rates of reinforcer expiry. When predators compete, prey that are not attacked quickly may not be attacked at all, and be left either to maintain the prey population, or to feed a competitor. Rational impulsiveness becomes irrational only in modern times, when banks and groceries ensure longevity of some reinforcers. Other reinforcers, such as sexual partners, maintain selection pressure for alacrity.

Smaller-sooner versus larger-later? Temporal inconsistencies may rarely occur naturally, but may be manufactured by the "smaller-sooner versus larger-later" (SSLL) experimental paradigm. SSLL appears to be designed to model the choices between temptations and virtues that modern human-regulated environments offer, yet it fails to capture the relevant features of the environment of nonhuman organisms – an environment similar to that in which human behavior presumably evolved. Foragers, for instance, sequentially exploit depleting food patches, continuously facing the choice between staying and leaving the currently exploited patch. Blue jays nearly maximize food intake when these contingencies are simulated in the laboratory, but not when facing an economically equivalent SSLL preparation (Stephens & Anderson 2001).

Whether shaped through natural selection or provoked through experimentation, hyperbolic discounting may be the expression of more general – and adaptive – behavioral processes. Let us consider two candidate processes: sensory/temporal discriminability, and Pavlovian conditioning.

Relative discriminability. Adding a constant delay to both reinforcers may turn an "impulsive" subject into a "self-controlled" one; it also makes both delays less discriminable, while the discriminability of amount remains unchanged. Human participants in temporal discounting experiments may be subject to an analogous – verbally mediated – phenomenon: Most humans can distinguish goods attainable with $10 versus $100, but not with $910 versus $1,000. The role of the discriminability of reinforcement delays and amounts on intertemporal choice is yet unknown, but the ubiquity of Weber's law suggests its potential significance.

Approach gradients? A key property of signals of reinforcement is that they become both conditioned reinforcers, or CRs, and conditioned stimuli, or CSs, eliciting approach (Hearst & Jenkins 1974). Is this the behavioral substrate of desire, of appetite? If so, then Ainslie's hyperbolic interests, Skinner's CRs, and Pavlov's CSs are the same entity. A theory of one is a theory of all. Ainslie offers the opportunity to embed the hoary study of CRs in the neomodern study of genetic algorithms and artificial life. This would reinvigorate conditioning theory, promising a much richer and more certain application to the human condition.

The irrationality of hyperbolic discounting may be an awkward by-product of the demands of (a) rapid cultural evolution on human behavior, and (b) experimental attempts to generate such distortions in nonhuman behavior. What would life be like if we discounted exponentially? Hyperbolae may remain the best way to wire *homo economicus* in this post-modern era of Internet foraging. Simulations, anyone? The winners in such round-robin Axelrod competitions will probably find themselves equipped with modules that owe more to Skinner than to Freud; but it would not be inappropriate to call the modules "Ainslies."

ACKNOWLEDGMENTS

Preparation of this commentary was supported by NSF IBN 0236821 and NIMH 1R01MH066860 grants.

On the coexistence of cognitivism and intertemporal bargaining

Keith E. Stanovich

Department of Human Development and Applied Psychology, University of Toronto, Toronto, Ontario M5S 1V6, Canada. **kstanovich@oise.utoronto.ca**
http://tortoise.oise.utoronto.ca/~kstanovich/index.html

Abstract: Although Ainslie rejects cognitivism as providing an explanation of willpower, a type of nonhomuncular cognitivism is hiding in his own proposal. The key mental mechanism of aggregating individual decisions (bundled reframings) involves representation and decoupling operations encompassed within the analytic system of dual-process mental architectures.

Like others (e.g., Dennett 2003), I am immensely impressed with Ainslie's subtle interpretation of the concept of will as intertemporal bargaining (Ainslie 2001). I would like to explore the idea of marrying his ideas of intertemporal bargaining with cognitivism in the form of dual-process theories of reasoning (Evans 2003; Kahneman & Frederick 2002; Sloman 1996; Stanovich 1999; 2004).

In his book, Ainslie is perhaps rightly scornful of traditional cognitivist views of the will that posited little more than a homunculus standing outside the emotions. But more contemporary dual-process accounts share more features with Ainslie's intertemporal bargaining view.

