

Intertemporal Bargaining in Habit

George Ainslie 

Received: 9 September 2016 / Accepted: 22 November 2016
© Springer Science+Business Media Dordrecht (outside the USA) 2016

Abstract Lewis ascribes the stubborn persistence of addictions to habit, itself a normal process that does not imply lack of responsiveness to motivation. However, he suggests that more dynamic processes may be involved, for instance that “our recurrently focused brains inevitably self-organize.” Given hyperbolic delay discounting, a reward-seeking internal marketplace model describes two processes, also normal in themselves, that may give rise to the “deep attachment” to addictive activities that he describes: (1) People learn to interpret current choices as test cases for how they can expect to choose in the future, thus recruiting additional incentive (willpower) against a universal tendency to temporarily prefer smaller, sooner to larger, later rewards. However, when this incentive is not enough, the same interpretation creates incentive to abandon the failed area, leading to the abstinence violation effect and a localized weak will. (2) Normal human value does not come entirely, or even mainly, from expectation of external rewards, but is generated *endogenously* in imagination. Hyperbolic discounting provides an account of how we learn to cultivate the hedonic importance of occasions for endogenous reward by building appetite. In this account, expectations of the far future have to be rewarded endogenously if they are to be as

important as currently rewarded alternatives; and this importance is prone to collapse. Both will and hedonic importance are recursive and thus hard to study by controlled experiment, but do represent modelable, reward-based hypotheses about the dynamic nature of habit.

Keywords Habit · Hyperbolic discounting · Reward bundling · Willpower · Endogenous reward · Abstinence violation effect

In *Addiction and the Brain*, Lewis argues that addiction is a pattern of choice, rather than involuntary behavior imposed by a disease. He describes the neural changes in addicts’ brains that have been held to demonstrate the disease model, and points out that all changes in people’s behavior must have neural substrates. The changes seen in addiction just reflect “recurrent desire for a single goal,” and are reversible. He acknowledges the role of dispositional and environmental factors in making some people more liable to addiction, but points out that everyone has a tendency to overvalue imminent rewards. He ascribes addictions’ especial resistance to change to the “deep learning” of the relevant neural connections. He argues that this depth is the product of habit, so belief in a disease model of addiction is mistaken.

His points are all well taken. We are learning more about the great plasticity of neural connections. The fact that most addicts give up this behavior without therapy argues that, whatever neural changes may have

G. Ainslie (✉)
University of Cape Town, 151 Veterans Affairs Medical Center,
Coatesville, PA 19320, USA
e-mail: ga@picoeconomics.org

G. Ainslie
e-mail: George.Ainslie@va.gov

energized the habit, they have not eliminated the basic process of choice. Conversely, his point that people all have an innate susceptibility to addictive choices is supported by research that, if we define addiction broadly, half of Americans suffer from it [1]. He also avoids invoking an overarching “self,” whose choices are merely influenced, rather than determined, by reward. However, just to say that “addiction is an outcome of learning... that has been accelerated and/or entrenched through recurrent pursuit of highly attractive goals” is not to deal with the entrenchment process itself. Why do addictive activities form trenches, “like the ruts carved by rainwater in the garden,” while other activities shift when the reward for them shifts? Why, indeed, do addictive habits sometimes shift with great suddenness, either to sobriety or to relapse [2]? The word *habit* describes only their persistence itself, but it is a place to begin.

Three kinds of habits are familiar. All imply a pattern of reward, since choice without reward leads to response alternation or exploratory behavior rather than repetition [3]. Call them routine habits, good habits, and bad habits.

Routine habits are subroutines that you learn for navigating familiar paths to reward with a minimum of attention. Repeated rewarded behaviors get more and more efficient and require less and less attention. We learn many of these to form words, ride a bicycle, and drive to work while thinking of something else. This has been studied experimentally by seeing how long it takes subjects to change choices as contingencies of reward change, or the extent to which subjects ignore how initial choices affect opportunities at second or third choice points [4]. As Lewis describes, the development of habits is accompanied by a shift of neural activity in midbrain striatal areas from “planning” or “voluntary” to “habitual” systems [5]. A similar shift has been described from “goal-directed” or “model-based” to “model-free” systems [6].¹

Some authors propose routine habits as an explanation for why addictions persist in the face of contrary incentives (e.g. [7]). In making frequent choices to get small amounts of money in the laboratory, people with either addictions or obsessive-compulsive disorder (OCD) have been shown to respond to changed cues more slowly than normals [6]. However, this is not a

promising hypothesis. Although routinely habitual behaviors are sometimes called automatic or robotic, “mindless” would be a better word. It is easy to call off the subroutine when you have to stop at the grocery store on the way to work. Although in some animal experiments routine habits persist despite nonreward [8], and ablation of the brain center that switches to model-based choice can increase this persistence [9], these subjects have not been confronted with the strong disincentives that might be comparable to the costs of addiction; and in human experiments the effects of the relevant brain damage have been moderate to slight [10] (discussed in [11]).

Good habits (or hard habits) are those that take effort—keeping a diary every night or jogging every day or getting out of bed when the clock radio plays a certain theme every morning—and seem to require an excuse to break on a particular day, lest they be harder to begin again. The test of a good habit might be that you feel better off afterward, or even that you feel a slight rush of pleasure when an external circumstance prevents you from doing it today. This rush of pleasure is evidence that the habit is not something you simply prefer; but nevertheless abandoning or “breaking” the habit feels like a loss. The behaviors of sobriety may be routine habits for someone who is not tempted by alcohol, but good habits of great significance for an alcoholic, and a great loss if they are broken.

