The key idea is that hyperbolic discounting, by devaluing future rewards and punishments (negative rewards) proportionately to their delay, lets "utility theory move beyond its stalemate with cognitivism" (p. 38). So-called "dynamic inconsistency" is really a side effect of the fact that the discount curves for different rewards can cross: "people will naturally go for smaller, earlier over later, larger rewards. . . . Akrasia is just maximizing expected reward, discounted in highly bowed curves" (p. 39). In explaining how preferences can temporarily reverse, this simple model eliminates the apparent mystery of how we can act against our better judgment and against our "true" interests. Ainslie contends that hyperbolic discounting can explain a host of phenomena, including impulsiveness, addiction, compulsion, ambivalence, procrastination, and back-sliding. It can explain "the irony of smart people doing stupid things or having to outsmart themselves in order not to" (p. 27), by adopting "personal rules," cultivating good work habits, and making commitments that increase the cost of yielding to temptation. Consideration of future rewards doesn't take us outside the realm of reward and require higher-level judgments. There is just the ongoing competition among the rewards themselves.

It may seem an exaggeration to treat a person's different values as autonomous agents engaged in intertemporal bargaining with one another. Ainslie himself recognizes that his seemingly schizophrenic model of "the self as a population" (p. 39) makes it puzzling how, in a dog-eat-dog world of competing bargaining agents, "a marketplace of hyperbolically discounted choices [could] ever come to look like a single individual" (p. 40) rather than a kennel. What most worries me, though, is something else: At a given time, the motivational force of a desire (drive, urge, goal, value, or whatever you want to call the members of this population) is not a function merely of the reward associated with it, even as adjusted for the odds of success and the cost of failure and as temporally discounted.

Here are three other factors that can contribute to the motivational force of a particular desire: (1) the frequency and persistence of the desire's coming to mind, (2) the desire's degree of insatiability, and (3) its resistibility to the second-order desire to get rid of it. Each of these factors can vary even as the perceived reward of what is desired stays the same. For example, (1) something you want but deem of little importance can be more zealously pursued simply because the thought of it keeps occurring to you and capturing your attention. Playing another video game when you are trying to finish writing a commentary does not seem all that important, but its urgency is enhanced just because the thought of doing it keeps occurring to you. This doesn't make it seem like a better thing to do. Rather, you think you had better do it. What is rewarding is not playing the game but eliminating the clamoring desire to play it. Even worse, sometimes (2) the thought of playing the game does not go away after you play it. You did what you wanted to do, but now you want to do it again, just as if you hadn't done it in the first place. It's not that you want it more but that you want it again. In such a predicament, you may desire to make this desire go away, but (3) try as you may, it keeps rearing its head, keeping you from concentrating on that commentary. Now you're back to square (1), and vulnerable to (2) and (3) all over again.

You could have the opposite problem, say with a long-term project that requires irregular but frequent attention over time. As much as you value the ultimate reward of cultivating your garden, for example, (1) the thought of doing even a little puttering does not occur as frequently as it should. Not only that, (2) when it does occur and you act accordingly, the mere satisfaction of doing a little something makes you feel as though you have made significant progress even though you haven't. And, to make matters worse, (3) the thought of cultivating your garden resists staying in mind even when you want it to.

It might seem that these three factors are reducible to the perceived size of the desired reward. However, to suppose that would confuse the assessed size of the desired reward with these three independent dimensions of strength of the desire itself.

I have three further worries about the explanatory depth and breadth of the hyperbolic discount model. First, it ignores the distinction between something's being desired because it is rewarding and something's being rewarding because it is desired. Second, this model does not explain the magnitudes assigned to particular rewards in the first place. And third, the fact, if it is a fact, that discounting is hyperbolic itself cries out for an explanation.

Hyperbolas and hyperbole: The free will problem remains

Bruce Bridgeman

Department of Psychology, University of California–Santa Cruz, Social Sciences 2, Santa Cruz, CA 95064. **bruceb@ucsc.edu**
http://psych.ucsc.edu/faculty/bruceb/

Abstract: Hyperbolic theories have the fatal flaw that because of their vertical asymptote they predict irresistible choice of immediate rewards, regardless of future contingencies. They work only for simple situations. Theories incorporating intermediate unconscious choices are more flexible, but are neither exponential nor hyperbolic in their predictions. They don't solve the free will paradox, which may be just a consistent illusion.

