our understanding of a variety of psychiatric disorders. In the spirit of reflecting on the significance and scope of their research, I briefly develop the idea that their framework can also contribute to improving our understanding of the pervasive problem of procrastination.

Starting from the idea that addiction involves "the continued making of maladaptive choices, even in the face of the explicitly stated desire to make a different choice" (target article, sect. 1), Redish et al. seek to develop a unified framework for addiction by (1) focusing on research concerning action selection and decision making, and (2) identifying failure points in our decision-making system. As they suggest, this approach may be fruitful for understanding not just addiction, but a variety of psychiatric disorders. I suspect that they are correct, and I want to develop a somewhat different but related suggestion, namely, that their approach can contribute to an improved understanding of the pervasive problem of procrastination.

Although procrastination is more common than addiction, it can figure as a crucial obstacle to realizing intentions to quit engaging in harmful addictive behavior. This fits neatly with the plausible conception of procrastination according to which it involves putting off an action that one should, given one's ends and information, perform promptly.

Even more so than addiction, which is still popularly cast as, at least in part, the product of powerful cravings that disable agents from acting voluntarily and in accordance with their decisions, procrastination is generally assumed to be the product of voluntary choices, and so the failures-in-decision-making approach that Redish et al. employ in their work seems particularly appropriate with respect to understanding procrastination. What better place to look for an understanding of self-defeating but voluntary delays than in research on failure points in our decision-making system?

The most established model of procrastination connects procrastination to problematic discounting processes (O'Donoghue & Rabin 1999a; 1999b; 2001), which is one of the vulnerabilities that Redish et al. discuss in their work. Like other animals, humans seem to discount future utility in a way that sometimes prompts preference reversals (Ainslie 2001; Kirby & Herrnstein 1995; Millar & Navarick 1984; Solnick et al. 1980). This can result in an agent's voluntary acting in a way that he or she planned against and will come to regret. The agent may, for example, keep making exceptions to his or her ongoing plan to cut down on indulgent purchases in order to save money for retirement. Discounting-induced preference reversals can thus foster procrastination.

A recent, complementary model of procrastination focuses on another vulnerability, one that is not directly discussed by Redish et al. but fits very well with their approach, namely, our vulnerability to intransitive preferences (Andreou 2007). Intransitive preferences (where, in particular, one cannot rank a set of options from most preferred to least preferred because there is a circularity in one's preferences) are often prompted by choice situations in which indulgences with individually negligible effects (such as smoking a cigarette) have momentous cumulative effects. Consider an agent who enjoys smoking but also values decent health. Someone in this situation may prefer, for all n, quitting after $n + 1$ cigarettes to quitting after n cigarettes, but also prefer quitting after a relatively low number of cigarettes to quitting after a very high number of cigarettes. This agent has intransitive preferences, and is vulnerable to intransitivity-induced procrastination (Andreou 2005).

Other interesting ideas concerning procrastination might fit comfortably within and be illuminated by Redish et al.'s framework. Consider, for example, the familiar idea that procrastination may be prompted by fear of failure, which may, in different cases, be the product of different vulnerabilities. For example, in some cases, fear of failure may result from the overvaluation of the expected value of stability; while in other cases, it may result from excessively (and perhaps obsessively) focusing on one possible outcome rather than appropriately distributing one's attention over the range of outcomes associated with a situation.

Consider next the idea that procrastination is strongly associated with the pursuit of "ephemeral pleasures" and "ephemeral chores" (Silver & Sabini 1981). Ephemeral pleasures and ephemeral chores are often more immediately gratifying or at least less aversive than the goal-directed actions that are called for by long-term projects. Moreover, ephemeral pleasures and ephemeral chores are often individually compatible with one's long-term projects, though they can accumulate in a way that interferes with these projects. These points suggest a connection between procrastination mediated by the pursuit of ephemeral pleasures and ephemeral chores, on the one hand, and problematic discounting processes or intransitive preferences, on the other.