Bad habits (or lazy habits) are those that resist effort, either because the effort feels too great or is too hard to focus. Bad habits make you feel less well off afterwards. I have the bad habit of continuing to read news after I’ve finished the stories I care about. I believe it persists because the one type shades imperceptibly into the other. But I know I’ve done it when I get the stale feeling of wasted time. This is a trivial example; Lewis’ addictive habits lie at the other end of a continuum, as he points out. Although someone may call an activity that she actually prefers a bad habit—cracking her knuckles or putting her feet on the furniture, or even smoking or drinking—the term has motivational meaning only when a converse good habit has been broken. You make an effort and it is not enough. Bad habits are the junk heap of failed good habits.

Good and bad habits are familiar, but are not well modeled in the motivational literature. To begin with, they clearly have more to them than plain repetition. The observation that they often gain force with repetition needs explanation itself. Lewis seems unsure how much

¹ From ventral to dorsal striatum in rats, or the analogous dorsomedial to dorsolateral striatum in humans [3].

to rely simply on rewarded repetition as the explanation for why addictive habits resist new learning. In the book he introduces the intervening variable of compulsion: “The brain is certainly built to make any action, repeated enough times, into a compulsion. But the emotional heart of addiction—in a word, *desire*—makes compulsion inevitable, because unslaked desire is the springboard to repetition, and repetition is the key to compulsion” (p. 33). But is repetition sufficient for compulsion? He seems to be saying yes, but elsewhere he often hints at something more dynamic: “a brain that changes itself”; “habits link with other habits”; “bad habits self-organize like any other habits”; “habits that become self-perpetuating and self-stabilizing”; and perhaps most significantly, “Our *recurrently-focused* brains inevitably self-organize” (my italics).

Hyperbolic Discounting—The Basic Reward Pattern

I will suggest how our recurrently focused brains self-organize, although using mostly terms of motivation rather than neural activity. Lewis brings up delay discounting, which is well recognized; but a key property of delay discounting is that the inborn form is hyperbolic in shape, so value is (or starts out) inversely proportional to delay (Formula (1); [12]). People have shown hyperbolic delay discounting over periods of weeks and months for real rewards [13, 14] and over years and even decades for hypothetical rewards [15, 16]. Neural activity showing hyperbolically discounted value can be seen with brain imaging [17].

$$\text{Present value} = \text{Value}_0 / [1 + (k \times \text{Delay})] \quad (1)$$

where Value_0 = value if immediate and k is degree of impatience [18].

Hyperbolic discounting creates conflict between options that will be rewarded imminently and alternatives that will be rewarded more, but later.² Choices that have been shaped by later, larger (LL) rewards reflect the “objective” value of the rewards, but are apt to be overturned by the lure of smaller, sooner (SS) rewards as the pairs of alternative outcomes draw nearer (right-most pair in Fig. 1). The SS rewards could be said to be the basis of an *interest* that opposes an LL interest, by

² Even pains and negative emotions must compete for attention by a positive value up front, experienced as an urge [65].




Fig. 1 Summed expected values of 3 choices between an SS and an LL reward, made all at once: LL will be preferred at all choice points in advance of the first pair. If the choice were only between the last pair, the SS option would come to be preferred just before it was due

analogy to economic interests in markets. Trivially, any behaviors that are learned to get a reward could be called its interest, but the term is useful only where one interest has an incentive interfere with another one. The key property of these interests is that an SS interest does not die out if its goal is not chosen at a distance, because its relative influence will grow as it draws closer. The SS interest will not be extinguished even if it fails to get its goal after many tries, as long as it has a chance of succeeding. This conflict can be demonstrated with nonhuman animals in the laboratory [19], but is probably not significant in nature because their long term interests—to hoard, migrate, build dams—are served by inborn instincts that reward the necessary behaviors immediately. In effect, hoarding is fun. Humans, by contrast, have to learn their long term interests. Innate instincts clearly keep some influence—hoarding may still be fun— but these incentives are likely to be distractions from the kinds of goals that can be seen over months or years, and so to be the basis of short term interests.

Reports of hyperbolic discounting and its variants (for instance hyperboloid discounting, [20, 21]) have led to widespread awareness of precommitting behaviors that serve long term interests, which can be demonstrated in elementary form in nonhumans [22, 23] and with more sophistication in humans [20, 24]. Acting in your long term interest there are simple ways to forestall temptations: keep your attention away from them so they do not enter consideration, inhibit the relevant appetite or emotion before it gets too strong, find external influences (for instance Alcoholics Anonymous) or commitments (for instance disulfiram = Antabuse). The wishes of your family and friends are major incentives in normal self-control. But all of these methods have

serious limitations: Attention is hard to divert for long; appetites and emotions are rewarding in their own right; other people have their own agendas, and neither they nor physical commitments may be available when you need them.

The most effective impulse control has a large internal component. People do learn to choose consistently in practice, at least when dealing with money transactions, lest someone else who has learned this skill take advantage of their impulsiveness. But the learning is easier in some topics than in others, and some people become more skillful at it than others. Most importantly, its practice is not ahistorical, so someone's record at practicing it in the past affects her potential to do it now.