Will a hyperbolic theory of reward discounting solve the persistent problems of the role of will in governing behavior? Ainslie (2001) makes a case for that idea, but hyperbolic theories have problems of their own. The hyperbola is defined by two perpendicular asymptotes, in this case one at the baseline of zero reward value and the other at the time of reward. Inevitably, a hyperbolic theory must predict reward value approaching infinity as the function moves toward the moment of reward. Rewards should become irresistible, no matter what other conditions apply, when the moment of reward gets very close. The exponential curve, in contrast, always has a finite value at a given time. True, the hyperbola has a predictive advantage over most of its range, where it predicts lower reward value (more discounting of future rewards) than exponential models, but the choice of either conic section is more a matter of mathematical convenience than theory, because no reasonably developed theory of reward motivates either model.

Another problem with hyperbolic theories became clear to me when I applied the only hyperbolic model that I have developed during my career, an ideal-observer model that predicts reading rate at any rate of display flicker, as on a CRT monitor. It was a harrowing experience, because once the positions of the two asymptotes have been established (in this case at zero display frequency and reading speed at infinite frequency) there is only one free parameter left, a scaling parameter. It's not much to go on. The model worked well enough, though, enabling us to predict reading speed at one frequency with great precision, given only the reading speed at another frequency (Montegut et al. 1997).

The hyperbolic model of reading rate worked because it dealt with a low-level stimulus sampling issue, not with deep psychological issues of reward and choice. The hyperbola simply doesn't leave enough room to take complexities into account. Recognizing this, Ainslie proposes that reward-delay decisions involve a whole cascade of hidden intermediate decisions, each with its own hyperbolic function that is replaced by another hyperbolic function at the time of an intermediate decision. The resulting predictions of future reward value are neither exponential nor hyperbolic, but depend on the timings of the intermediate decisions. Each intermediate decision brings with it a new free parameter, making the model more predictive but less theoretically useful, because the number of free parameters expands faster than the number of predicted actions.

Ironically, the decision cascade idea highlights an essential paradox of linear decision theories – decisions are themselves nonlinearities, places where everything that has gone before is applied to a single binary choice: you either accept the reward or you

https://doi.org/10.1017/S0140525X05240116 Published online by Cambridge University Press

don't. After each binary choice, represented by the vertical lines in Ainslie's diagrams, the data that previously went into the choice are discarded, and the process starts again. This is the sense in which decisions are nonlinear. And it's difficult to make a linear theory of nonlinear behavior.

Even this expanded theory, however, fails to resolve the deepest paradox of reward theory, the issue of what seems like free will. Even if the intermediate decisions are driven by utility theory or by some other deterministic algorithm, we still do not know how that last decision gets made. Substituting intertemporal bargaining within an individual adds flexibility, but there still must be rules governing each bargaining agent. The paradox of free will remains, the two alternatives of choice determined either by some unknown influence or by uncaused action. In the first alternative, free will is an illusion, a feeling of free choice in an environment where the decision is in fact determined by unconscious information-processing in the brain. The second flies in the face of everything we know about the physical world. Indeterminacy from random or chaotic processes doesn't solve the problem, for it only adds a bit of noise to the reasoning process, whether conscious or unconscious. And noise is not the same as free will; even in introspection, we don't confuse lack of control with freedom. The only alternative left is the uncomfortable first alternative, that is, that free will is indeed an illusion, but since it is a consistent illusion, it is accepted as reality.

The idea that consistent illusions are perceived as reality has precedents from more prosaic, but better-defined, domains in perceptual research. Illusions can be defined as situations that change upon closer inspection, whereas reality remains the same upon closer inspection. Length illusions, such as the Müller-Lyer arrows-in versus arrows-out figure, for example, can be tested easily by measuring the two lines in question, or by superimposing them for direct comparison. The illusion becomes obvious, but even after years of experience the figures still appear to be of distorted lengths. Other illusions are more difficult to expose as illusions. The slopes of hills, for instance, are grossly overestimated by most people, who will go through their lives believing that the steepest streets in San Francisco are perhaps 45 degrees, when they are actually about 10 degrees. If no one corrects them on this, the 45 degrees is reality for them, with none of the conflict that pertains to illusions when they are exposed. Analogously, if the unconscious information-processing that goes into decision-making is never exposed, people can go through their lives believing that their thought processes are guided by free will, and never be confronted by the paradoxes of uncaused action or hidden determinism. Still, a theory that can help us to predict behavior in serious situations, such as addictions, is a pragmatic step forward.

