

Motivation must be momentary

George Ainslie

As it appeared in
Understanding Choice, Explaining Behaviour
Essays in Honour of Ole-Jørgen Skog

Edited by Jon Elster, Olav Gjelsvik,
Aanund Hylland and Karl O. Moene

Unipub forlag, Oslo Academic Press, 2006. Pp. 11-28.

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, and is thus not subject to copyright. The opinions expressed are not those of the Department of Veterans Affairs or of the US Government.

Utility theory has represented motivation as a consistent function of expected value, but it relies on three commonsense assumptions which, on closer examination, are probably false: 1. Utility is seen as somewhat inertial, so that a ranking of preferences made at one time would not change at another time in the absence of new information. 2. Utility is seen as a cognitive judgment, a matter of information processing, rather than as the biological factor that determines behaviour from moment to moment. 3. Utility is seen as originating in environmental events, interpreted but not created within the individual. These assumptions have made utility theorists aspire to more regularity in predicting choice than is probably possible. Ole-Jørgen Skog has been a pioneer in criticizing the rational choice assumptions of utility theory (e.g. Skog, 1999a). This chapter further explores the obstacles to predicting choice as a function of future incentives.

1. Hyperbolic discounting leads to volatile preferences and intertemporal conflict

The effect of timing on value is controversial. Utility-based theories generally take notice of the universal tendency to discount future prospects, but disagree about whether this is rational, and, if rational, whether it is also obligatory i.e., whether choice is constrained by this judgment. One school of thought has always held it to be irrational to devalue delayed events for any reason other than reduced probability (e.g. Pigou, 1920), but as increasingly efficient marketplaces over the centuries have insisted on devaluation simply for delay, that school is now little heard from. These marketplaces have determined that delayed events should be devalued with an exponential function—change by a constant percentage per unit of time—since all other

functions will make a person into a *money pump*: They will make her preferences inconsistent over time, so that competitors who devalue the future according to exponential curves will be able to buy from her when she devalues a good and sell it back to her when her estimation of it climbs (Ainslie, 1991). However, even among exponential functions shallow curves have a competitive advantage over steep ones in competitive markets, so that someone consistently willing to lend at a given rate will wind up impoverishing someone willing to borrow at that rate. The fact that history has not long ago crowned some society that arrived at zero discounting (corrected for uncertainty) reveals that a person's discount preferences are not a matter of dispassionate choice, but rather that low discount rates are in some way hard, so that achieving them is a feat. Something in nature makes us strongly present-oriented.

Marketplaces where goods with long-term value are traded are an artifact of human society. Of course, it may have been that nonhuman individuals that could tolerate more present deprivation or pain wound up surviving and reproducing more than others; but since a species' rewards are shaped by natural selection, optimal levels of this tolerance would have eventually become the most rewarding. Hoarding nuts became rewarding for squirrels and defending territory became rewarding for many species *in the short run*, so that, as instinct was shaped by the odds, immediate obedience to it stayed the most adaptive choice. Only we humans have outdistanced our evolution. That is, traits such as intelligence and imagination--which must have been shaped by competition with other hominids for the mundane resources that had always determined survival--have created an environment of competitions very different from the ones that shaped these traits. For the first time we create goods and bads with enduring value, but evolution has not had the time, or perhaps does not have the capacity, to change the present-weighted valuation process that kept our ancestors obeying their instincts.

A large body of parametric research with both human and nonhuman subjects has shown the discount function that describes our spontaneous choices among delayed prospects to be not exponential but hyperbolic (Ainslie, 2006; Green & Myerson, 2004; Kirby, 1997). This finding is now widely but not universally accepted, and some of its implications have been explored. This work is available elsewhere

(Ainslie, 2001, 2005), and will be only summarized here: Highly bowed curves give future contingencies radically different properties than exponential curves do. The most obvious difference is that curves from rewards of different amounts (as defined by their aptness to be chosen when available simultaneously) at different delays may cross, causing preference to shift as a function only of time and creating an incentive at one time to influence your own expected motivation at a later time. Hyperbolic valuations are disproportionately high when events are imminent, a shape that should make a hyperbolic discounter a *money pump* for any competitor that achieved exponential discounting. Our instinct is to seize the reward at hand, and resisting this instinct is hard. The conflict between present and future creates incentives for the self at one time to behave strategically toward selves at other times, so that both reported preferences and observed choices are apt to reflect the outcome of strategic thinking, not spontaneous preference.