“Habit” from Intertemporal Bargaining The few existing models of internal self-control have suggested that a person either 1) has a separately motivated motivational faculty that exerts its “strength” [25] or 2) avoids weighing her incentives after the moment of change [26]. I have detailed the problems with these models elsewhere [12]. Alternatively, with the relatively flat tails of hyperbolic discount curves, just making a whole series of SS/LL choices at once gives a boost to the LL interest (Fig. 1), a phenomenon that consistency-maintaining (exponential) discount curves would not produce. The expected additive effect has been found in both nonhumans [27] and humans [28, 29].

To see how this additive effect sometimes commits a person to make a whole series— or *bundle*— of SS/LL choices just like her current one, we need to look at the relationship that hyperbolic curves create between her present and expected future selves. This could be described as limited warfare [30]. At each point she can expect herself to want the same long term outcome (say, being a sober person) but to indulge her appetite in the immediate future (whoop it up tonight). Or, more realistically in the case of someone deep in addiction, she always wants to stand the pain of withdrawal, anxiety, anhedonia etc. in the future, but to get relief right now. Assuming, as Lewis and I do, that she has no separately motivated, overarching faculty of self-control, the best evidence she gets of how she will decide when the choice comes up again is how she decides this time.

Once a person notices this connection, every current choice becomes a *test case* for how she can expect to make the whole bundle of similar choices in the future. If she stays up too late one night despite knowing that she will feel groggy the next day, she can expect to keep

doing it under similar circumstances. If she gets drunk despite the prospect of a hangover, she will probably do it again. To the extent that she notices how her current choice between an SS and LL reward predicts similar choices in the future, she creates a bundle of expectations that depend at least somewhat on the current choice—and which thus motivate that choice. Seeing her current choice as a test case creates a variant of repeated prisoner's dilemma (RPD) with her expected future selves, and her moves in this game over time establish *personal rules* for when she will count a choice of SS as a defection. Evidence that people see RPDs in SS/LL choices is reviewed in [31]. It includes how well the RPD fits common descriptions of willpower, how subjects behave in interpersonal RPDs, and how the RPD solves ostensibly paradoxical thought experiments about SS/LL choice. For instance, in Kavka's problem, a person is highly rewarded for *intending* to undergo an intensely aversive experience, but can back out and still get the reward once she has been found to have seriously intended it. (There are real life variants where such proof is possible.) Subjects' seemingly irrational feeling that they should not back out becomes rational if they see the need for serious intention as a recurring situation—an RPD—and expect damage to their ability to form such intentions if they defect.

When you have noticed how your current choice augurs for future choice, the cost of eating a serving of a forbidden food will only slightly be its effect on your weight or health, and will mostly be its damage to the credibility of your diet. The struggle between impulse and control now turns not so much on how close you get to a temptation—although this remains a factor—but on whether you expect a later self to see your current choice as a defection and thus have less reason in turn not to defect. The same logic governs a repeated choice of which the bad effects are distant: Each time you smoke you do not noticeably increase your risk of cancer, but you may decrease your expectation that you will stop. This will be especially true if you have been actively trying not to smoke.

Intertemporal bargaining lays down a history of choices that have turned out either well or badly, with the good choices becoming the basis of modifications in your personal rules. The logic is the same as for how court decisions over time have formed the English and American common laws. Relevant choices then have more import than the outcomes that are literally at stake,

leading an observer who does not have a theory of intertemporal bargaining to conclude that the extra motivation comes from “force of habit.” A person develops force of habit just as a society develops its habits, ways of doing things that are determined by its history of truces among competing interests—and the failures of those truces.

Limitations of Intertemporal Bargaining Probably only humans use our own past and present behaviors as cues to predict what we will do in the future. Even for humans it seems like a jury-rigged method for planning, not shaped specifically by the needs of long term consistency—in the way, for instance, that animals have evolved longer memories for flavors than for other information so as to identify a poison they ate hours before becoming sick (“bait-shyness” [32]). Self-control by identifying intertemporal prisoners’ dilemmas is a kludge. As we might expect, it has major limitations. Intertemporal bargains only compensate for the underlying hyperbolic discount function, rather than changing it. Recruiting incentive by interpreting current choices as test cases creates resolve that is often effective but is also brittle. In the face of strong temptations this tactic is apt to backfire:

- *Compulsiveness.* When current choices become less important in their own right than as test cases for bundles of expected future choices, it is harder to live in the here-and-now. A person becomes lawyerly, “inauthentic” in the existentialists’ terminology, less influenced by situations in themselves. Test cases for various temptations readily become seen as relevant to each other, so you might interpret a failure to exercise on schedule as less reason to believe that you will avoid smoking. Someone who depends heavily on this kind of self-control may feel burdened by how much is at stake in small decisions, a condition that is probably a better use of the word “compulsive” than just being strongly motivated (see [33]). Compulsiveness might seem to be a worthwhile price to pay for controlling a serious addiction, but having to avoid too many opportunities for reward may create an urge for relief that wins out at some point, either by finding multiple exceptions to the rule (rationalizations) or abandoning it altogether.

- *Circumscribed dyscontrol.* The greater stake assembled by bundling expected rewards together creates more incentive for LL choices, but also the risk of a greater loss after you detect a lapse. A recovering addict’s apparent collapse of will after a defection is the “abstinence violation effect” [34, 35], which has also been described after breach of a diet [36]. Having your will overwhelmed in one sphere has a predictive effect on other spheres as well. So a serious defeat creates an incentive to discriminate that temptation from other kinds and to give up trying in the damaged sphere, so as to preserve the credibility of your will elsewhere. Giving in becomes “automatic,” not in the sense of “mindless” as in routine habits, but because the LL alternative has such a long shot at winning. You feel that you *can’t* avoid getting drunk on New Year’s Eve, or maybe avoid it ever. The failed bargain has entrenched a bad habit.

