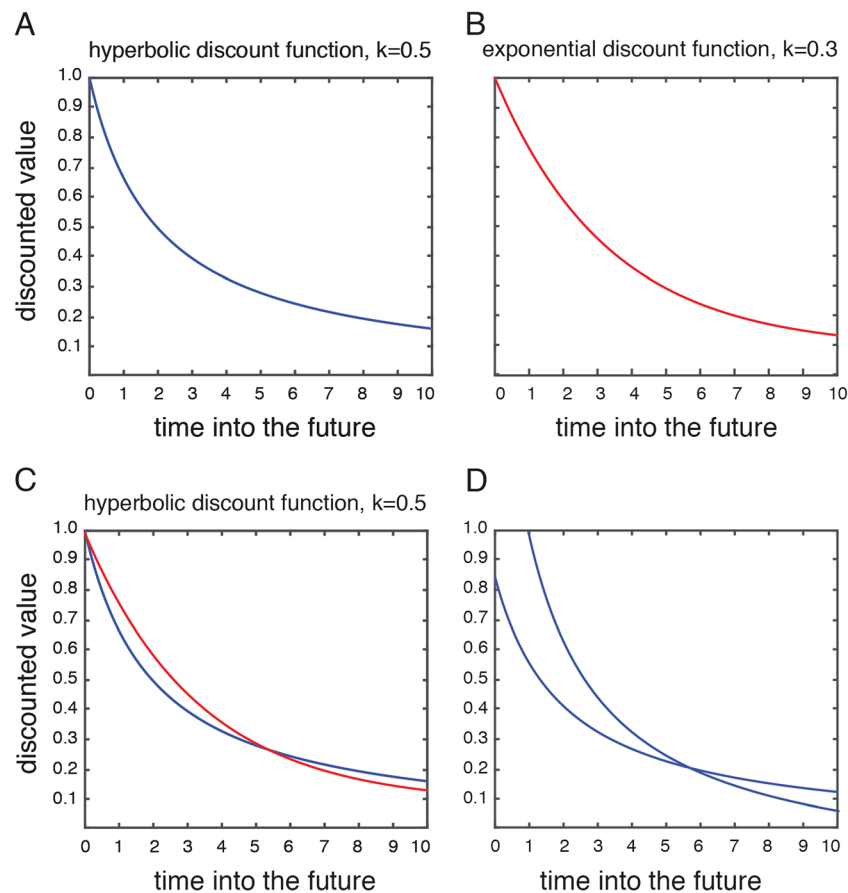


Fig. 2 (A) Hyperbolic discount curve. Discount curves are theoretical constructs that indicate the subjective value of a reward as a function of its delay into the future. (B) Exponential discount curve. (C) The curves from panels A and B overlaid. Even when they are similar, hyperbolic discount curves are steeper at the immediate present (and thus, value

declines faster) and then become flatter in the future (and thus, value declines more slowly). (D) Hyperbolic discount curves can cross over: Because the curve declines steeply early on, a delayed larger reward can drop in value fast enough that it crosses the value of the smaller sooner reward

(Stephens & Anderson, 2001). However, these costs are constant with time, so most economists believe that normative discount curves should decay at a constant proportional rate and describe an exponential decay curve.

Exponential curves, however, do not predict the preference reversals that pigeons and other animals show. The reason for this is that the decay of value has a constant proportionality. For this reason, Ainslie (1975) proposed a hyperbolic discount curve. This curve is somewhat similar in shape to an exponential curve (and indeed, aside from preference reversals, it is quite difficult in practice to distinguish these curves, even with large datasets; see Hwang et al., 2009, and cf. Fig. 2A and B). However, the hyperbolic curve has a steeper drop-off and a heavier tail. This means that, for hyperbolic discounters, value declines relatively quickly for short delays, and then more slowly for long delays.

In other words, for a hyperbolic discounter, time has more effect on choices when delays are short and less influence when delays are long. This in turn allows for a crossover in value, so that a smaller reward can be worth more than a large

one if the delay on the smaller one is short and the delay on the large one is long. Much subsequent research has replicated the preference reversal finding and the hyperbolic fit (Green & Myerson, 2004; Hwang et al., 2009; Kim, Hwang, & Lee, 2008; Louie & Glimcher, 2010; Mazur, 1987; Monterosso & Ainslie, 1999; Richards, Mitchell, De Wit, & Seiden, 1997). Note that other shapes can produce the same crossover effect; these include hyperboloid shapes (shapes resembling hyperbolas), beta–delta models, and even exponential curves if the coefficient varies with magnitude (Green & Myerson, 2004; McClure et al., 2004; Mitchell, Wilson, & Karalunas, 2015).

Although the hyperbolic and exponential curves look superficially similar, the distinction between the two is extremely important, because only the hyperbolic curve predicts a preference reversal (Fig. 2). In other words, the hyperbolic curve is a putative indicator of self-control processes; the exponential curve could arise from rate-maximizing, collection risk, and so on. Notably, humans also show hyperbolic discount curves in the human intertemporal choice task; thus, humans and animals show a similar characteristic shape in

their discount curves. This fact suggests to many scientists that the discount curves in both species arise from similar mental operations. It is proposed below that this similarity is coincidental.

Although humans' and pigeons' discount curves have similar shapes, they differ greatly in their size. Pigeons discount on the order of seconds, whereas humans discount on the order of months (Frederick et al., 2002; Jimura, Myerson, Hilgard, Braver, & Green, 2009; Mazur, 1987; Stevens & Stephens, 2009). At the time of Ainslie's (1975) review, the great variety of observed discount rates obscured that discrepancy somewhat.

Ainslie (1975) did consider the alternative possibility that animals' failures to choose the rate-maximizing reward in intertemporal choice tasks was a by-product of the animals' incomplete learning of the task. However, he dismissed this possibility, arguing that "the animals' behavior came to equilibrium over a large number of trials making incomplete learning unlikely" (Ainslie, 1975). The central argument here is that this possibility actually has some merit.

Evolutionary theory predicts neutral time preferences

Choice of the SS option in an intertemporal choice task is costly in the long run—just like poor self-control more generally. Although animals demonstrably choose SS options, this preference pattern runs against strong evolutionary pressure to use long-term rate-maximizing strategies (MacArthur & Pianka, 1966; Stephens & Krebs, 1986). A broad survey reveals that for most animals performing intertemporal choice tasks, the reward value has a half-life of a few seconds (Hwang et al., 2009; Mazur, 1987; Stevens & Stephens, 2009). A few seconds is roughly how long it takes a monkey to grab a fruit off a tree and put it in its mouth.

Consider a hungry foraging monkey in a tree that has consumed the best fruit and should, according to foraging theory, now leave this tree and find a new, richer one. Even with an atypically low discount factor of $k = 0.1s^{-1}$, monkeys could never work up the motivation to climb a new tree across the field. If that action takes a minute, the discounted value of the food in the new tree will be less than 2 % the value of the present supply. Yet monkeys routinely travel much farther multiple times every day, and forage quite well. Longer-term behaviors, like caching of food or fighting for access to mates, are plainly impossible to reconcile with steep discounting. It is theoretically possible to justify various behaviors by endowing animals with a variety of discount factors, but this begs the question of how general intertemporal choice data are. Moreover, it would be difficult to think of natural contexts in which the steep discounting measured in the lab would

apply, thus raising doubts about the external validity of discount measures.

Behavioral ecologist David Stephens has said, understatedly, "it is hard to imagine plausible processes that could justify such severe discounting" (Stephens & Anderson, 2001). The disjunction between the observed discount factors and the fundamental facts of foraging behavior has been particularly salient to field biologists, ethologists, and foraging theorists, because they are in the habit of thinking in evolutionary terms (Kacelnik, 2003; Pavlic & Passino, 2010; Stephens & Anderson, 2001; Stephens, Kerr, & Fernández-Juricic, 2004). Because severe discounting ought to be strongly disfavored by natural selection, they are reluctant to accept at face value the notion that animals make such obviously poor choices in their daily lives.

One common explanation for steep discounting rates in the intertemporal choice task is that the animals are sensitive to collection risk or interruption risk (Green & Myerson, 1996; Henly et al., 2007; Kagel, Green, & Caraco, 1986; McNamara & Houston, 1987). The gist of this argument is that animals who wait for some food risk losing it all to a rival or to a prey escaping. Despite the pedigree of this idea, it seems unlikely that interruption risk plays a major role in laboratory intertemporal choice tasks. First, it is implausible to imagine much interruption risk in the very small number of seconds associated with discounting in most laboratory experiments (Henly et al., 2007; Stephens & Anderson, 2001; Stephens et al., 2004). Second, empirical testing of this idea has marshaled evidence opposing it (Henly et al., 2007). Third, collection risk cannot explain the hyperbolic shape of the discount curve; it predicts an exponential shape. Fourth, animals are found to readily adapt their foraging strategies to changing environmental conditions, making it unlikely that "hardwired" foraging strategies would be unresponsive to the lack of predation risk found in the lab, especially when those strategies are so costly (e.g., Blanchard, Wolfe, Vlaev, Winston, & Hayden, 2014; Giraldeau & Caraco 2000; Hayden, Pearson, & Platt, 2011; Sih & Christensen, 2001). Some behaviors are hardwired, but there is little reason to think that time preferences are so fixed.