Hyperbolic valuations are higher than exponential ones at long delays also, which gives a farsighted individual the opportunity to predict and forestall her own future urge, in the manner of Ulysses facing the Sirens. The conflict between the individual's own short and long range interests, probably never important for nonhumans outside of a psychology lab, creates in humans an intertemporal bargaining situation. The expectable growth of processes to deal with short-term urges (*impulses*) provides a bottom-up mechanism for the development of "higher mental functions," thus avoiding the theorist's need to postulate the top-down operation of such cognitive organs as the ego, the conscience, and the will. For instance, insofar as a person notices that her current choice in a situation is evidence of how she is apt to choose when similar situations recur, she would be expected to add the value of the better long range reward in these situations to its reward in this particular situation, thus creating a stake that arguably motivates willpower. It was Ole-Jørgen Skog who worked out the mathematical proof of the necessary summation property, which exponential curves do not have (1999b). This property has been found empirically in animals (Mazur, 1986), but may be more complex in humans, as we shall see.

2. Utility is more than information

In the conventional image of the future each sequence of choosable events is drawn on its own line extending from the present, and the events have heights and lengths representing their reward values and durations. Given a belief that alternative events can be represented this way at all—some authors and perhaps more nonspecialists believe that not all possibilities are commensurable (e.g. Schwartz, 1986)—motivation is seen as some function of these heights and lengths. Since the formalization of utility theory into rational choice theory (RCT) it has been popular to assume that individuals at a given choice point always choose the option that promises the greatest discounted product of height times length—when they are rational, if the theory is applied normatively, or at all times, if the theory is held to describe the fundamental basis of choice. The discovery that the basic discount function is hyperbolic undermined RCT's description of choice as inertial, but did not call into question the primacy of the graph of value times duration.

Such a graph has connotations of a psychometric analysis, the basis, perhaps, of the comprehensive choices imagined by the philosophy of mind, with “all things considered.” However, an important implication of hyperbolic discounting is that the self is a population of reward-seeking processes (Ainslie, 2001, pp. 39-47), so that utility is not primarily a pattern that some of these processes discern but the very factor that selects them to begin with. Intelligent individuals develop cognitive valuation skills, of course, and with such skills construct concepts of the future, but these skills do not thereby stand outside of the selection process, which operates strictly in the present moment. That is, although experimenters ask people to estimate future values, and people often engage in similar cognitive exercises in everyday life, future calculations will be meaningless unless converted to immediate motivational heft.

We humans have not lost our animal orientation toward maintaining present mood, so that expectations of the future will have motivational impact only insofar as they affect this mood. A strong current incentive can overpower the sum of all future incentives, as in the Biblical case of Esau, who became so hungry that he sold his birthright for a mess of pottage, Melville's Captain Ahab who made

sure he would not be obeyed if he ordered the amputation of his leg to stop, or, of course, Ulysses when he heard the Sirens. As a thought experiment the reader might imagine being offered a bet: A prize of ten times a year's income if you can hold your breath until you pass out (without tricks) as verified by EEG, but loss of half a year's income if you fail. Is it a bet you would take? Holding your breath that long will not hurt you, and your breathing will begin automatically once you pass out, but there may not exist enough incentive to make you confident of getting past the strong urge to breathe. (Saving someone's life? World peace? It will still be hard to be sure.) All plans, however great and extensive, have to pass through the narrow neck of present willingness.

The question for motivational science then becomes, how does the prospect of future goods and bads create the present experiences that select our courses of action? In particular, what makes the present willingness diverge from future evaluation? Even small differences in timing matter. For a big enough reward, you might be willing to let someone else strangle you to unconsciousness, even though you would not be sure of holding your breath that long. You might be more able to refrain from signaling the person to stop strangling you than to refrain from breathing, even though the difference in time it took to get a breath of air might be no longer than a second or two. A great factor in present willingness is clearly timing, as hyperbolic discounting predicts. But there is reason to believe that the hyperbolic function that has been observed with delays varying from seconds to decades is not one long, continuous curve, but has different mechanisms as the orders of magnitude shift.