With an abstinence violation effect that is limited by sickness, as in drinking or bulimic eating, the indulgence may stabilize in a binge pattern—consume until you have to stop, then stay abstinent for as long as you can. This may get called a habit, but the cause is a bargaining situation. With alcohol in particular, other impulsive interests such as for sexual excess, rage, or cruelty, may take advantage of binges to gain limited expression, licensed “while the alcohol is talking.”

- *Cognitive distortions.* The extra incentive to not lose a bundle of expectations is also an incentive not to see your present choice as a lapse. If you can’t manage to make it a believable exception to your personal rule, it may still be possible to skew your audit of your choices so as not to detect it—the Freudian defenses of denial and repression. But of course, the more you are able to use these evasions, the less will be the general credibility of the bundles of LL reward you hope for. Denial may be the most important mechanism in postponing the acknowledgment that an addictive behavior is out of control [37].

The latter two side effects can be expected to limit the power of the will, leaving long term interests at the mercy of delay discounting.

Apparent Habit Also Comes from Hedonic Importance

Addictions are especially apt to be ascribed to habit when they persist despite progressive increase in their prospective harm. Prospecion is often said to be blunted in addicts, but this is not an adequate explanation. An addiction often leaves a person's factual expectations intact, or at least no less accurate than other people's—She often expects to get cancer, HIV, jail time or death. As for predicting unhappiness, most people are poor at predicting how their actual moods will be only a few weeks hence. Subjects' predictions of future experience are distorted by being “essentialized” (fail to take account of detail), do not allow for fatigue, and do not imagine changing circumstances [38]. The short answer seems to be that addicts can do the same job at predicting future unhappiness as others, but care less about it. However, this observation confronts us with our lack of knowledge about how valuation of the future normally takes place.

A reward-based analysis in particular is complicated by the apparent steepness of the inborn discount rate. The data we have only suggest the nature of this process, but they still make it clear that a straightforward delay discounting model is not adequate. I propose a modification that incorporates the detachment of forward-looking motivation from objective evidence, while maintaining the assumption of strict determination by a single reward-comparing mechanism. However, I do abandon the behaviorist discipline that reward must come from external events.

When current comfort is at stake, its demands tend to overwhelm other motives. Most studies of delayed gratification deal with surplus value. Whether an experimental subject chooses \$50 now or \$100 in a year, her resources for sustaining a good mood over the following few hours will remain the same. The discount rate for actual comfort vs. discomfort is much steeper. The single digit annual discount rates that are adequate to sell people financial investments clearly apply only to surplus wealth—that beyond what is needed to sustain current hedonic tone. Four year old children, who can metarepresent others' beliefs [39] and tell distances to past events, still have difficulty waiting a few minutes to get two marshmallows instead of one [40]. Even adult subjects who strongly prefer six bits of food to two when both are immediate will usually not choose the six bits if they will be delayed two minutes—in marked contrast

to the same subjects' patience for sums of money [41]. In the extreme, the desire for continuing a crack cocaine high has often been enough to bankrupt someone, a modern equivalent of the Biblical Esau selling his birth-right for a mess of pottage. More mundanely, we have little tolerance for the boredom of a bad lecture or getting stuck in traffic, times when our usual supply of entertainment is interrupted. Volunteer subjects will often choose not to wait two minutes to quadruple their access to a video game [42] and are similarly impatient to get relief from unpleasant noise [43]. Playwrights notoriously have to design not just a plot that develops over two hours or so, but smart dialogue that provides payoffs from minute to minute—cf. the “flip value” required of novelists. At the discount rates implied by people's impatience with actual discomfort, the conventional exponential formula makes the value of an experience that is even a few days away infinitesimal. And yet people often deprive themselves seriously for distant goals, even resist torture—or give up addictions. How can we understand this contrast?

The relatively high tails of hyperbolic discount curves raise the value of distant events relative to what it would be with exponential curves, but this would still not be enough for events that are expected after days to compete with events that are expected after minutes [44]. Call the realm that is distant enough that expected options cannot compete, even with the help of bundles, the *far future*. Expectations for the far future have to bring into the present not only the picture of future events but also a significant share of their likely motivational impact. Beliefs about the risks of smoking, for instance, must create a significant fraction of the incentive created by facing the diagnosis of cancer if they are to compete with an immediate nicotine sensation. In effect, the discounted value of distant prospects has to be amplified to compete with attractions at hand. To compete in real time with these attractions—even to avoid the boredom of waiting for distant outcomes—the models that govern expectations of the future must pay off currently, that is, must be games that pay in the same league as video games or tasty snacks. Since the reward in these *scenarios* is evidently not discounted continuously from the far future, it must be *endogenous*—generated in imagination.