External validity of the intertemporal choice task

No studies appear to have demonstrated the external validity of the discount rates measured from the intertemporal choice task in animals. On the contrary, whereas animals are impatient in intertemporal choices, they are quite patient in foraging tasks (Blanchard & Hayden, 2014; Giraldeau & Caraco, 2000; Hayden et al., 2011; Kacelnik, 1984; Mellgren, Misasi, & Brown, 1984; Stephens & Anderson, 2001; Stephens & Krebs, 1986; Sih & Christensen, 2001). This approximately optimal behavior in foraging contexts extends to naturally

observed behaviors, suggesting that laboratory foraging tasks and not intertemporal choice tasks may measure the true time preferences (Stephens, Brown, & Ydenberg, 2007).

In particular, the discount factors derived from intertemporal choice tasks are poor at predicting how animals will behave in foraging contexts such as the patch-leaving task. In the patch-leaving task, animals choose between staying and leaving (in practice, these options can just be two levers or saccade targets). Staying provides a small reward whose value decreases each time it is chosen. Leaving produces a long delay and no reward, but resets the value of the staying reward to its initial state. Imagine a bear hunting fish in a lake; as the fish deplete, the rate declines. Leaving the lake imposes a time cost (called the *travel time*) but leads to a new lake fully stocked with fish. The critical feature of this task for the present discussion is that, very much like the intertemporal choice task, it involves a choice between two options that vary in reward size and delay.

In their seminal study, Stephens and Anderson (2001) directly compared blue jays' time preferences in tasks that have a standard intertemporal choice structure and a patch-leaving structure (Fig. 3). Specifically, when a bird receives its first small reward, it has a choice to make: It can leave and start the next trial, and get another reward after the first delay of the next trial (travel time), or it can stay and get a second reward

identical in size to the first after an additional delay (search time). A critical element of this paradigm is that the two tasks are structured so that the delays can match up, leaving as the only difference the foreground–background structure. By keeping the overall reinforcement rate constant, the authors can isolate the effects of structure. Despite the similarity of the two tasks, Stephens and Anderson found highly dissimilar time attitudes in the same subjects in the two tasks; subjects were much more patient in the foraging task. Notably, these results show that there is no need to carefully replicate natural environments to trigger qualitative changes between the discounting mode and foraging mode of behavior; instead, nearly optimal foraging appears to be robustly observable even in constrained artificial situations, just not in intertemporal choice tasks.

A similar pattern has been found in a monkey implementation of the patch-leaving task (Fig. 4; Blanchard & Hayden, 2015; Hayden et al., 2011). Monkeys are nearly optimal in patch-leaving, harvesting about 97 % of the reward that would be obtained by an optimal chooser with a matched stochastic bias (Hayden et al., 2011). These studies directly compared monkeys' behavior in a patch-leaving task and an intertemporal choice task in interleaved blocks (Blanchard & Hayden, 2015). Preferences in the foraging task were used to infer the monkeys' time preferences, assuming optimal foraging plus some amount of discounting. Monkeys were much more patient in the patch-leaving task than in the intertemporal choice task. Moreover, their behavior in the patch-leaving task was fit better by parameter-free optimal foraging equations (specifically, Charnov's marginal value equations: Charnov, 1976; Stephens & Krebs, 1986) than by a model based on hyperbolic discounting using parameters fit from the discounting task. These foraging equations assume no discounting at all (neither hyperbolic nor exponential), so they describe the behavior of an infinitely patient, or long-term rate-maximizing, decision-maker. In addition, the day-to-day variation in discount factor, as measured by the intertemporal choice task, did not predict the variation in time attitudes measured by the foraging choice task. These data indicate that discount factors measured in the intertemporal choice task fail to show external validity to a foraging task, even when the tasks are nearly the same in their superficial aspects.

When evaluating sequences of rewards, humans generally prefer increases in quality over time; this finding is difficult to reconcile with the idea of temporal discounting (Frederick et al., 2002; Loewenstein & Prelec, 1993). In a previous study from the present lab, we examined monkeys' preferences for sequences of five rewards (Blanchard et al., 2014). Monkeys were given a series of rewards whose sizes varied, and then they chose between repeating the sequence and taking a standard of known size (Fig. 5). The results showed that monkeys, like humans, prefer increasing sequences. Interpreting this result in the context of discounting model results requires

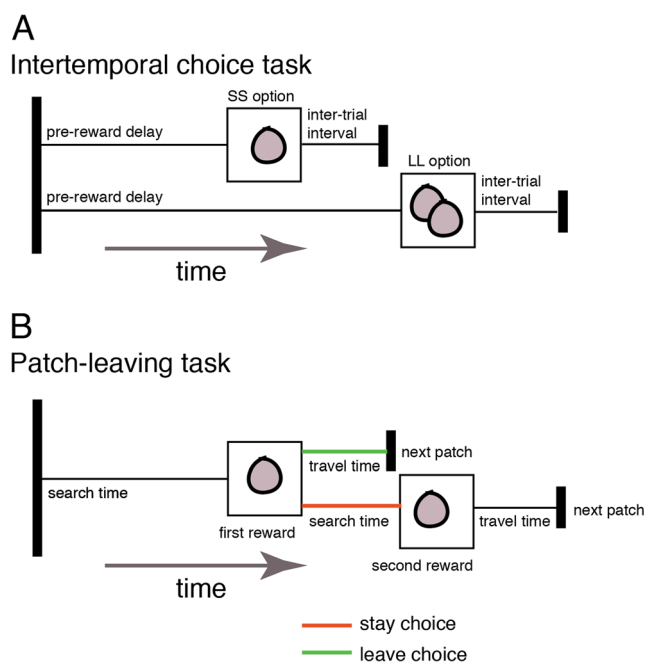


Fig. 3 Schematic of the tasks used in Stephens and Anderson (2001). (A) In the intertemporal choice task (which they called a *self-control task*), jays choose between a smaller reward offered sooner (SS; the rewards were food pellets) and a larger reward offered later (LL). Then an intertrial interval followed, and the next trial began. (B) On starting the trial, the jay got a reward after a fixed delay and then could leave the patch and start a new trial, or stay in the patch and get another small reward. The tasks are structured so that the overall reward rates are matched across conditions

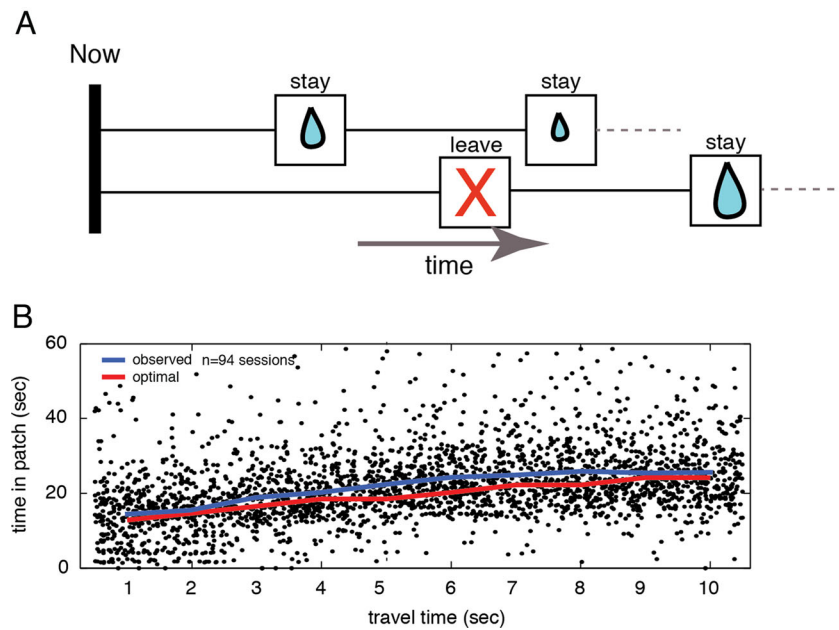


Fig. 4 (A) Schematic of the structure of a typical patch-leaving task, which produces nearly optimal (meaning, almost no discounting) preferences. On each trial, the subject chooses between a *stay* option, which gives a short delay and decrements a known reward, and a *leave* option, which yields a long delay and no reward, but resets the known

reward to a high value. (B) In this task, monkeys' performance (cool line) is nearly optimal (warmer line), although only stochastically so (black dots show the individual trial data). They show slight overstay, suggesting a weak but positive discounting (Hayden et al., 2011)

negative discount factors—something that has not previously been observed. Indeed, these monkeys showed significantly positive discount factors in a standard intertemporal choice task. (It is important to note that monkeys do not show a preference for increasing reward amounts in all contexts; Xu, Knight, & Kralik, 2011.) We speculated that monkeys' preferences in our task might reflect biases in the way that monkeys learn and remember sequences, rather than pure time preferences (Blanchard et al., 2014). If this speculation is true, then this dataset illustrates the importance of extraneous psychological biases (i.e., factors beyond time preference)

influencing what appear to be simple choices. Further discussion of the importance of learning biases in measures of time preferences appears below.