The range of moods that determine present willingness is not nearly as wide as the range of wealth and poverty that can be rationally calculated. We have a tendency to imagine units of pleasure (utiles, say) as analogs of currency, so that, if my experience this minute is worth one utile, it might be worth half a million utiles for me to keep it up for a year. But hedonic tone cannot be observed straightforwardly in the way brightness or loudness can, and what sense we have of it can register only so many distinctions between the deepest unpleasure and the highest pleasure. No one has devised a practical measure for utiles; but the number of just-noticeable-differences in our most finely tuned sensory modalities is limited to

about 100, and must be less in the proximoreceptor modalities that are probably closer analogs of our pleasure-detectors, say for saltiness or the intensity of a smell. If asked whether we would rather have a current good mood for thirty days or thirty-one, have a hundred dollars or a hundred and five, even enjoy a million utiles or a million and five, we can give decisive, consistent answers, but clearly not by weighing the feel of one option against the feel of the other. We have learned many procedures for solving quantitation problems—the process that Piaget documented in its early phases (1937/1954)—and are particularly facile in pitting one set of numbers against another and determining which wins—but the resulting hedonic sensation has to do with the outcome of these tests rather than the raw impact of the quantities involved.

We sometimes laugh at our ability to peg current mood to distant predictions:

Fred: It says here that the universe will end in ten billion years.

Ned: That's awful, how can we live with that knowledge?

Fred: Well, ten billion years is a very long time.

Ned: Oh, ten billion? That's a relief. I thought you said ten million.

We set up tests with not only numbers, but heights, distances, durations—any modality that we can line up along a scale. In the case of money, Lea and Webley have recently attributed the difference between scalar value (dollar amount) and felt valuation to a “drug effect,” (as opposed to the “tool effect,” instrumental value—in press), but I have argued that this drug effect is only the most observable example of a general phenomenon. For instance, when McClure et.al. awarded subjects “immediate” Amazon.com book coupons they observed fMRI activity in brain reward centers that did not occur if the coupons were to be delayed by two or four weeks, even though the only good for which the coupons could be exchanged would have to be mailed to them (2005; see commentary by myself and Monterosso, 2005). People have been reported to form expectations of reward that are not the same as the experiences of reward on which these expectations are based. Kahneman has summarized evidence that subjects undergoing uncomfortable experiences use a combination of their greatest and their most recent reported levels when rating the experiences in retrospect (2000). (Note that I am treating aversive

experiences as reward-nonreward sequences, and thus as on the same dimension of value as reward—see Ainslie, 2001, pp54-61). The Kahneman findings suggest that even in a parametric experiment the expectation of reward may not be proportional to the experience on which it is based. The height and length of the plot of expected reward, if that is a meaningful representation, cannot be simply copied from the height and length of the reward(s) that the subjects themselves have reported in this situation, a finding that complicates the simple summation process reported by Mazur (1986, v.s.).

Another reflection of the departure of humans' observed choice from the elementary discounting process may be the great variability in their discounting parameter, k . People report consistent preferences among different amounts of money at delays of years or decades, and these preferences obey the same hyperbolic curves that describe students' preferences for fruit juice or pigeons' preferences for grain; but the k of preferences for money differs among individuals by factors of hundreds, while that among animals choosing food that will be delivered in seconds differs only by single digits (Monterosso & Ainslie, 1999). As long as the situation does not suggest the need for *money pump* precautions, the exercises that people have learned for quantitating remote events remarkably generate hyperbolic shapes whatever the range of times. I would expect people to report hypothetical preferences for “an imaginary good measured in X” hyperbolically as a function of delay, or even “an imaginary entity with a good feature measured in X and a bad feature measured in Y” to be rated hyperbolically as a function of Y. But in such exercises the representation of the future has clearly diverged from prediction of hard motivational force, the kind that determines whether we choose the “mess of pottage” or not—even though the form of both, where it can be discerned, is hyperbolic.