Far Future Expectations Depend on Endogenous Reward The motivational effect of scenarios varies with a mental process that may be independent of its

predictive accuracy. Conventional bookkeeping makes the reward value of any option depend ultimately on the extent to which it predicts hardwired rewards, a set not restricted to food, comfort, drugs, and sex but still innately configured, non-assignable [45, 46].³ A soft currency of imagined rewards might seem to need backing by a hard currency that is outside of a person's control, lest she short-circuit the reward process and divert reward from its adaptive purpose. But there are many concrete cases of eating food or gratifying sexual desire where we can consume rewards at will, and in those cases the constraint is our appetites for them. An analogous constraint, combined with a hyperbolic impatience to gratify the appetite before the best time, permits a model of endogenous rewards that stand on their own [47, 48]. It can be argued that the great majority of "secondary" rewards in a wealthy society do not predict hardwired primary rewards, but occur by an endogenous process.

According to this model, we set up endogenous rewards in imagination by the same process as in daydreaming, but controlling the hyperbolically discounted urge to cash them in early, as it were, by attaching them to infrequent and unpredictable occasions. This attachment is a betting process, the terms of which are enforced by the same recursive self-prediction as personal rules—The cost of cheating at solitaire, or of saying that a frightening movie is "only a story," is to reduce the stake of similar bets in the future. Likewise, the cost of laxness in testing the reality of far future expectations, or of not imagining the future at all, is to reduce our stake in predictive evidence—the occasions for endogenous reward in scenarios. This stake could be called *hedonic importance*, which we experience when we "give importance" to a project and "find it important" in turn. With endogenous reward, it is the aptness of occasions to pace appetite that determines their value as external goods. This aptness for being bet on is thus the counterpart of hardwired rewardingness. The aptness has its own determinants, both the probability pattern of the occasions and the importance that they have recursively developed [47, 48]. These determinants overlap with realism, but must ultimately serve the need of far future scenarios to be good stories in order to compete with imminent alternatives. Given the limitless potential of endogenous reward, the

importance of the occasions on which it hinges can cumulate to enormous values—spending thousands to spot a white tiger in nature, climbing Mount Everest, perhaps dying a martyr, or just obtaining portents of a full, satisfying life in the future. Conversely, investment in hardwired drug effects or in the challenges that build appetite for gambling or video games can undermine the hedonic importance of future prospects. Since this importance is not determined by discounting the objective value of future prospects, but by current imagination that is at most inspired by such value, it is subject to the same drift or collapse as the importance of sports teams or romantic quests. That is, the notion that future goods are "really" worth present sacrifice is not a perception but a construct: recursively determined hedonic importance.

The neural correlates of scenarios in self-control are just beginning to be visible: More patient choice has been found to be correlated with activity in the ventromedial prefrontal cortex (PFC) when subjects imagine future events [49]. Similarly, presenting subjects with words naming their own expected future events during an intertemporal choice task causes more patient choice, accompanied by activity in the ventromedial PFC and anterior cingulate gyrus (an "episodic imagery network") and increased coupling between this gyrus and the hippocampus [50]. These findings are tantalizing, but the motivational contingencies that induce and constrain the activity of imagination cannot themselves be seen.

Even activities that pay off in the very short term may be rewarded endogenously, as shown in some experiments with what is ostensibly secondary reward. In one, subjects chose between amounts of expected money ranging from \$0.01 to \$0.24 at delays ranging from 2 s to 64 s. The value of delayed amounts declined as a hyperbolic function of the delay, even though the money itself would not be delivered until the end of the experiment, and obviously could not be spent until even later [51]. As the authors pointed out, the substantial behavioral impact of these meaningless delays indicated that the reward announcements were valued for their own sake, like points in a video game. In a similar experiment using winnings signaled by pictures, the ventral striata of half the subjects showed discounting, but the brains of the other half did not, the latter half realistically reflecting the uniform delay to actual delivery [52]. The observation that half of subjects did not show discounting also suggests that this valuation did not

³ The "intrinsic" rewards that roboticists have begun to model are still inborn, "inherently interesting or enjoyable" [66, 67].

occur through a passive process such as secondary reward. As for the impact of these short term endogenous rewards on welfare, the addictive potential of immediate endogenous rewards from video games and other apps is just becoming apparent [53, 54].

Hyperbolic Discounting Is still Present in Scenarios

Although it is clearly impossible that far future prospects are discounted continuously over their expected delays, their value in scenarios still tends to be rated hyperbolically—for instance, in reported preferences on the order of \$4000 now vs. \$10,000 in ten years [55]. The same hyperbolic pattern is seen when subjects value future health [56], climate change [57], or procrastination [58]. However, the fact that the impatience factor k in subjects' discount formulas (Formula (1)) varies by hundreds-fold (e.g. [17]) indicates that they are not reporting the raw feel of the various outcomes; they are apparently imagining discount factors according to personal meanings. (The variation of k among members of a nonhuman species is in the single digits [59].) Another telling finding, if replicated, is that brain imaging subjects who evaluate amounts versus delays of money with respect to a future reference point have been reported to discount hyperbolically in relation to that point as if to the present. That is, if offered \$20 in 60 days vs. \$30–60 at 180 days, the shape of their curve to the 60 day point is the same as that of the curve to the present moment when they choose \$20 now vs. \$30–60 at 120 days, albeit with amplitudes reduced proportionately [60]. This result suggests that scenarios allow subjects to move a make-believe “now” to a future point, and thus that subjects in other discounting studies might feel similarly detached from the literal contingencies of delay. The hyperbolic discounting of outcomes in the far future seems to be somewhat constructed, rather than anchored in raw anticipation.