Intertemporal choices appear to lack external validity when compared to other self-control tasks. The *delay-of-gratification task* is a self-control task that is related to the well-validated marshmallow task (Beran & Evans, 2009; Evans & Beran, 2007; Mischel et al., 1972; Mischel et al., 1989; Reynolds, de Wit, & Richards, 2002). In the delay-of-gratification task, animals can wait for a large reward or, at any time while they wait, defect and obtain an immediate small reward. This opportunity

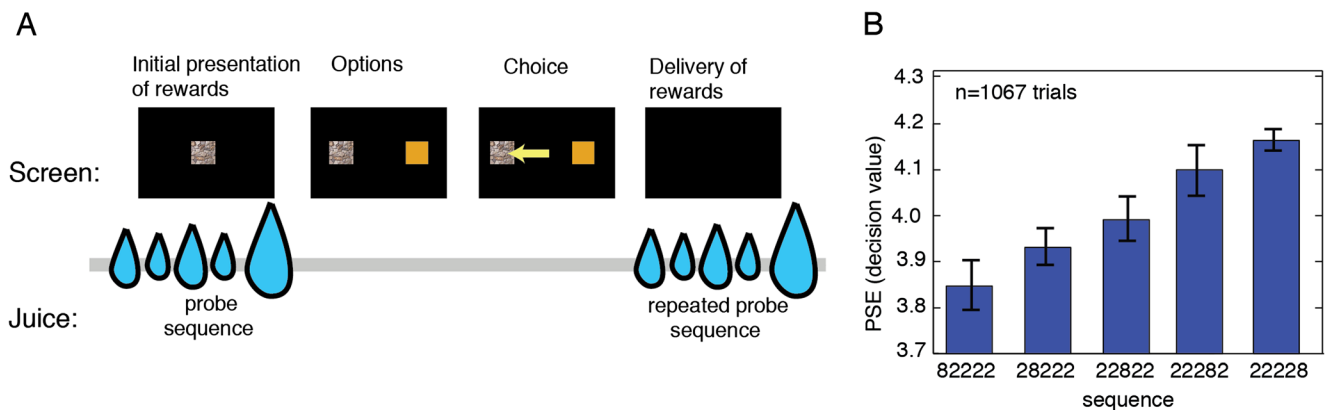


Fig. 5 (A) Task used to study preferences for sequences of rewards in monkeys. On each trial of the reward repeat task, monkeys first receive a probe—a series of rewards (juice drops)—and then choose whether to repeat the sequence or to have a standard; the ratio of choices for the standard indicates the subjective value of the probe sequence. (B)

Monkeys prefer sequences with larger rewards later in the sequence (e.g., a sequence of size 2, 2, 2, 2, and then 8 units of juice) over sequences of the same total value but with rewards early in the sequence (e.g., a sequence of 8, then 2, 2, 2, and 2 units of juice). The data are from Blanchard et al. (2014)

to give up is the key difference between the delay-of-gratification and intertemporal choice tasks: The task requires sustained commitment to a decision rather than just a one-time choice (cf. Rachlin, 2000). Animals appear to be quite good at delay of gratification (although their delay-of-gratification performance is context-dependent; Stevens et al., 2011). Monkeys, for example, will wait for a delayed large reward for tens of seconds or minutes (Anderson & Woolverton, 2003; Evans & Beran, 2007; Pelé, Micheletta, Uhlrich, Thierry, & Dufour, 2011). Good self-control is also observed in long-tailed macaques in exchange tasks (i.e., tasks in which they are given a small reward and, if they do not eat it, they can later exchange it for a larger one; Pelé, Dufour, Micheletta, & Thierry, 2010).

When intertemporal choice and delay of gratification are directly compared, animals have been found to be more patient in the delay-of-gratification than in the intertemporal choice task (Reynolds et al., 2002; Reynolds & Schiffbauer, 2005). Data from macaques agree with these findings about the difference between persistence and choice (Blanchard & Hayden, 2014). We devised a delay-of-gratification task, with delays of up to 10 s, based on the diet selection problem in foraging theory (Stephens & Krebs, 1986). Monkeys' average performance was nearly optimal and was fit very well by Krebs's "zero-one rule" equations with additional stochasticity. (Similar results were reported in Krebs, Erichsen, Webber, & Chamov, 1977.) The zero-one rule assumes no discounting. Further evidence for the questionable external validity of intertemporal choice tasks, albeit for very different reasons than the ones proposed below, can be found in other recent studies (Addessi et al., 2013; Paglieri, Addessi, Sbaffi, Tasselli, & Delfino, 2014; Paglieri et al., 2013).

One area in which intertemporal choice data do appear to have some external validity is in the contexts of drug manipulations and susceptibility (e.g., Dalley et al., 2007; Paine, Dringenberg, & Olmstead, 2003; Perry, Larson, German, Madden, & Carroll, 2005). In many contexts, drugs that increase impulsivity in humans have corresponding effects in animals, and animals that are impulsive prior to administration are more susceptible to drug abuse. This work provides some evidence for the validity of the intertemporal choice task; however, it is not clear how much. First, the effects of drugs in animals and humans are different, sometimes subtly, and often more so; more often still, dose equivalencies are unclear. Using the intertemporal choice task to justify a correspondence can lead to circular reasoning in the absence of other measures of impulsivity. More importantly, no matter the psychological basis of intertemporal choice behavior, it requires unbiased perception and evaluation of the time and reward, and an appropriate trade-off between them; drugs that affect temporal discounting generally also affect the perception and evaluation of both time and drugs (Heilbronner & Meck, 2014). Changes to these processes may spuriously appear to be changes in self-control or impulsivity.

Together, these data suggest that something about the intertemporal choice task motivates animals to choose the SS option even when it is very costly to do so. According to foraging theory, delays are important to the extent that they impose opportunity costs. For example, a polar bear waiting at one ice hole cannot wait at another one, cannot pursue alternative food sources, and cannot engage in mating. Thus, delays can be costly even if they pose no special self-control problem. For this reason, several scholars have suggested that animals may avoid delays in intertemporal choice because they wish to avoid the opportunity cost they carry, and not because they lack self-control (Bateson & Kacelnik, 1996; Blanchard et al., 2013; Kacelnik, 2003; Pavlic & Passino, 2010; Stephens & Anderson, 2001).

The problematic postreward buffer

Does something about the structure of the intertemporal choice task motivate choice of the SS in the absence of impulsive preferences? Several authors have pointed to the postreward delay (Bateson & Kacelnik, 1996; Blanchard et al., 2013; Pavlic & Passino, 2010; Pearson, Hayden, & Platt, 2010; Stephens & Anderson, 2001). In the animal intertemporal choice task, the overall trial length is generally kept constant by lengthening the postreward delay following SS choices (Fig. 1). This aspect of task design is critical. If the postreward delays were identical, the SS choice would lead more quickly to the next choice, and animals would be able to maximize long-term reward rate by choosing the SS option—and what appeared to be an impulsive preference for the SS would in fact be a prudent long-term rate-maximizing strategy (Fig. 6).

Just as essential as this differential postreward time buffer is the fact that subjects *understand* that their choice affects the postreward delay. If animals do not understand the structure of the task, and incorrectly believe that they can maximize long-term reward by choosing the SS, then SS choices cannot be




Fig. 6 Schematic illustrating why a perfectly patient decision-maker might choose a smaller sooner (SS) option. A choice of SS can lead to another choice right away; overall, the reward rate might match or even surpass that for the LL. Crucially, if a perfectly patient animal believed that this is the task structure, it would choose the SS even if the task used postreward buffers

interpreted as evidence of impulsivity. Few studies have checked whether animals understand this aspect of the task. As was noted above, Ainslie did consider the possibility, but reasoned that this was unlikely because of the large amount of training they received (Ainslie, 1975). Despite this, many studies have confirmed that even well-trained animals are either less attentive or completely inattentive to any delays occurring after the reward, as compared to delays occurring before it (Bateson & Kacelnik, 1996; Blanchard et al., 2013; Goldschmidt, Lattal, & Fantino, 1998; Green, Fisher, Perlow, & Sherman, 1981; Lea, 1979; Logue, Smith, & Rachlin, 1985; Mazur, 1987, 1989; Mazur & Logue, 1978; Mazur & Romano, 1992; Mazur, Snyderman, & Coe, 1985; Pearson et al., 2010; Snyderman, 1983; see also Gallistel & Gibbon 2000). These data suggest that animals do not fully consider the postreward delay in their choices, and thus it does not fulfill its purpose. Such down-weighting, which is distinct from a preference for immediate rewards, would motivate SS choices, and thus would inflate measures of discounting.

Inattention to the postreward delay may itself be a symptom of poor self-control, but it may also reflect a failure to understand the task (Logue et al., 1985). Several pieces of evidence suggest that a failure to understand the task—what Stephens calls *informational constraint*—is what actually accounts for a good deal of the effect (Blanchard et al., 2013; Kacelnik, 2003; Stephens, 2002; Stephens & Anderson, 2001; for a slightly different view, see Namboodiri, Mihalas, Marton, & Hussain Shuler, 2014; Namboodiri, Mihalas, & Shuler, 2014).

In another study, we measured monkeys' intertemporal preferences in a special version of the intertemporal choice task, in which we explicitly cued the postreward delay (Fig. 7A; Pearson et al., 2010). We had previously trained the monkeys to understand that the height of gray rectangles indicated delays and that selecting one rectangle caused the bar to shrink at a constant rate, like a fuse. The bottom of the bar had a colored band; its color indicated the size of the reward the monkey would get when the bar disappeared. Following training, we added a twist to the task: An additional segment of bar was sometimes displayed below the colored reward band (the "explicit cue"). The postreward delays were all the same in both versions.