Remote valuations clearly provide material for use by the process that maintains present mood, but this process does not just pass these evaluations through. There is another step that may or may not preserve sheer hyperbolic discounting. The properties of this step are far from clear at this point, but they seem to depend on the nature of emotional reward—the kind of reward that least resembles the consumption of an external commodity.

Higher mental processes. Before discussing emotional reward, let me suggest how the “higher” mental processes or ego functions--the ones so often depicted as dispassionate cognitive faculties--may be conceived as natural growths from the behaviours by which we forage for reward in our immediate environments. As long as an individual responds to whatever cues come from moment to moment, without calculation, her choices will be shaped by whatever factors govern optimal timing of these responses. If she is walking and sees a ditch full of water across her path, she will begin her run to jump it at the moment when her experience has shown that it will be most efficient. She will learn not to run unnecessarily or start too late, and even if she only half wants to jump over the ditch she will learn not to half jump over it. If a cue predicts food she will start generating appetite at the moment that it will maximize the reward for eating; if the cue predicts pain she will learn when to best make her avoidance response, if possible, and when to narrow her openness to experience in order to endure it at the least cost, if not. Pain cues will also shape the emotions of fear or anger, which, I have argued, function in the same way as appetites (Ainslie, 2001, pp. 65-69). These again will come to be timed so as to best produce reward and reduce loss of reward.

Intelligence will permit prediction of such contingencies from greater and greater distances, potentially months and years in humans. It will also support the development of processes that are no longer specific to a particular appetite, but that search for appetite/satisfaction combinations—the functions that Piaget called “tertiary circular reactions (1937/1954),” and which are now more apt to be called metacognitions (Flavell, 1976). After that it will support the development of processes that examine these search processes and select among them. However, all of these processes still depend on moment-to-moment reward. The higher processes function like the writer or director of a play, who has not only to create an overall emotional design that will justify an evening at the theater, but also to translate this design into an adequate string of involving moments. The exposition that sets the action up, for instance, must be energized by enough ongoing by-play that it does not demand too much patience of the audience, lest the audience maintain its own hedonic tone by withdrawing its involvement and thinking about something else. Similarly, a person may design a day or a lifetime so as to fit her

ideals or other theories, but must make sure that there will be enough hedonic tone—cash flow, as it were—from moment to moment to keep her plan from derailing. Higher processes predict and manipulate more short-sighted ones, broker them, but are ultimately at their mercy if foresight errs. This is a *bottom-up* model, which does not assume that higher processes are inborn, although it does permit inborn preparedness to learn them.

Concrete images often help, however homespun. Here is one for the bottom-up development of ego functions: Think of a clearing burned away by a forest fire. The first opportunists to fill it are rapidly growing weeds. In competition for the limited area of sunlight the weeds are gradually overgrown by bushes, which grow more slowly but can support higher leaves. Gradually the bushes are passed by birches and the birches by still taller trees, oaks perhaps, until the competition occurs thirty meters or more above the ground. However, each tree in its infant stages must be able to survive the competition of the faster-growing weeds. This is the math of increasingly foresighted mental processes, except that the growth is not upward but earlier in time, and further from the reward that selects them rather than closer to the sun. Like most analogies it is an inadequate model. Trees that grow toward their “reward” can completely monopolize it, extinguishing growth below them. Mental processes that grow *away* from their rewards can only partially cull processes that are close to rewards, and this by art rather than strength, for they are progressively weaker as distance increases.

Of course mental processes adapt in other ways besides foresight—suitability for particular drives, learnability, interactive patterns with other processes—just as trees benefit from drought-resistance, cold-resistance, etc. besides height. But *ceteris paribus*, more foresighted processes that can function at a distance from their rewards will get a jump on those that cannot, avoiding them, forestalling them, or exploiting them as the distant incentives dictate. In the evolution of species the next step beyond height was mobility. Animals emerged to forage among plants, just as processes that are effective at getting reward can be exploited by rootless processes that search for that very property, the Piagetian tertiary circular reactions. Ultimately these search processes learn to behave strategically toward each other, analogously to the carnivores that feed on other animals—

a quaternary circular reaction, if you will, which was never postulated before hyperbolic discount curves described an incentive for such a process.