Certainly the hyperbolic value function shows up in studies that do not involve delay of payment, or even payment at all—For instance in the report that volunteer subjects value a hypothetical past prize as a hyperbolic function of its supposed recency [61], or where subjects report willingness to make altruistic gifts as a hyperbolic function of a wholly dimensionless attribute, “social distance” [62]. The hyperbolic shape seems to suggest itself to people's scenarios involving quantity. The important point is that the hyperbolic discounting reported in choices about the far future appears to be just a widespread feature of scenarios, not the presumably

innate psychophysical discount function itself (such as that demonstrated in nonhumans, for instance [19]). It seems to be learned readily, perhaps by simple analogy, but is nevertheless elective just as an exponential pattern of discounting is.⁴

Summary

A person chooses behaviors for three kinds of incentive, all of which are well known, and the first two of which are well studied: The behavior may be intrinsically rewarding; it may be instrumental in getting other rewards; or it may acquire hedonic importance through endogenous reward. A professional athlete is rewarded by physical sensations including endorphins as she performs her activities, by pay, and by the occasions for endogenous reward provided by events in the play. A fan watching the athlete is rewarded only by endogenous reward, which, however, can reach great intensity as her history of fandom increases the play's hedonic importance. This importance grows or shrinks by recursive self-prediction. Any of these incentives can lead to choices that reduce long term reward—impulses that sometimes become addictions. When the incentives are perceived as RPDs they are apt to give rise to intertemporal bargains (personal rules), which are also enforced by recursive self-prediction [48]. Subsequent defections may make the impulses worse, sometimes entrenching an addiction. Discriminating the two recursive mental processes—intertemporal bargaining and hedonic importance—by controlled experiment is probably impossible, though it might be parametrically modeled, Turing fashion. My point here is that addictive “habits” are not explained by repetition, but by something more dynamic.

⁴ People sometimes value even recent experiences by some means other than the summation of momentary values found over multiple trials with nonhumans [68]. In a pioneering project to observe directly how people evaluate visceral experiences, Kahneman and his co-workers found that “decision utility” is not the integral of momentary experiences [69]. That is, a subject's estimate of how painful a just-passed laboratory procedure was is the sum of her most extreme and most recent memories of it. So, for instance, adding a period of lesser discomfort at the end of a colonoscopy leads subjects to rate it less aversive. The subjects seem to have been sampling their component experiences rather than adding them up, and doing so without regard to their durations.

Rewarded repetition itself does not insulate an activity against extinction; in fact at the most elementary level, unpredictability is what does this [63]. I have argued here that Lewis' "deep trench" is based on two recursive phenomena:

- intertemporal bargains that have gone bad. Bundling choices together creates willpower but stabilizes failures of this willpower.
- the progressive withdrawal of hedonic importance from prospects in the far future, because of competition by both hardwired drug effects and challenging tasks that accompany—or create—addictive activity. As in Becker & Murphy's consumption capital [64], hedonic importance grows with practice, whether "playing the piano, baking bread, or smoking crack."

These two factors are apt to be how "synaptic networks are not only self-reinforcing but mutually reinforcing...to form a web that holds addiction in place." As Lewis points out, addictive choices are still motivated.

Acknowledgements This material is the result of work supported with resources and the use of facilities at the Department of Veterans Affairs Medical Center, Coatesville, PA, USA. The opinions expressed are not those of the Department of Veterans Affairs or of the US Government.