On trials in which we used these explicit cues, monkeys adjusted their strategies and began to strongly prefer the LL option, as if they now realized that the total time cost of these options was the same as for the SS options. The two versions of the task were identical aside from the cueing, suggesting that the cue itself was the critical factor. These results are consistent with the hypothesis that even monkeys with copious training fail to understand the postreward delay when it is uncued, but can readily do so if it is cued.

In a subsequent experiment, we asked how simple this guidance could be (Fig. 7B; Blanchard et al., 2013). We found

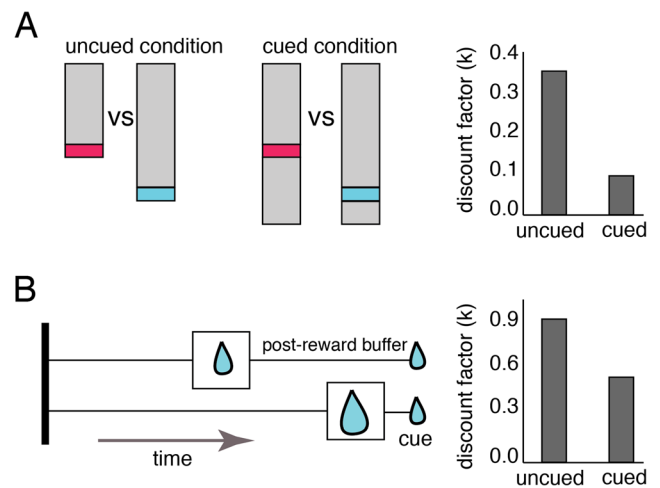


Fig. 7 (A) Schematic of cues used in a standard (i.e., uncued) and a cued condition of the intertemporal choice task. When postreward delays are cued, the observed discount factors drop significantly. (B) Schematic of the structure of a second-cue condition of the task. Even giving a second cue (a small reward) to highlight the delay reduced discount factors significantly

that even a small nudge, something as minimal as giving a small second reward at the end of the buffer (to make the timing more salient) was sufficient to reduce the measured discount rates. This finding suggests that exogenously driving attention to the end of the delay is sufficient to partially fix the bias toward SS, and implicates attention (and possibly its cousin, learning) in animals' poor intertemporal choices (see also Kacelnik, 2003). Note that for both manipulations, the discounting rates remained above zero. It is not yet clear whether additional cueing and training would reduce the measured rates further, or whether the new rates were the underlying true discount factors.

Animals' troubles with the postreward delays can come in two possible forms: (1) failure to link the duration of the buffer to their choices (failure of *contingent association*), and (2) underestimation of the buffer duration. We found evidence that both factors play roles in monkeys (Blanchard et al., 2013). First, we randomized the relationship between choices and the duration of postreward delays, thus severing the contingency between them. If monkeys fully understood the contingency, they would notice this change and shift their strategy toward the SS to exploit it; however, even with long training (over 10,000 trials per monkey), they failed to adjust their strategies in any measurable way. One simple explanation for this failure to adjust strategy is that monkeys did not notice the randomization because they already treated the postreward delays as if they were randomly related to their choices (Blanchard et al., 2013).

In another experiment from the same study, the postreward delays did not depend on choice but varied in long blocks (Fig. 8A). If monkeys entirely ignore postreward delays, they will not adjust their choices to changes in those delays.

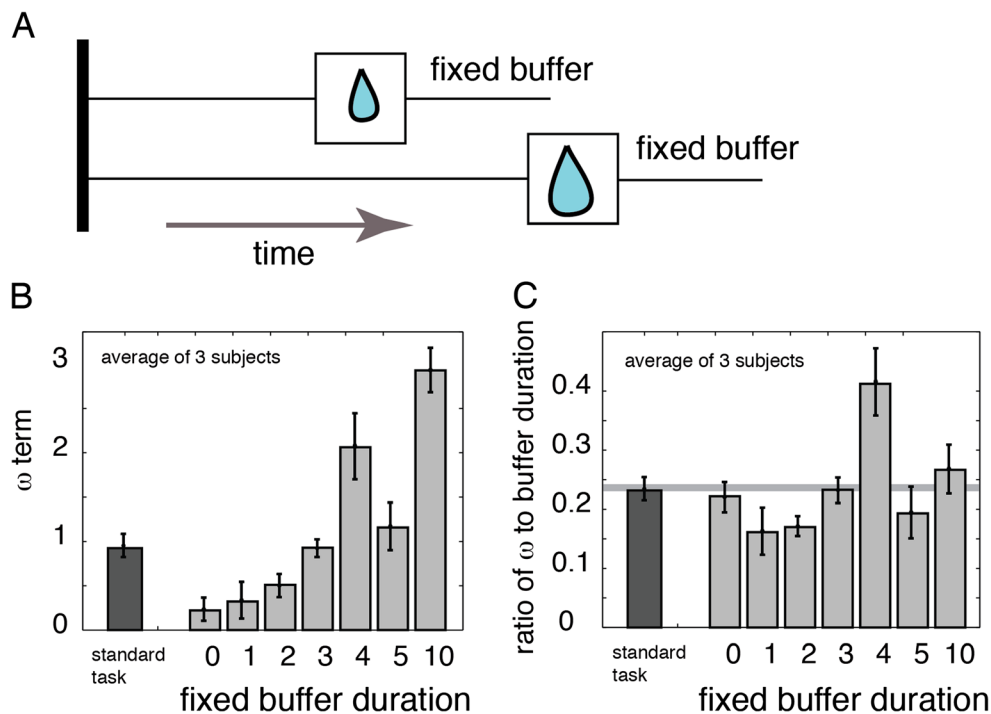


Fig. 8 (A) Schematic of a fixed-buffer version of the intertemporal choice task. (B) Plot of the best-fit ω term for each of seven buffer lengths (and in a standard task with an average 6-s buffer length). (C) An equivalent plot, but expressed as a ratio of ω to the actual buffer

length. In the fixed-buffer version of the task, monkeys behave as if they are optimal long-term choosers who estimate the duration of the postreward delay to be about 0.25 of its actual length

However, in this version monkeys did adjust their choices to the buffer length; thus, they were not entirely blind to its duration. Nonetheless, their behavior was not reward-maximizing; instead, they preferred the SS more than they should have. If we assume that monkeys sought to rate-maximize, but underestimated the postreward delays, we might then ask what duration they estimated the postreward delays to have (Fig. 8B). (The symbol ω is used to represent our estimate of the monkeys' internal estimate of postreward delay.) Monkeys acted as if they estimated the postreward delays to be about 25 % of their actual duration (Fig. 8C). The idea that animals suboptimally incorporate postreward delays into their strategies is supported by other studies, as well (Lea, 1979; Logue et al., 1985; Smethells & Reilly, 2014; Stephens & Dunlap, 2009).

Why is the postreward buffer difficult?

Perhaps it should not be surprising that animals fail to understand the timing and contingent nature of the postreward delay: Delays that occur after rewards present a difficult credit assignment problem (Dickinson, 1980; Kacelnik, 2003). Animals and other learners are generally better at learning about events that precede and predict rewards than about events that follow (and thus do not predict) rewards. This does not mean that animals cannot learn about postreward delays—

evidently they can when they are cued—just that they benefit from extra cueing. Indeed, there is empirical support for the idea that postreward time periods are less well understood than prereward time periods (Bateson & Kacelnik, 1996; Green et al., 1981; Snyderman, 1983).

Beyond the general difficulty of learning about events that occur after rewards, postreward buffered delays are biologically unlikely (Kacelnik, 2003; Stephens et al., 2004). In natural foraging situations, handling-time costs occur before consumption (Stephens & Krebs, 1986). To the extent that choice delays occur after the prey is consumed, they are generally independent of the size of the prey or grow with its size. For example, food must be digested, and during digestion it may be more difficult to pursue other prey. So larger prey may provide longer postreward delays. But in practice, larger rewards seldom lead to *shorter* postconsumption handling times.

Perhaps then, animals seek to maximize long-term reward rate (not just per-trial reward rate), but instead of doing the math, they use convenient rules of thumb that produce rate-maximizing behavior in many, but not all, situations (Stephens et al., 2004). A great deal of evidence suggests that animals do indeed use heuristics that are easier to compute in many situations (Hutchinson & Gigerenzer, 2005; Janetos & Cole, 1981; Stevens, 2010; Stevens & Stephens, 2009). Bateson and Kacelnik (1996) identified a heuristic that approximates rate-maximizing behavior in many situations

and that describes animals' choices quite well: value divided by prereward delay (which they call the *expectation of rates*, or EoR). EoR produces good behavior in many naturalistic contexts but produces poor behavior in the intertemporal choice task, because it ignores postreward delays (see also Pavlic & Passino, 2010; Stephens & Anderson, 2001; Stephens et al., 2004). The truth is likely to be more complex. Stephens found that subjects do prefer an option that offers a second reward after the first one, a violation of strict EoR (Stephens & Dunlap, 2009). We proposed the following alternative (Blanchard et al., 2013):

The heuristic model: $v_{\text{observed}} = \frac{v_{\text{actual}}}{\omega + D}$.

Rate is reward divided by time, so this equation just states that animals maximize rate, but are not particularly accurate when estimating delays after the reward. The term ω is defined as the animal's internal estimate of the postreward delay (including the intertrial interval), and in practice it is likely to be lower than the actual value (see also Namboodiri, Mihalas, & Shuler, 2014). The assumption is that ω is affected by learning and attention processes, which is why discount factors are so variable across and within studies. Because this equation has only one degree of freedom, it gives precisely the same risk of overfitting as the parameter k in the standard hyperbolic discounting equation.