The vulnerabilities I have been focusing on are vulnerabilities in the planning system, which is only one part of our decision-making system. As Redish et al. stress, problematic decisions can also result from vulnerabilities in the habit system or from vulnerabilities in the interaction of the planning system and the habit system. In the case of procrastination, it seems clear that planning-based vulnerabilities can foster habit-based vulnerabilities as well. If, for example, one's intransitive preferences prompt one to repeat individually negligible but cumulatively destructive actions, a habit-based vulnerability may flourish atop one's planning-based vulnerability. Soon enough, automatic indulgence will replace rationalized indulgence.

Relatedly, coping with procrastination often involves dealing with both planning-based vulnerabilities and habit-based vulnerabilities. Again, consider the agent whose intransitive preferences prompt intransitivity-induced procrastination. Once the agent's problematic indulgences are supported by habit as well, overcoming procrastination will involve (1) dealing with the planning system failure by, for example, adopting a plan that draws some bright lines in order to stop oneself from sliding down the slippery slope along a self-destructive path; and (2) overhauling one's habits so that acting accordingly becomes second nature.

In short, in addition to contributing to our understanding of addiction, Redish et al.'s failures-in-decision-making approach is suggestive with respect to the related, but more pervasive problem of procrastination. Indeed, it is probably less controversial to propose that the approach is well suited to providing a unified framework for procrastination than to propose that it is well suited to providing a unified framework for addiction.

Computing motivation: Incentive salience boosts of drug or appetite states

doi:10.1017/S0140525X08004767

Kent C. Berridge,[a] Jun Zhang,[a] and J. Wayne Aldridge[b]
[a]*Department of Psychology, University of Michigan, Ann Arbor, MI 48109;* [b]*Departments of Neurology and Psychology, University of Michigan, Ann Arbor, MI 48109.*
berridge@umich.edu
http://www-personal.umich.edu/~berridge/junz@umich.edu
http://www.lsa.umich.edu/psych/junz/jwaynea@umich.edu
http://sitemaker.umich.edu/aldridgelab/home

Abstract: Current computational models predict reward based solely on learning. Real motivation involves that but also more. Brain reward systems can dynamically generate incentive salience, by integrating prior learned values with even novel physiological states (e.g., natural appetites; drug-induced mesolimbic sensitization) to cause intense

desires that were themselves never learned. We hope future computational models may capture this too.

Redish et al. provide a valuable and comprehensive analysis of addiction models. They deserve gratitude from the field. Their sophisticated assessment of alternative models and explanatory mechanisms is admirably wide-ranging and thoughtful. In their fine scholarly effort, we would like to highlight what they note remains a significant gap unfilled by any available computational model. As the authors put it, "[C]omputational models of the planning system are insufficiently detailed to lead to specific predictions or explanations of the mechanisms by which outcomes are overvalued" (sect. 3.2.1, para. 4). We think this touches a central problem of dynamic motivation: the *generation of dynamic incentive salience motivation from static learned values*.

Current computational models that Redish et al. elegantly describe are powerful, but they are based purely on associations and memories. They act solely on what they know. In most reinforcement-based models, motivation is encoded as a part of an environmental state associated with the learned value of rewards, based on previous experiences. Motivational states may serve as occasion-setting contexts to modulate the value of rewards that have been previously experienced, and they may also modulate the unconditioned impact of a reward via alliesthesia. But the learned incentive value of the reward is never directly and dynamically modulated by the motivational state of the animal, without an additional learning process to intermediate.

However, evidence indicates the brain does something more when controlling desire. It dynamically generates motivation too, sometimes in surprising ways, by integrating static learned values with changing neurobiological states, some of which may never have before been experienced. Modulation of incentive value by new physiological/pharmacological states can be very potent – even the first time a relevant state occurs (Fudim 1978; Tindell et al. 2005a; 2005b; Zhang et al. 2005).

From the computational point of view, novel integrations of learned cues with new physiological states requires a kind of coupling that has not yet been satisfactorily modeled. Such coupling must connect associative values that have been previously learned and stabilized with current physiological/pharmacological states that can change from moment to moment. This falls into their *Vulnerability 4*, the one that we feel is particularly relevant to drug addiction. In particular, we concur with their assertion that new models are needed to address this point.