References

1. Sussman, Steve, Nadra Lisha, and Mark Griffiths. 2011. Prevalence of the addictions: a problem of the majority or the minority? *Evaluation & the Health Professions* 34: 3–56.
2. Miller, William R., and J. C'de Baca. 2001. *Quantum change: when epiphanies and sudden insights transform ordinary lives*. New York: Guilford.
3. Dolan, R.J., and P. Dayan. 2013. Goals and habits in the brain. *Neuron* 80(2): 312–325. p. 219
4. Kleinsorge, T. 1999. Response repetition benefits and costs. *Acta Psychologica* 103(3): 295–310.
5. Everitt, B.J., and T.W. Robbins. 2013. From the ventral to the dorsal striatum: devolving views of their roles in drug addiction. *Neuroscience & Biobehavioral Reviews* 37(9): 1946–1954.
6. Voon, V., K. Derbyshire, C. Rück, M.A. Irvine, Y. Worbe, J. Enander, L.R.N. Schreiber, C. Gillan, N.A. Fineberg, B.J. Sahakian, T.W. Robbins, N.A. Harrison, J. Wood, N.D. Daw, P. Dayan, P. Grant, and E.T. Bullmore. 2015. Disorders of compulsivity: a common bias towards learning habits. *Molecular Psychiatry* 20(3): 345–352.
7. Everitt, B.J., and T.W. Robbins. 2005. Neural systems of reinforcement for drug addiction: from actions to habits to compulsion. *Nature Neuroscience* 22: 3312–3320.
8. Yin, H.H., and B.J. Knowlton. 2004. Contributions of striatal subregions to place and response learning. *Learning and Memory* 11(4): 459–463.
9. Robbins, T.W., and B.J. Everitt. 2007. A role for mesencephalic dopamine in activation: commentary on Berridge (2006). *Psychopharmacology* 191: 433–437.
10. Fellows, Lesley K., and Martha J. Farah. 2005. Different underlying impairments in decision-making following ventromedial and dorsolateral frontal lobe damage in humans. *Cerebral Cortex* 15: 58–63.
11. Ainslie, George. 2016. Palpating the elephant; Current theories of addiction in the light of hyperbolic delay discounting. In *Addiction and choice: rethinking the relationship*, ed. Nick Heather and Gabriel Segal. 236, note#4. Oxford U.
12. Ainslie, George. 2012. Pure hyperbolic discount curves predict "eyes open" self-control. *Theory and Decision* 73: 3–34.
13. Johnson, Matthew W., and Warren K. Bickel. 2002. Within-subject comparison of real and hypothetical money rewards in delay discounting. *Journal of the Experimental Analysis of Behavior* 77: 129–146.
14. Kirby, Kris N. 1997. Bidding on the future: evidence against normative discounting of delayed rewards. *Journal of Experimental Psychology: General* 126: 54–70.
15. Green, Leonard, and Joel Myerson. 2004. A discounting framework for choice with delayed and probabilistic rewards. *Psychological Bulletin* 130: 769–792.
16. Cropper, Maureen L., Sema K. Aydede, and Paul R. Portney. 1992. Rates of time preference for saving lives. *American Economic Review* 82: 469–472.
17. Kable, Joseph W., and Paul W. Glimcher. 2007. The neural correlates of subjective value during intertemporal choice. *Nature Neuroscience* 10: 1625–1633.
18. Mazur, James E. 1987. An adjusting procedure for studying delayed reinforcement. In *Quantitative analyses of behavior V: the effect of delay and of intervening events on reinforcement value*, ed. M.L. Commons, J.E. Mazur, J.A. Nevin, and H. Rachlin. Hillsdale: Erlbaum.
19. Ainslie, George, and Richard J. Herrnstein. 1981. Preference reversal and delayed reinforcement. *Animal Learning and Behavior* 9: 476–482.
20. Laibson, David. 1997. Golden eggs and hyperbolic discounting. *Quarterly Journal of Economics* 62: 443–479.
21. McClure, S.M., D.L. Laibson, G. Loewenstein, and J.D. Cohen. 2004. The grasshopper and the ant: separate neural systems value immediate and delayed monetary rewards. *Science* 306: 503–507.
22. Ainslie, George. 1974. Impulse control in pigeons. *Journal of the Experimental Analysis of Behavior* 21: 485–489.
23. Deluty, M.Z., W.G. Whitehouse, M. Millitz, and P. Himeline. 1983. Self-control and commitment involving aversive events. *Behavioral Analysis Letters* 3: 213–219.
24. O'Donoghue, Ted, and Matthew Rabin. 1999. Doing it now or later. *The American Economic Review* 89(1): 103–124.
25. Baumeister, Roy F., Gailliot Matthew, C. Nathan DeWall, and Megan Oaten. 2006. Self-regulation and personality: how interventions increase regulatory success, and how