ACKNOWLEDGMENT

Preparation of this commentary was supported by a Faculty Research Grant from the Academic Senate of the University of California, Santa Cruz.

Regret and the control of temporary preferences

Terry Connolly[a] and Jochen Reb[b]

[a]*Department of Management and Policy, The Eller College, University of Arizona, Tucson, AZ 85721;* [b]*Singapore Management University, Lee Kong Chian School of Business, Singapore 178899.* **connolly@email.arizona.edu jochenreb@smu.edu.sg**

Abstract: Regret is often symptomatic of the defective decisions associated with "temporary preference" problems. It may also help overcome these defects. Outcome regret can modify the relative utilities of different payoffs. Process regret can motivate search for better decision processes or trap-evading strategies. Heightened regret may thus be functional for control of these self-defeating choices.

In Ainslie's taxonomy of "temporary preference" problems, the defining feature of addictions is that "the imminent prospect of such activities is strongly rewarding but they're avoided if foreseen from a distance *and regretted afterward*" (Ainslie 2001, p. 49; italics added). In compulsions, similarly, "regret may still occur . . . and the person may even expect the regret while indulging in the behavior" (p. 50), but the behavior still persists. Regret, then, may accompany, or even define, the problematic behaviors. What Ainslie appears to overlook is the possibility that regret may help to control them.

The "temporary preference" problem involves, minimally, the integration of two payoffs, one of which arrives earlier than the other, and is thus prone to the immediacy or hyperbolic discount effect. Viewed from a suitable distance before or after the choice, the later option is superior. Close to the decision point the earlier option is (temporarily) more attractive, a phenomenon that, as Ainslie notes (p. 198), "can't be changed by insight per se." The addict surrenders to the overwhelming desire for a fix, and the compulsive for yet another hand-wash, despite a clear intellectual understanding that she would, on balance, prefer not to and that the temporary urge will pass. Intellection is overpowered by emotion, System 2 by System 1 (Kahneman 2003).

Although emotions commonly involve both cognition and feeling (Frijda 1988), regret does so to an unusually large extent (Landman 1993). Asked if we feel regret over our choice of job, spouse, or vacation, most of us would reply "Let me *think*." We comfort a friend torn by feelings of regret by offering consoling *thoughts*: "How could you have known?" "You made a careful choice, there's nothing to blame yourself for." This intimate interweaving of thought and feeling has made regret the variable of choice for decision theorists interested in emotions. Perhaps it has a role to play in the complete understanding of temporary preference problems.

It is useful to distinguish two targets of decision-related regret: (1) regret associated with the outcome of a decision, and (2) regret associated with the choice process itself (Connolly & Zeelenberg 2002). The former seems to be essentially a reference-point phenomenon, in which the value of an outcome is reduced (or, if positive, enhanced: e.g., rejoicing) by comparison with some reference point, commonly the outcome of a foregone alternative (see Bell 1982; Loomes & Sugden 1982; Mellers et al. 1999). The second seems to involve a stronger component of self-blame or remorse, and is tied to the feeling that the decision made or the process used in making it was, in retrospect, insufficiently justified. The two regrets may be compounded, as for a mother who feels both outcome regret at the sickness of her small child, and self-blame regret at not having thought more carefully about his medical care (Reb & Connolly 2005). The failed dieter regrets both the additional weight gained and the poor decision about the chocolate cake.

There is abundant evidence that anticipated regret can influence decisions in a variety of domains, including medical care (Connolly & Reb 2003), consumer decisions (Simonson 1992), and negotiations (Larrick & Boles 1995). Richard et al. (1996) report some success in one temporary preference context, curbing unsafe sexual behavior, by shifting time-frames and making regret salient. They asked their respondents about either their feelings *about* unsafe sex or the feelings they would anticipate *after having had* unsafe sex. Participants in the second condition reported "safer" behavioral expectations immediately, and less actual risky sexual behavior in the six months following the experiment. How, exactly, might such a manipulation of regret salience achieve this promising result?

Two mechanisms might be suggested, paralleling the two sorts of regret described above. One possibility is that the manipulation enhanced outcome regret associated with the smaller, sooner behavior (unsafe sex), lowering its payoff value, and/or increased the larger, later payoff value by adding a component of rejoicing. It is clear from the conventional portrayal of discounted payoff values (e.g., Ainslie 2001, Fig. 4, p. 63) that modestly lowering the ear-

https://doi.org/10.1017/S0140525X05240116 Published online by Cambridge University Press