Standard hyperbolic discounting equation: $v_{\text{discounted}} = \frac{v}{1 + k \cdot D}$.

Both the hyperbolic discounting equation and the heuristic model are essentially ratios of reward over time, with a fit-scaling parameter. Indeed, these two equations are, formally speaking, scaled versions of each other (for a mathematical proof, see the supplementary material of Blanchard et al., 2013). In other words, they can fit any intertemporal choice dataset equally well, and there is therefore no way that any set of intertemporal choice preference data alone can favor one or the other.

This mathematical equivalence is important: It reveals that the well-known hyperbolic discounting equation is a camouflaged version of short-term rate-maximizing. However, although the two are equivalent, they have entirely different psychological meanings. The conjecture here is that the success of the hyperbolic model in explaining animals' preferences reflects its coincidental similarity to animals' true choice functions, given by the heuristic model.

Benefits of the heuristic model

Our heuristic model is one of several possible ones that provide an impulsivity-free explanation for choice of the SS option in

an intertemporal choice task (Bateson & Kacelnik, 1996; Kacelnik, 2003; Namboodiri, Mihalas, & Shuler, 2014; Pavlic & Passino, 2010; Stephens et al., 2004). Certainly, the evidence does not conclusively favor our model. However, these models, as a class, do provide several appealing features when compared to the discounting model.

First, they explain why foraging tasks, persistence tasks, exchange tasks, and so on give a very different portrait of an animal's time attitudes than do intertemporal choice tasks. The alternative tasks generally have a foreground/background structure, in which postreward delays are fixed, brief, and do not depend on choice (Blanchard & Hayden, 2014; Hayden et al., 2007; Krebs et al., 1977; Pavlic & Passino, 2010; Stephens & Anderson, 2001; Stephens et al., 2004; Stephens & Krebs, 1986). In the intertemporal choice task, by contrast, the short delay on the SS leads to a need to buffer the postreward delay. This buffer poses a learning challenge.

Second, they explain why animal discounting studies produce discount factors that are several orders of magnitude larger than those produced in human studies. Human intertemporal choices are based on description, not experience, so there is no need for postreward delays and no possibility for confusion. (Of course, the human intertemporal choice task has its own problems, which are ignored here; these problems are reviewed in Frederick et al., 2002; Scholten & Read, 2010; Van den Bos & McClure, 2013; see also Stevens, 2015). This explanation is more consistent with evolutionary theory—which predicts patience—than is the idea that animals discount on the order of seconds. Speculatively, our model also provides a possible explanation for why even within the domain of human studies, discount factors are typically much larger for small primary rewards than for money and in foraging-like tasks than in intertemporal choice tasks (Bixter & Luhmann, 2013; Carter, Pedersen, & McCullough, 2015; Estle, Green, Myerson, & Holt, 2007; Hayden et al., 2007; Jimura et al., 2009; Luhmann, Chun, Yi, Lee, & Wang, 2008). Maybe humans do not really understand the rules when trials are strung out in sequence, or do understand the rules, but fail to apply them.

Taking this idea to its conclusion means that human and animal intertemporal choice tasks measure unrelated psychological processes. This discontinuity can explain a *third* striking difference between animal and human discounting patterns. Humans show a clear magnitude effect in their behavior (reduced discounting for larger stakes), whereas animals appear to have none (Green, Myerson, Holt, Slevin, & Estle, 2004; Green, Myerson, & McFadden, 1997; Jimura et al., 2009; Kirby & Maraković, 1996; Richards et al., 1997). This puzzling inconsistency makes more sense if the human and animal tasks are psychologically unrelated.

Fourth, these models provide an explanation for the hyperbolic shape of the discount curve observed in animal studies. In standard approaches, the hyperbolic curve is empirically

derived; no principle explains its shape. In contrast, the heuristic model provides a straightforward explanation: If animals seek to maximize rate (even if boundedly), they should compute the rate, or the ratio of reward to cycle duration (i.e., the sum of pre- and postreward delays; Blanchard et al., 2013). Thus, the argument here is that discounting is not intrinsically hyperbolic, but that hyperbolic discount curves are an artifact of the heuristic algorithm that animals use to generate behavior in intertemporal choice tasks. One implication is that the shared hyperbolic shape of the discount curves in animals and humans may be coincidental. (Of course, the evidence that humans have straightforward hyperbolic discount curves may be overstated, as well, although that topic is beyond the scope of the present work; Carter et al., 2015; Frederick et al., 2002; McClure et al., 2004; Schweighofer et al., 2006; Stevens, 2015; Van den Bos & McClure, 2013.)

Relatedly, these models provide an alternative explanation for why pigeons will precommit (Ainslie, 1974). The willingness of animals to select an option that reduces future choice is problematic for simple economic maximization theories of choice, but fits with the idea that animals recognize that they may struggle and lose the battle for self-control. However, the self-control interpretation of precommitment is problematic. It suggests that animals have a metacognitive awareness of changes in future motivational states that differ from the present one. However, it is not clear than most animals have such an ability (Naqshbandi & Roberts, 2006; Raby & Clayton, 2009; Roberts, 2002). Conversely, their behavior is readily explainable without recourse to prospection or metacognition through our rate-maximization model: Animals precommit early because it provides the larger average reward intake, but then switch their choices later when prereward delays shrink relative to total trial lengths.

Fifth and finally, the models explain the variability of the discount curves across studies, within studies, and even with subjects. For example, some studies have shown strikingly low discount rates in monkeys (Szalda-Petree, Craft, Martin, & Deditus-Island, 2004; Tobin et al., 1996). In our lab's own work on risky choice and intertemporal choice, we are often struck by how consistent a monkeys' risk attitudes are across very different paradigms, but how variable time preferences are from day to day within the same task (Blanchard et al., 2013; Heilbronner & Hayden, 2013; Heilbronner, Hayden, & Platt, 2011). To give one example, training procedure has a large effect on intertemporal choices (Mazur & Logue, 1978; Pearson et al., 2010), but has very little effect on risky choices (Heilbronner & Hayden, 2013). Our proposed model holds that attention and learning have a critical influence on the observed behavior in intertemporal choice (Kacelnik 2003; Pavlic & Passino, 2010; Stephens et al., 2004; see also Monterosso & Ainslie, 1999, who make a similar argument). Thus, factors that bias attention, such as cueing, or learning,

may strongly alter the observed discount factors (Mazur & Logue, 1978).

Suggestions for future research

These arguments should not be taken to support the idea that animals do not discount the future. Indeed, there are many reasons why animals should have positive discount factors, including collection risk, failure to prospect, chance of changing needs, the possibility of starvation, and even poor self-control. Nonetheless, the arguments here suggest that, to the extent that animals have stable time preferences, the intertemporal choice task produces only a biased measure thereof. Future studies will be needed to determine animals' true time preferences and how best to measure them.

A second goal for future studies will be to determine whether other approaches may be able to measure self-control and impulsivity better. One possibility would be to borrow tasks that are more structurally complex but that show clear external validity in humans, such as the BART task or the delay-of-gratification task (Evans & Beran, 2007; Lejuez et al., 2002; MacLean et al., 2014). Of course, translating a valid task from humans to animals is not always straightforward. These tasks also tend to involve multiple competing cognitive processes. However, if the goal is to measure self-control and impulsivity, then simplicity may be more dispensable than validity. Progress in this direction will also require an acknowledgement that self-control itself is multifactorial, and that different tasks will be needed to measure different aspects of self-control (Evenden, 1999).

In any case, one simple recommendation from this review is that time preferences should be studied using tasks with a foreground/background structure: accept-reject tasks. This set includes tasks based on the basic problems of foraging theory, including the patch-leaving problem, the diet selection problem, the central place foraging problem, and so forth (Stephens & Krebs, 1986). It also includes stopping problems and other classic optimization problems, such as the *k*-arm bandit problems, horizon problems, and change point detection problems (Pearson, Hayden, Raghavachari, & Platt, 2009; Wilson, Geana, White, Ludvig, & Cohen, 2014; Wilson, Nassar, & Gold, 2013). Indeed, it may also include variants of the intertemporal choice task in which the postreward delays are clearly cued (Pearson et al., 2010). Such tasks will provide a different viewpoint on animals' attitudes toward time, one that is more stable, more predictive across contexts, and more closely related to self-control and impulsivity.

Author Note Thanks to Daeyeol Lee, Sarah Heilbronner, Maya Wang, Jeff Stevens, John Pearson, and Tommy Blanchard for helpful discussions. This work was supported by an award from the John Templeton Foundation to the author.