First, dual-process models are closer to multi-entity bargaining views than it may seem, because such models are really about a multitude of processes, not just two. What has been termed, perhaps misleadingly, System 1 (Kahneman & Frederick 2002; Stanovich 1999), is really a whole *set* of autonomous processes – processes given the term autonomous because: (1) their execution is rapid, (2) their execution is mandatory when the triggering stimuli are encountered, (3) they are not dependent on input from high-level control systems, and (4) they can operate in parallel.

In contrast to the autonomous set of systems stands the nonautonomous system(s) – sometimes going under the name of analytic processing – with the opposite set of properties. Analytic cognitive processes are serial, rule-based, often language-based, and computationally expensive. The analytic system carries out two critical processes. First, it is responsible for sustaining the decoupling of representations from the world so that cognitive simulations can be run which test the outcomes of imaginary actions (Currie & Ravenscroft 2002; Dienes & Perner 1999; Nichols & Stich 2003; Perner 1991). Secondly, the analytic system can override responses triggered by the autonomous system.

Hyperbolic discounting provides a wonderful and heretofore missing explanation of where the analytic system gets the motivational force to override the autonomous systems that are responsible for preference reversals in conflicts between short-term and long-term interests. The bundling of rewards extended over time turns hyperbolic curves into exponential ones (at least temporarily, subject, of course, to the many other subtle bargaining effects that Ainslie's book details so ingeniously). However, there is much reference in Ainslie's book to bundling choices together, to choosing on the basis of principle, and to aggregating choices into categories. Now this kind of talk hides a lot of cognitivism – but in this case I think it is the good kind of cognitivism and not the bad kind that Ainslie spurns at the beginning of his book. Bundles, categories, principles – these are representations, often formed using the quasi-linguistic representational formats of the analytic system. The bundles or principles represent an alternative framing of the situation (one that yields exponential discounting rather than hyperbolic discounting). If the bundled frame dominates, then the long-term interest wins out. If the unbundled, single-choice frame dominates, then the short-term interest wins out. This type of representational redescription (see Karmiloff-Smith 1992) is much facilitated by language (and the discrete categories it provides). Using linguistic rules to rapidly implement new goal hierarchies is a quintessential function of the nonautonomous analytic system.

Likewise, I suspect that metarepresentation (another unique

https://doi.org/10.1017/S0140525X05340118 Published online by Cambridge University Press

analytic system function not within the capability of autonomous systems) would be implicated if we fully unpacked the mechanisms behind bundling, aggregating, or implementing a rule. At some point before the rule is fully instantiated, the first-order preference for the short-term reward has undergone a critique (the person realizes that they prefer to prefer otherwise).

Finally, Ainslie is clear that the reward bundling process critically implicates a self-prediction process ("a big part of this picture is her expectation of how she'll choose at later times," p. 131). However, these self-predictions are cognitive representations used in simulations in what some investigators call the Possible World Box (PWB; Nichols & Stich 2003), and sustaining the PWB is cognitively effortful (Glenberg 1997; Stanovich 2001; 2004). Ainslie has given us an insightful explanation of the motivational force that sustains the PWB.

I suspect that such dual-process explanations generalize to situations and effects not involving intertemporal discounting (e.g., probability matching, conjunction fallacies, various non-Bayesian responses). I have argued (Stanovich 2004) that many fallacies in the heuristics and biases literature arise from conflicts between autonomous systems more geared to genetic fitness maximization (i.e., to subpersonal entities not maximizing at the level of the individual) and analytic system procedures geared toward personal utility maximization. Ainslie edges toward this view when he posits (p. 45) that hyperbolic discounting might have been fitness maximizing at the expense of the individual organism, and when he allows for bundling to be facilitated by "executive processes that find bright lines" (p. 99).

Ainslie's view is consistent with dual-process theorists (Evans & Over 2004; Stanovich 2004) who emphasize that the bundling of choices into rules is evolutionarily unprecedented ("certainly a novelty in nature . . . a rather artificial process unlikely to have arisen in lower animals"; Ainslie 2001, p. 146), but he parts company with most of us in his view that "we apparently used the same intelligence that created this impulse problem to find a way around it"; p. 146). Autonomous processes create the problem of making us prone to responses that serve ancient interests of genetic fitness but not current personal utility maximization, and they lack a host of features necessary to overcome the problem, including: the ability to decouple from primary representations, the ability to create possible worlds and run simulations in them, and the ability to metarepresent. These processes – necessary to perceive an alternative world that better serves personal long-term interests and to take actions that bring about that world – are fundamentally different from the ballistic and automatically triggered features of autonomous systems.