- depletion moderates the effects of traits on behavior. *Journal of Personality* 74: 1773–1801.
26. Fudenberg, D., and D. Levine. 2006. A dual-self model of impulse control. *American Economic Review* 96: 1449–1476.
 27. Ainslie, George, and John Monterosso. 2003. Building blocks of self-control: increased tolerance for delay with bundled rewards. *Journal of the Experimental Analysis of Behavior* 79: 83–94.
 28. Kirby, Kris N., and Barbarose Guastello. 2001. Making choices in anticipation of similar future choices can increase self-control. *Journal of Experimental Psychology: Applied* 7: 154–164.
 29. Hofmeyr, André, George Ainslie, Richard Charlton, and Don Ross. 2010. The relationship between addiction and reward bundling: an experiment comparing smokers and non-smokers. *Addiction* 106: 402–409.
 30. Schelling, Thomas C. 1960. *The strategy of conflict*, 53–80. Cambridge: Harvard University Press.
 31. Ainslie, George. 2001. *Breakdown of will*, 105–140. Cambridge: Cambridge University.
 32. Revusky, S., and J. Garcia. 1970. Learned associations over long delays. *Psychology of Learning and Motivation* 4: 1–84.
 33. Ainslie, George. 2016. Palpating the elephant; current theories of addiction in the light of hyperbolic delay discounting. In *Addiction and choice: rethinking the relationship*, ed. Nick Heather and Gabriel Segal, 236. Oxford: Oxford University.note#4
 34. Marlatt, G. Allen, and Judith R. Gordon. 1980. Determinants of relapse: implications for the maintenance of behavior change. In *Behavioral medicine: changing health lifestyles*, ed. Park O. Davidson and Sheena M. Davidson, 410–452. Oxford: Pergamon.
 35. Curry, S., G.A. Marlatt, and J.R. Gordon. 1987. Abstinence violation effect: validation of an attributional construct with smoking cessation. *Journal of Consulting and Clinical Psychology* 55: 145–149.
 36. Polivy, J., and C.P. Herman. 1985. Dieting and bingeing: a causal analysis. *American Psychologist* 40: 193–201.
 37. Pickard, Hanna, and Serge Ahmed. 2016. How do you know you have a drug problem? The role of knowledge of negative consequences in explaining drug choice in humans and rats. In *Addiction and choice: rethinking the relationship*, ed. Hanna Pickard and Serge H. Ahmed. London: Routledge.
 38. Gilbert, Daniel T., and Timothy D. Wilson. 2007. Prospection: experiencing the future. *Science* 317: 1351–1354.
 39. Sabbagh, M.A., and M.A. Callanan. 1998. Metarepresentation in action: 3-, 4-, and 5-year-olds' developing theories of mind in parent–child conversations. *Developmental Psychology* 34(3): 491.
 40. Mischel, H.N., and W. Mischel. 1983. The development of children's knowledge of self-control strategies. *Child Development* 54: 603–619.
 41. Rosati, A.G., J.R. Stevens, B. Hare, and M.D. Hauser. 2007. The evolutionary origins of human patience: temporal preferences in chimpanzees, bonobos, and human adults. *Current Biology* 17: 1663–1668.
 42. Millar, A., and D.J. Navarick. 1984. Self-control and choice in humans: effects of video game playing as a positive reinforcer. *Learning and Motivation* 15: 203–218.
 43. Navarick, D.J. 1986. Human impulsivity and choice: a challenge to traditional operant methodology. *Psychological Record* 36(3): 343–356.
 44. Ainslie, George. 2006. Motivation must Be momentary. In *Understanding choice, explaining behaviour: essays in honour of ole-Jorgen Skog*, ed. J. Elster, O. Gjelsvik, A. Hylland, and K. Moene, 11–28. Oslo: Unipub Forlag.
 45. Wilson, James Q., and Richard J. Herrnstein. 1985. *Crime and human nature*, 45. New York: Simon & Schuster.
 46. Baum, William M. 2005. *Understanding Behaviorism*. 2d ed. Blackwell.
 47. Ainslie, George. 2013. Grasping the impalpable: the role of endogenous reward in choices, including process addictions. *Inquiry* 56: 446–469.
 48. Ainslie, George. in press. *De gustibus disputare: Hyperbolic delay discounting integrates five approaches to choice*. *Journal of Economic Methodology*.
 49. Mitchell, Jason P., Jessica Schirmer, Daniel L. Ames, and Daniel T. Gilbert. 2011. Medial prefrontal cortex predicts intertemporal choice. *Journal of Cognitive Neuroscience* 23: 857–866.
 50. Peters, J., and C. Büchel. 2010. Episodic future thinking reduces reward delay discounting through an enhancement of prefrontal-mediotemporal interactions. *Neuron* 66(1): 138–148.
 51. Wittman, Mark, Kathryn L. Lovero, Scott D. Lane, and Martin P. Paulus. 2010. Now or later? Striatum and insula activation to immediate versus delayed rewards. *Journal of Neuroscience, Psychology and Economics* 1: 15–26.
 52. Gregorios-Pippas, Lucy, Philippe N. Tobler, and Wolfram Schultz. 2009. Short-term temporal discounting of reward value in human ventral striatum. *Journal of Neurophysiology* 101: 1507–1523.
 53. King, D.L., P.H. Delfabbro, and M.D. Griffiths. 2011. The role of structural characteristics in problematic video game play: an empirical study. *International Journal of Mental Health and Addiction* 9: 320–333.
 54. Ascher, M.S., and P. Levounis. 2015. *The behavioral addictions*. Washington, DC: American Psychiatric Publishing.
 55. Green, Leonard, Astrid Fry, and Joel Myerson. 1994. Discounting of delayed rewards: a life-span comparison. *Psychological Science* 5: 33–36.
 56. Chapman, G. 2002. Your money or your health: time preferences in trading money for health. *Medical Decision Making* 22: 410–416.
 57. Gollier, Christian, and Martin L. Weitzman. 2010. How should the distant future be discounted when discount rates are uncertain? *Economics Letters* 107(3): 350–353.
 58. Schouwenburg, H.C., and J.T. Groenewoud. 2001. Study motivation under social temptation: effects of trait procrastination. *Personality and Individual Differences* 30: 229–240.
 59. Monterosso, John, and George Ainslie. 1999. Beyond discounting: possible experimental models of impulse control. *Psychopharmacology* 146: 339–347.
 60. Glimcher, Paul William, Joseph Kable, and Kenway Louie. 2007. Neuroeconomic studies of impulsivity: now or just as soon as possible? *American Economic Review* 97: 142–147.

61. Moody, L., and W.K. Bickel. 2015. Symmetrical discounting of the future and the past in heavy smokers and alcohol drinkers. *Drug and Alcohol Dependence* 156: e157.
62. Jones, B., and H. Rachlin. 2009. Delay, probability, and social discounting in a public goods game. *Journal of the Experimental Analysis of Behavior* 91: 61–73.
63. Brown, R., and R.J. Herrnstein. 1975. *Psychology*, 146. Boston: Little Brown.
64. Becker, G., and K. Murphy. 1988. A theory of rational addiction. *Journal of Political Economy* 96: 675–700.
65. Ainslie, George. 2010. The core process in addictions and other impulses: hyperbolic discounting versus conditioning and cognitive framing. In *What is addiction?* ed. Don Ross, Howard Kincaid, David Spurrett, and Peter Collins, 211–245. Cambridge, Massachusetts: MIT.
66. Ryan, R.M., and E.L. Deci. 2000. Self-determination theory and the facilitation of intrinsic motivation, social development, and well-being. *American Psychologist* 55: 68–78.
67. Singh, S., R.L. Lewis, and A.G. Barto. 2009. Where do rewards come from. Proceedings of the Annual Conference of the Cognitive Science Society 2601–2606.
68. Mazur, James E. 1986. Choice between single and multiple delayed reinforcers. *Journal of the Experimental Analysis of Behavior* 46: 67–77.
69. Kahneman, Daniel. 2000b. Evaluation by moments: past and future. In *Choices, values, and frames*, ed. Daniel Kahneman and Amos Tversky, 693–708. Cambridge, UK: Cambridge University Press.