References

- Abeyesinghe, S. M., Nicol, C. J., Hartnell, S. J., & Wathes, C. M. (2005). Can domestic fowl, *Gallus gallus domesticus*, show self-control? *Animal Behaviour*, 70, 1–11. doi:[10.1016/j.anbehav.2004.10.011](https://doi.org/10.1016/j.anbehav.2004.10.011)
- Addressi, E., Paglieri, F., Beran, M. J., Evans, T. A., Macchitella, L., De Petrillo, F., & Focaroli, V. (2013). Delay choice versus delay maintenance: Different measures of delayed gratification in capuchin monkeys (*Cebus apella*). *Journal of Comparative Psychology*, 127, 392–398.
- Ainslie, G. W. (1974). Impulse control in pigeons. *Journal of the Experimental Analysis of Behavior*, 21, 485–489.
- Ainslie, G. (1975). Specious reward: A behavioral theory of impulsiveness and impulse control. *Psychological Bulletin*, 82, 463–496. doi:[10.1037/h0076860](https://doi.org/10.1037/h0076860)
- Anderson, K. G., & Woolverton, W. L. (2003). Effects of dose and infusion delay on cocaine self-administration choice in rhesus monkeys. *Psychopharmacology*, 167, 424–430.
- Bateson, M., & Kacelnik, A. (1996). Rate currencies and the foraging starling: The fallacy of the averages revisited. *Behavioral Ecology*, 7, 341–352.
- Bechara, A. (2005). Decision making, impulse control and loss of willpower to resist drugs: A neurocognitive perspective. *Nature Neuroscience*, 8, 1458–1463.
- Beran, M. J., & Evans, T. A. (2009). Delay of gratification by chimpanzees (*Pan troglodytes*) in working and waiting situations. *Behavioural Processes*, 80, 177–181. doi:[10.1016/j.beproc.2008.11.008](https://doi.org/10.1016/j.beproc.2008.11.008)
- Bickel, W. K., & Marsch, L. A. (2001). Toward a behavioral economic understanding of drug dependence: Delay discounting processes. *Addiction*, 96, 73–86.
- Bixter, M. T., & Luhmann, C. C. (2013). Adaptive intertemporal preferences in foraging-style environments. *Frontiers in Neuroscience*, 7, 93. doi:[10.3389/fnins.2013.00093](https://doi.org/10.3389/fnins.2013.00093)
- Blanchard, T. C., & Hayden, B. Y. (2014). Neurons in dorsal anterior cingulate cortex signal postdecisional variables in a foraging task. *Journal of Neuroscience*, 34, 646–655.
- Blanchard, T. C., & Hayden, B. Y. (2015). Monkeys are more patient in a foraging task than in a standard intertemporal choice task. *PLoS ONE*, 10, e0117057. doi:[10.1371/journal.pone.0117057](https://doi.org/10.1371/journal.pone.0117057)
- Blanchard, T. C., Pearson, J. M., & Hayden, B. Y. (2013). Postreward delays and systematic biases in measures of animal temporal discounting. *Proceedings of the National Academy of Sciences*, 110, 15491–15496.
- Blanchard, T. C., Wolfe, L. S., Vlaev, I., Winston, J. S., & Hayden, B. Y. (2014). Biases in preferences for sequences of outcomes in monkeys. *Cognition*, 130, 289–299.
- Cardinal, R. N., Winstanley, C. A., Robbins, T. W., & Everitt, B. J. (2004). Limbic corticostriatal systems and delayed reinforcement. *Annals of the New York Academy of Sciences*, 1021, 33–50.
- Carter, E. C., Pedersen, E. J., & McCullough, M. E. (2015). Reassessing intertemporal choice: Human decision-making is more optimal in a foraging task than in a self-control task. *Frontiers in Decision Neuroscience*, 6, 95. doi:[10.3389/fpsyg.2015.00095](https://doi.org/10.3389/fpsyg.2015.00095)
- Chamov, E. L. (1976). Optimal foraging, the marginal value theorem. *Theoretical Population Biology*, 9, 129–136.
- Cheng, K. E. N., Peña, J., Porter, M. A., & Irwin, J. D. (2002). Self-control in honeybees. *Psychonomic Bulletin & Review*, 9, 259–263.
- Chung, S.-H., & Herrnstein, R. J. (1967). Choice and delay of reinforcement. *Journal of the Experimental Analysis of Behavior*, 10, 67–74.
- Dalley, J. W., Everitt, B. J., & Robbins, T. W. (2011). Impulsivity, compulsivity, and top-down cognitive control. *Neuron*, 69, 680–694. doi:[10.1016/j.neuron.2011.01.020](https://doi.org/10.1016/j.neuron.2011.01.020)
- Dalley, J. W., Fryer, T. D., Brichard, L., Robinson, E. S. J., Theobald, D. E. H., Lääne, K., ... Robbins, T. W. (2007). Nucleus accumbens D2/3 receptors predict trait impulsivity and cocaine reinforcement. *Science*, 315, 1267–1270. doi:[10.1126/science.1137073](https://doi.org/10.1126/science.1137073)
- Dickinson, A. (1980). *Contemporary animal learning theory*. Cambridge, UK: Cambridge University Press.
- Dufour, V., Pelé, M., Sterck, E. H. M., & Thierry, B. (2007). Chimpanzee (*Pan troglodytes*) anticipation of food return: Coping with waiting time in an exchange task. *Journal of Comparative Psychology*, 121, 145–155.
- Estle, S. J., Green, L., Myerson, J., & Holt, D. D. (2007). Discounting of monetary and directly consumable rewards. *Psychological Science*, 18, 58–64. doi:[10.1111/j.1467-9280.2007.01849.x](https://doi.org/10.1111/j.1467-9280.2007.01849.x)
- Evans, T. A., & Beran, M. J. (2007). Delay of gratification and delay maintenance by rhesus macaques (*Macaca mulatta*). *Journal of General Psychology*, 134, 199–216. doi:[10.3200/GENP.134.2.199-216](https://doi.org/10.3200/GENP.134.2.199-216)
- Evenden, J. L. (1999). Varieties of impulsivity. *Psychopharmacology*, 146, 348–361.
- Frederick, S., Loewenstein, G., & O'Donoghue, T. (2002). Time discounting and time preference: A critical review. *Journal of Economic Literature*, 40, 351–401.
- Gallistel, C. R., & Gibbon, J. (2000). Time, rate, and conditioning. *Psychological Review*, 107, 289–344. doi:[10.1037/0033-295X.107.2.289](https://doi.org/10.1037/0033-295X.107.2.289)
- Giraldeau, L. A., & Caraco, T. (2000). *Social foraging theory*. Princeton, NJ: Princeton University Press.
- Goldschmidt, J. N., Lattal, K. M., & Fantino, E. (1998). Context effects on choice. *Journal of the Experimental Analysis of Behavior*, 70, 301–320.
- Green, L., Fisher, E. B., Perlow, S., & Sherman, L. (1981). Preference reversal and self control: Choice as a function of reward amount and delay. *Behavioral Analysis Letters*, 1, 43–51.
- Green, L., & Myerson, J. (1996). Exponential versus hyperbolic discounting of delayed outcomes: Risk and waiting time. *Integrative and Comparative Biology*, 36, 496–505.
- Green, L., & Myerson, J. (2004). A discounting framework for choice with delayed and probabilistic rewards. *Psychological Bulletin*, 130, 769–792. doi:[10.1037/0033-2909.130.5.769](https://doi.org/10.1037/0033-2909.130.5.769)
- Green, L., Myerson, J., Holt, D. D., Slevin, J. R., & Estle, S. J. (2004). Discounting of delayed food rewards in pigeons and rats: Is there a magnitude effect? *Journal of the Experimental Analysis of Behavior*, 81, 39–50. doi:[10.1901/jeab.2004.81-39](https://doi.org/10.1901/jeab.2004.81-39)
- Green, L., Myerson, J., & McFadden, E. (1997). Rate of temporal discounting decreases with amount of reward. *Memory & Cognition*, 25, 715–723. doi:[10.3758/BF03211314](https://doi.org/10.3758/BF03211314)
- Hayden, B. Y., Parikh, P. C., Deaner, R. O., & Platt, M. L. (2007). Economic principles motivating social attention in humans. *Proceedings of the Royal Society B*, 274, 1751–1756.
- Hayden, B. Y., Pearson, J. M., & Platt, M. L. (2011). Neuronal basis of sequential foraging decisions in a patchy environment. *Nature Neuroscience*, 14, 933–939. doi:[10.1038/nn.2856](https://doi.org/10.1038/nn.2856)
- Hayden, B. Y., & Platt, M. L. (2007). Animal cognition: Great apes wait for grapes. *Current Biology*, 17, R922–R923. doi:[10.1016/j.cub.2007.08.061](https://doi.org/10.1016/j.cub.2007.08.061)
- Heilbrunner, S. R., & Hayden, B. Y. (2013). Contextual factors explain risk-seeking preferences in rhesus monkeys. *Frontiers in Neuroscience*, 7, 7. doi:[10.3389/fnins.2013.00007](https://doi.org/10.3389/fnins.2013.00007)
- Heilbrunner, S. R., Hayden, B. Y., & Platt, M. L. (2011). Decision salience signals in posterior cingulate cortex. *Frontiers in Neuroscience*, 5, 55. doi:[10.3389/fnins.2011.00055](https://doi.org/10.3389/fnins.2011.00055)
- Heilbrunner, S. R., & Meck, W. H. (2014). Dissociations between interval timing and intertemporal choice following administration of fluoxetine, cocaine, or methamphetamine. *Behavioural Processes*, 101, 123–134. doi:[10.1016/j.beproc.2013.09.013](https://doi.org/10.1016/j.beproc.2013.09.013)
- Henly, S. E., Ostdiek, A., Blackwell, E., Knutic, S., Dunlap, A. S., & Stephens, D. W. (2007). The discounting-by-interruptions