In short, Ainslie adds the motivational component missing in most dual-process theories; whereas the latter deal extensively with the linguistic bracketing, cognitive decoupling, and simulation operations that are the mechanisms lying underneath the bundling and aggregating of choices. I see great potential for enriching dual-process theory with the insights from Ainslie's very nuanced attempt at a scientifically valid elucidation of the ancient concept of willpower.

"To do or not to do?" Modeling the control of behavior

John D. Swain[a] and James E. Swain[b]

[a]*Physics Department, High Energy Group, Northeastern University, Boston, MA 02115-5000;* [b]*Child Study Center, Yale University School of Medicine, New Haven, CT 06520.* **john.swain@cern.ch** **james.swain@yale.edu**

Abstract: The author of this fascinating book explores the problem of decision-making. As a basis, he uses hyperbolic discounting theory to discuss many basic assumptions related to self-control. In an accessible conversational tone, he succeeds in capturing many current problems in decision science and presents a rational framework for further work.

Rewards in the far future tend to be valued less than rewards in the more immediate future. The way in which future value is decreased seems to constitute a "breakdown of will," in that choices are often made that seem not to favor the overall best outcome. Assignment of reduced value to a delayed reward is expressed in terms of "discounting," and arguments are made that rational discounting should be done according to an exponential function. Such a distribution is well known in essentially all the sciences and follows from expressions of the form "rate of change of V with respect to time is proportional to the value of V itself at that time." It embodies an assumption that there is no change in the form of the solution if time is delayed by some constant.

Ainslie (2001) points out the rather striking fact that people do not discount rewards according to a relationship of this sort.

> If I asked a roomful of people to imagine that they've won a contest and can choose between a certified cheque for $100 that they can cash immediately and a postdated certified check for $200 that they can't cash for three years, more than half of the people usually say they would rather have the $100 now. If I then asked what about $100 in six years versus $200 in nine years virtually everyone picks the $200. But this is the same choice seen at six years' distance. (p. 33)

An exponential discounting function is incompatible with this behavior, so clearly, some other sort of discounting must take place favoring early rewards over later ones more strongly. The beauty of this sort of quantitative measure of reward and time is that the discounting function can actually be computed (or, more properly, estimated) from experiments. The result is that the data are well described, not by an exponential discounting, but by a hyperbolic one, that is, the perceived value at time t from now falls off as $1/t$.

Such a hyperbolic discounting has significant qualitative differences both at early and late times. At very early times, it in fact explodes, requiring a modification to render it finite, so that $1/t$ must be replaced by an expression of the form $1/(t + a)$. It very strongly favors immediate rewards. At very late times, however, it falls off more slowly than an exponential decay, making very long-term rewards more appealing than they would be for an exponential discounting. These twin features, argues Ainslie, can go a long way towards explaining both the fact that we tend to favor immediate gratification over mid-term gratification, and yet we can also plan for events far into the future.

From the point of view of the physical sciences, it is interesting to examine the origins of the distributions in order to see why hyperbolic discounting might be favored. Any discounting formula must come from an expression of the form "rate of change of value is some function of value and time." Exponential discounting, as described earlier, assumes no special time (so the function referred to is independent of time), and it also assumes that there is no special value – the rate of discounting is proportional to the value already there. Hyperbolic discounting assumes, on the other hand, that discounting is proportional to the square of the value, and although it introduces no special time, the solution in fact favors the special time $t = 0$, or "now."

Under conditions of great stability, where there is no special time, and "now" is much the same as "tomorrow" or "next year," and there is no special urgency attached to a large versus small value, one might well expect a simple exponential discounting, and indeed, this is what one sees in finance with respect to interest rates. Constant interest rates reflect stable times, whereas interest rates that change with time represent deviations from strict exponential discounting.

In nature, for the most part, there is little reason to expect this sort of stability. Lifespan is certainly finite, and most people would have little interest in waiting for 1,000 years versus 1,003 for the experiment with the cashiers' cheques. Even more pressing is the evolutionarily critical ability to reproduce, and Ainslie does suggest that a gene-driven urge to somehow reproduce "now," for "tomorrow" we could be dead, would seem to motivate hyperbolic discounting. Ainslie also hints at the role of nonlinearities, and although he only begins the discussion of what could be a very wide

https://doi.org/10.1017/S0140525X05340118 Published online by Cambridge University Press