Addiction is recognized to usurp natural reward mechanisms, and even natural appetites offer dramatic demonstrations of dynamic generation of motivation (Berridge 2004; Toates 1986). Salt appetite is especially exemplary, because it can be produced as a novel state, as most humans today and most laboratory animals have never experienced a sodium deficiency. Salt appetite transforms the value of an intensely salty taste that normally tastes nasty (such as a NaCl solution that is triple the concentration of seawater). The intense saltiness becomes nice in a sodium-depleted body state, and the same triple-seawater becomes as hedonically positive as a sucrose solution (Tindell et al. 2006).

But what if a rat were not given the newly liked salt taste on its first day in a salt appetite state? What if it were instead given only a Pavlovian conditioned stimulus (CS) that had previously been paired with a salt unconditioned stimulus (UCS) when it was nasty? Cues for the previously nasty salt should have no incentive value according to most models. All the learning models described by Redish et al. make the same prediction here: the CS should elicit only negative reactions on the first trials of sodium deficiency. Any cached value of the stimulus-response (S-R) habit system obtained via a temporal difference mechanism must remain strongly negative, established by the previous pairings with punishing saltiness. Contextual knowledge does not yet exist about the potential goodness of salt in a sodium-deficient state. Even a cognitive-tree search mechanism has no way to infer the new value: its cognitive tree contains only memories of unpleasant saltiness. It lacks a branch for "liked" saltiness, at least until the rat is allowed to taste NaCl in its new physiological state.

Yet data from our lab and others show clearly that the incentive value of relevant cues can be modified on-the-fly based on homeostatic state. Indeed, we find the motivational value of the CS for salt is transformed to positive on the first day in the new state, even before saltiness is experienced: the cue becomes avidly approached and consumed, and it becomes able to fire limbic neurons like a cue for sweetness (Berridge & Schulkin 1989; Fudim 1978; Tindell et al. 2005b; 2006).

Bizarre as this reversal of cue valuation by a natural appetite may seem, nearly the same mechanism is exploited by drugs of abuse to cause addiction (Robinson & Berridge 1993; 2003). For example, other data from our laboratory show that drug-induced sensitization of mesolimbic systems, or acute amphetamine elevation of dopamine levels, causes certain relevant reward cues to dramatically become more "wanted," eliciting more incentive salience (Tindell et al. 2005b; Wyvell & Berridge 2001). The elevation in CS incentive value occurs before the UCS reward has ever been experienced while amphetamine was in the brain, or while the brain was in a sensitized state. In addicts, such sensitized "wanting" is posited to cause intense cue-triggered motivation for drugs that far outstrips their previously learned values (Robinson & Berridge 2003).

The implication of these examples is that desire is not reducible to memory alone. Brain mesocorticolimbic systems are designed to dynamically modulate previously learned incentive values, and they do not necessarily require new learning to do so. As Redish et al. have so admirably shown, current models give a fine account of how previously learned values of reward are recalled and coordinated to predict rewards based solely on previous experiences. We hope future models may also generate new motivation values of the sort we have described in order to more fully capture addiction.

ACKNOWLEDGMENT

The research mentioned in this commentary was supported by NIH grants DA-015188, DA-017752, and MH-63649.

Addiction science as a hedgehog and as a fox

doi:10.1017/S0140525X08004779

Warren K. Bickel and Richard Yi

Department of Psychiatry, Center for Addiction Research, University of Arkansas for Medical Sciences, Little Rock, AR 72205.

WBickel@uams.edu

RYi@uams.edu

Abstract: Redish et al. provide a significant advance in our understanding of addiction by showing that the various addictive processes are in fact all decision-making processes and each may undergird addiction. We propose means for identifying more central addiction processes. This recognition of the complexity of addiction followed by identification of more central processes would help guide the development of prevention and treatment.

In a classic essay, Sir Isaiah Berlin (1953) characterizes how individuals organize their subject matter by referring to the statement attributed to the ancient Greek poet Archilochus: "The fox know many things, but the hedgehog knows one big thing." At one extreme, some scientists may behave as foxes, treating their subject as a pluralistic conjunction of many diverse phenomena. At the other, some may approximate the hedgehog by viewing their subject as a monolith where one or a small number of phenomena play a central role in the subject of interest. The target article by Redish et al. proposes an interesting duality. On the one hand, they suggest that they are unifying the processes of addiction by proposing that the varied and