- hypothesis: Model and experiment. *Behavioral Ecology*, 19, 154–162. doi:[10.1093/beheco/arm110](https://doi.org/10.1093/beheco/arm110)
- Hull, C. (1943). *Principles of behavior*. New York, NY: Appleton-Century-Crofts.
- Hutchinson, J. M. C., & Gigerenzer, G. (2005). Simple heuristics and rules of thumb: Where psychologists and behavioural biologists might meet. *Behavioural Processes*, 69, 97–124. doi:[10.1016/j.beproc.2005.02.019](https://doi.org/10.1016/j.beproc.2005.02.019)
- Hwang, J., Kim, S., & Lee, D. (2009). Temporal discounting and intertemporal choice in rhesus monkeys. *Frontiers in Behavioral Neuroscience*, 3, 9. doi:[10.3389/neuro.08.009.2009](https://doi.org/10.3389/neuro.08.009.2009)
- Janetos, A. C., & Cole, B. J. (1981). Imperfectly optimal animals. *Behavioral Ecology and Sociobiology*, 9, 203–209.
- Jimura, K., Myerson, J., Hilgard, J., Braver, T. S., & Green, L. (2009). Are people really more patient than other animals? Evidence from human discounting of real liquid rewards. *Psychonomic Bulletin & Review*, 16, 1071–1075. doi:[10.3758/PBR.16.6.1071](https://doi.org/10.3758/PBR.16.6.1071)
- Kable, J. W., & Glimcher, P. W. (2007). The neural correlates of subjective value during intertemporal choice. *Nature Neuroscience*, 10, 1625–1633.
- Kacelnik, A. (1984). Central place foraging in starlings (*Sturnus vulgaris*) I. Patch residence time. *Journal of Animal Ecology*, 53, 283–299.
- Kacelnik, A. (2003). The evolution of patience. In G. Loewenstein, D. Read, & R. F. Baumeister (Eds.), *Time and decision: Economic and psychological perspectives on intertemporal choice* (pp. 115–138). New York, NY: Russell Sage.
- Kagel, J. H., Green, L., & Caraco, T. (1986). When foragers discount the future: Constraint or adaptation? *Animal Behavior*, 34, 271–283.
- Kim, S., Hwang, J., & Lee, D. (2008). Prefrontal coding of temporally discounted values during intertemporal choice. *Neuron*, 59, 161–172.
- Kirby, K. N., & Maraković, N. N. (1996). Delay-discounting probabilistic rewards: Rates decrease as amounts increase. *Psychonomic Bulletin & Review*, 3, 100–104. doi:[10.3758/BF03210748](https://doi.org/10.3758/BF03210748)
- Krebs, J. R., Erichsen, J. T., Webber, M. I., & Charnov, E. L. (1977). Optimal prey selection in the great tit (*Parus major*). *Animal Behaviour*, 25, 30–38. doi:[10.1016/0003-3472\(77\)90064-1](https://doi.org/10.1016/0003-3472(77)90064-1)
- Lea, S. E. G. (1979). Foraging and reinforcement schedules in the pigeon: Optimal and non-optimal aspects of choice. *Animal Behaviour*, 27, 875–886.
- Lejuez, C. W., Read, J. P., Kahler, C. W., Richards, J. B., Ramsey, S. E., Stuart, G. L., ... Brown, R. A. (2002). Evaluation of a behavioral measure of risk taking: The Balloon Analogue Risk Task (BART). *Journal of Experimental Psychology: Applied*, 8, 75–84. doi:[10.1037/1076-898X.8.2.75](https://doi.org/10.1037/1076-898X.8.2.75)
- Loewenstein, G. F., & Prelec, D. (1993). Preferences for sequences of outcomes. *Psychological Review*, 100, 91–108. doi:[10.1037/0033-295X.100.1.91](https://doi.org/10.1037/0033-295X.100.1.91)
- Logue, A. W., Smith, M. E., & Rachlin, H. (1985). Sensitivity of pigeons to preinforcer and postreinforcer delay. *Animal Learning & Behavior*, 13, 181–186.
- Louie, K., & Glimcher, P. W. (2010). Separating value from choice: Delay discounting activity in the lateral intraparietal area. *Journal of Neuroscience*, 30, 5498–5507. doi:[10.1523/JNEUROSCI.5742-09.2010](https://doi.org/10.1523/JNEUROSCI.5742-09.2010)
- Luhmann, C. C., Chun, M. M., Yi, D.-J., Lee, D., & Wang, X.-J. (2008). Neural dissociation of delay and uncertainty in intertemporal choice. *Journal of Neuroscience*, 28, 14459–14466. doi:[10.1523/JNEUROSCI.5058-08.2008](https://doi.org/10.1523/JNEUROSCI.5058-08.2008)
- MacArthur, R. H., & Pianka, E. R. (1966). On optimal use of a patchy environment. *American Naturalist*, 100, 603–609.
- MacLean, E. L., Hare, B., Nunn, C. L., Addessi, E., Amici, F., Anderson, R. C., ... Barnard, A. M. (2014). The evolution of self-control. *Proceedings of the National Academy of Sciences*, 111, E2140–E2148. doi:[10.1073/pnas.1325331111](https://doi.org/10.1073/pnas.1325331111)
- Mazur, J. E. (1987). An adjusting procedure for studying delayed reinforcement. In M. L. Commons, J. E. Mazur, J. A. Nevin, & H. Rachlin (Eds.), *Quantitative analyses of behavior: Vol. 5. The effect of delay and intervening events on reinforcement value*. Mahwah, NJ: Erlbaum.
- Mazur, J. E. (1989). Theories of probabilistic reinforcement. *Journal of the Experimental Analysis of Behavior*, 51, 87–99.
- Mazur, J. E., & Logue, A. W. (1978). Choice in a “self-control” paradigm: Effects of a fading procedure. *Journal of the Experimental Analysis of Behavior*, 30, 11–17.
- Mazur, J. E., & Romano, A. (1992). Choice with delayed and probabilistic reinforcers: Effects of variability, time between trials, and conditioned reinforcers. *Journal of the Experimental Analysis of Behavior*, 58, 513–525.
- Mazur, J. E., Snyderman, M., & Coe, D. (1985). Influences of delay and rate of reinforcement on discrete-trial choice. *Journal of Experimental Psychology: Animal Learning Processes*, 11, 565–575. doi:[10.1037/0097-7403.11.4.565](https://doi.org/10.1037/0097-7403.11.4.565)
- McClure, S. M., Ericson, K. M., Laibson, D. I., Loewenstein, G., & Cohen, J. D. (2007). Time discounting for primary rewards. *Journal of Neuroscience*, 27, 5796–5804.
- McClure, S. M., Laibson, D. I., Loewenstein, G., & Cohen, J. D. (2004). Separate neural systems value immediate and delayed monetary rewards. *Science*, 306, 503–507. doi:[10.1126/science.1100907](https://doi.org/10.1126/science.1100907)
- McDiarmid, C. G., & Rilling, M. E. (1965). Reinforcement delay and reinforcement rate as. *Psychonomic Science*, 2, 195–196.
- McNamara, J. M., & Houston, A. I. (1987). A general framework for understanding the effects of variability and interruptions on foraging behaviour. *Acta Biotheoretica*, 36, 3–22.
- Mellgren, R. L., Misasi, L., & Brown, S. W. (1984). Optimal foraging theory: Prey density and travel requirements in *Rattus norvegicus*. *Journal of Comparative Psychology*, 98, 142–143.
- Mischel, W., Ebbesen, E. B., & Zeiss, A. R. (1972). Cognitive and attentional mechanisms in delay of gratification. *Journal of Personality and Social Psychology*, 21, 204–218.
- Mischel, W., Shoda, Y., & Rodriguez, M. I. (1989). Delay of gratification in children. *Science*, 244, 933–938.
- Mitchell, S. H., Wilson, V. B., & Karalunas, S. L. (2015). Comparing hyperbolic, delay-amount sensitivity and present-bias models of delay discounting. *Behavioural Processes*, 114, 52–62. doi:[10.1016/j.beproc.2015.03.006](https://doi.org/10.1016/j.beproc.2015.03.006)
- Monterosso, J., & Ainslie, G. (1999). Beyond discounting: Possible experimental models of impulse control. *Psychopharmacology*, 146, 339–347. doi:[10.1007/PL00005480](https://doi.org/10.1007/PL00005480)
- Nambodiri, V. M. K., Mihalas, S., Marton, T. M., & Hussain Shuler, M. G. (2014a). A general theory of intertemporal decision-making and the perception of time. *Frontiers in Behavioral Neuroscience*, 8, 61. doi:[10.3389/fnbeh.2014.00061](https://doi.org/10.3389/fnbeh.2014.00061)
- Nambodiri, V. M. K., Mihalas, S., & Shuler, M. G. H. (2014b). Rationalizing decision-making: Understanding the cost and perception of time. *Timing and Time Perception Reviews*, 1, 1–40.
- Naqshbandi, M., & Roberts, W. A. (2006). Anticipation of future events in squirrel monkeys (*Saimiri sciureus*) and rats (*Rattus norvegicus*): Tests of the Bischof-Köhler hypothesis. *Journal of Comparative Psychology*, 120, 345–357. doi:[10.1037/0735-7036.120.4.345](https://doi.org/10.1037/0735-7036.120.4.345)
- Pagliari, F., Addessi, E., Sbaiffi, A., Tasselli, M. I., & Delfino, A. (2014). Is it patience or motivation? On motivational confounds in intertemporal choice tasks. *Journal of the Experimental Analysis of Behavior*, 103, 196–217.
- Pagliari, F., Focaroli, V., Bramlet, J., Tierno, V., McIntyre, J. M., Addessi, E., ... Beran, M. J. (2013). The hybrid delay task: Can capuchin monkeys (*Cebus apella*) sustain a delay after an initial choice to do so? *Behavioural Processes*, 94, 45–54.
- Paine, T. A., Dringenberg, H. C., & Olmstead, M. C. (2003). Effects of chronic cocaine on impulsivity: Relation to cortical serotonin mechanisms. *Behavioural Brain Research*, 147, 135–147.

- Pavlic, T. P., & Passino, K. M. (2010). When rate maximization is impulsive. *Behavioral Ecology and Sociobiology*, 64, 1255–1265. doi:10.1007/s00265-010-0940-1
- Pearson, J. M., Hayden, B. Y., Raghavachari, S., & Platt, M. L. (2009). Neurons in posterior cingulate cortex signal exploratory decisions in a dynamic multioption choice task. *Current Biology*, 19, 1532–1537. doi:10.1016/j.cub.2009.07.048
- Pearson, J. M., Hayden, B. Y., & Platt, M. L. (2010). Explicit information reduces discounting behavior in monkeys. *Frontiers in Psychology*, 1, 237. doi:10.3389/fpsyg.2010.00237
- Pelé, M., Dufour, V., Micheletta, J., & Thierry, B. (2010). Long-tailed macaques display unexpected waiting abilities in exchange tasks. *Animal Cognition*, 13, 263–271. doi:10.1007/s10071-009-0264-6
- Pelé, M., Micheletta, J., Uhlrich, P., Thierry, B., & Dufour, V. (2011). Delay maintenance in Tonkean macaques (*Macaca tonkeana*) and brown capuchin monkeys (*Cebus apella*). *International Journal of Primatology*, 32, 149–166. doi:10.1007/s10764-010-9446-y
- Perry, J. L., Larson, E. B., German, J. P., Madden, G. J., & Carroll, M. E. (2005). Impulsivity (delay discounting) as a predictor of acquisition of IV cocaine self-administration in female rats. *Psychopharmacology*, 178, 193–201.
- Raby, C. R., & Clayton, N. S. (2009). Prospective cognition in animals. *Behavioural Processes*, 80, 314–324.
- Rachlin, H. (2000). *The science of self-control*. Cambridge, MA: Harvard University Press.
- Rachlin, H., & Green, L. (1972). Commitment, choice, and self-control. *Journal of the Experimental Analysis of Behavior*, 17, 15–22.
- Renner, E. K. (1964). Delay of reinforcement: A historical review. *Psychological Bulletin*, 61, 341–361.
- Reynolds, B., de Wit, H., & Richards, J. B. (2002). Delay of gratification and delay discounting in rats. *Behavioural Processes*, 59, 157–168. doi:10.1016/S0376-6357(02)00088-8
- Reynolds, B., & Schiffbauer, R. (2005). Delay of gratification and delay discounting: A unifying feedback model of delay-related impulsive behavior. *Psychological Record*, 55, 439–460.
- Richards, J. B., Mitchell, S. H., De Wit, H., & Seiden, L. S. (1997). Determination of discount functions in rats with an adjusting-amount procedure. *Journal of the Experimental Analysis of Behavior*, 67, 353–366.
- Roberts, W. A. (2002). Are animals stuck in time? *Psychological Bulletin*, 128, 473–489. doi:10.1037/0033-2909.128.3.473
- Rosati, A., Stevens, J., Hare, B., & Hauser, M. (2007). The evolutionary origins of human patience: Temporal preferences in chimpanzees, bonobos, and human adults. *Current Biology*, 19, 1663–1668.
- Scholten, M., & Read, D. (2010). The psychology of intertemporal tradeoffs. *Psychological Review*, 117, 925–944. doi:10.1037/a0019619
- Schweighofer, N., Shishida, K., Han, C. E., Okamoto, Y., Tanaka, S. C., Yamawaki, S., & Doya, K. (2006). Humans can adopt optimal discounting strategy under real-time constraints. *PLoS Computational Biology*, 2, e152. doi:10.1371/journal.pcbi.0020152
- Siegel, E., & Rachlin, H. (1995). Soft commitment: Self-control achieved by response persistence. *Journal of the Experimental Analysis of Behavior*, 2, 117–128.
- Sih, A., & Christensen, B. (2001). Optimal diet theory: When does it work, and when and why does it fail? *Animal Behaviour*, 61, 379–390. doi:10.1006/anbe.2000.1592
- Smethells, J. R., & Reilly, M. P. (2014). Intertrial interval duration and impulsive choice. *Journal of the Experimental Analysis of Behavior*, 103, 153–165.
- Smith, D. V., Hayden, B. Y., Truong, T. K., Song, A. W., Platt, M. L., & Huettel, S. A. (2010). Distinct value signals in anterior and posterior ventromedial prefrontal cortex. *Journal of Neuroscience*, 30, 2490–2495. doi:10.1523/JNEUROSCI.3319-09.2010
- Snyderman, M. (1983). Delay and amount. *Journal of the Experimental Analysis of Behavior*, 39, 437–447.
- Stephens, D. W. (2002). Discrimination, discounting and impulsivity: A role for an informational constraint. *Philosophical Transactions of the Royal Society*, 357, 1527–1537.
- Stephens, D. W., & Anderson, D. (2001). The adaptive value of preference for immediacy: When shortsighted rules have farsighted consequences. *Behavioral Ecology*, 12, 330–339.
- Stephens, D. W., Brown, J. S., & Ydenberg, R. C. (2007). *Foraging: Behavior and ecology*. Chicago, IL: University of Chicago Press.
- Stephens, D. W., & Dunlap, A. S. (2009). Why do animals make better choices in patch-leaving problems? *Behavioural Processes*, 80, 252–260. doi:10.1016/j.beproc.2008.11.014
- Stephens, D. W., Kerr, B., & Fernández-Juricic, E. (2004). Impulsiveness without discounting: The ecological rationality hypothesis. *Proceedings of the Royal Society B*, 271, 2459–2465. doi:10.1098/rspb.2004.2871
- Stephens, D. W., & Krebs, J. R. (1986). *Foraging theory*. Princeton, NJ: Princeton University Press.
- Stevens, J. R. (2010). Rational decision making in primates: The bounded and the ecological. In M. L. Platt & A. A. Ghazanfar (Eds.), *Primate neuroethology* (pp. 97–116). Oxford, UK: Oxford University Press.
- Stevens, J. R. (2015). Intertemporal similarity: Discounting as a last resort. *Journal of Behavioral Decision Making*. doi:10.1002/bdm.1870.
- Stevens, J. R., & Mühlhoff, N. (2012). Intertemporal choice in lemurs. *Behavioural Processes*, 89, 121–127. doi:10.1016/j.beproc.2011.10.002
- Stevens, J., Rosati, A., Heilbronner, S. R., & Mühlhoff, N. (2011). Waiting for grapes: Expectancy and delayed gratification in bonobos. *International Journal of Comparative Psychology*, 24, 99–111.
- Stevens, J. R., & Stephens, D. W. (2009). The adaptive nature of impulsivity. In G. Madden & W. K. Bickel (Eds.), *Impulsivity: The behavioral and neurological science of discounting* (pp. 361–387). Washington, DC: APA Press.
- Szalda-Petree, A. D., Craft, B. B., Martin, L. M., & Deditus-Island, H. K. (2004). Self-control in rhesus macaques (*Macaca mulatta*) controlling for differential stimulus exposure. *Perceptual and Motor Skills*, 98, 141–146.
- Tangney, J. P., Baumeister, R. F., & Boone, A. L. (2004). High self-control predicts good adjustment, less pathology, better grades, and interpersonal success. *Journal of Personality*, 72, 271–314.
- Thaler, R. H., & Shefrin, H. M. (1981). An economic theory of self-control. *Journal of Political Economy*, 89, 392–406.
- Tobin, H., Logue, A. W., Chelonis, J. J., Ackerman, K. T., Brook, S., York, N., & Iii, J. G. M. A. Y. (1996). Self-control in the monkey *Macaca fascicularis*. *Animal Learning & Behavior*, 24, 168–174.
- Van den Bos, W., & McClure, S. (2013). Towards a general model of temporal discounting. *Journal of the Experimental Analysis of Behavior*, 99, 58–73.
- Wilson, R. C., Geana, A., White, J. M., Ludvig, E. A., & Cohen, J. D. (2014). Humans use directed and random exploration to solve the explore–exploit dilemma. *Journal of Experimental Psychology: General*, 143, 2074–2081. doi:10.1037/a0038199
- Wilson, R. C., Nassar, M. R., & Gold, J. I. (2013). A mixture of delta-rules approximation to Bayesian inference in change-point problems. *PLoS Computational Biology*, 9, e1003150. doi:10.1371/journal.pcbi.1003150
- Xu, E. R., Knight, E. J., & Kralik, J. D. (2011). Rhesus monkeys lack a consistent peak-end effect. *Quarterly Journal of Experimental Psychology*, 64, 2301–2315. doi:10.1080/17470218.2011.591936