How does reward function in shaping higher processes? The simplest theory would be that these processes are selected entirely on the basis of the greatly attenuated hyperbolic discount curves from the distant rewards that they plan for; but this theory is almost certainly inadequate, because the curve from an event even weeks away would be unphysiologically low. The curve from decades to seconds cannot be continuous. Indeed, this consideration predicts the need for the augmented valuations of money, book coupons, etc., in the artificial choice situations that we have just discussed. But valuations at distances of years still take a hyperbolic form. How can we understand an augmentation process that seemingly pulls reward out of nowhere but remains functionally hyperbolic?

To discuss such a process we need to recognize a fundamental property of reward that tends to be obscured by the study of external incentives: that it is an intrapsychic process, and often occurs without external occasions.

3. Value is created within the individual

Foresight is not passive. An organism is not simply acted upon by the rewards that attach to the options in a given situation. Even a rat seems actively to construct possible action sequences and test them for rewardingness before choosing, an observable process that the early psychologists called vicarious trial and error (VTE—Tolman, 1939). People imagine scenarios that are undoubtedly more complex, and less restricted to the elements of the current environment. Furthermore, foresight is not emotionally neutral. The estimation of future utility is not a weightless function that simply reports the size and likelihood of rewards, but rather a process that is rewarding in its own right. It would be hard to know whether a rat gets reward from VTE that adds anything to the remembered contingent rewards in the given situation, but for a person foresight is imagination, more or less constrained by the demands of a task. When we imagine scenarios we generate the emotions that they occasion. When we imagine food we are readily lured into appetite, a mistake that leads to pangs of hunger if we are

hungry and food is not available. We often conjure lust and rage in situations where they cannot be “gratified,” for these appetites (or emotions) are gratifying in their own right and the mainstays of the various fiction-producing industries. We pay even to experience fear and grief (in horror films and tear-jerkers), a phenomenon that I have cited to argue that even “negative” emotions must have a rewarding component (Ainslie, 2001, pp. 173-186). To fit the needs of foresight these imaginings must be disciplined by what psychotherapists call reality testing, but they still create temptations to distort prediction unless the cost of bad prediction will be felt imminently.

The conventional view of self-generated emotion is that it is a trick—we imagine a stimulus for the emotion and experience a conditioned response—and that it is thus limited by the stimulus pattern that we imagine. Furthermore, to generate a strong emotion we would need a strong conditioned stimulus, that is, we would have to actually expect the innately programmed, “unconditioned” stimulus. I have elsewhere criticized the conditioning theory of response choice (Ainslie, 2001, pp. 19-22), and will comment here only that if emotional responses are conditioned, then someone like an actor who learns to command them at will, should experience increasing difficulty with repetition because they extinguish, rather than increase ease as she succeeds in summoning them. In contrast to this view I have argued that emotions and (other) appetites are both reward-dependent and rewarding in their own right—that their compelling quality, which often but not always evades deliberate control, comes from the immediacy of this reward, and the negativity of some of them comes from obligatory nonreward that follows the reward cyclically. The experience of desirable emotions is limited not by their dependence on releasing stimuli, but by their tendency to habituate and produce diminishing returns when not paced by occasions of limited availability and predictability.

To illustrate this distinction in more concrete terms: The conventional economic model of choice is one of obtaining scarce commodities to repair environmental deprivations. (Naturally enough. The builder with only a hammer sees only nails.) The infant learns to find the breast when hungry and to cling to the mother when cold or scared; the toddler learns other needs and satisfactions according to the same model, and so on to the consuming adult. Value is thought

of as arising from innate drive-reducers or appearances associated with them. However, young children of all people are driven least. Most of the time they are not hungry or in pain or scared, and they maintain their mood perfectly well, merely noticing minor surprises in their world. When there are no surprises they languish a bit, and ultimately their intelligence suffers, but it takes only someone rolling a ball to them or the discovery of a step to learn to jump off to provide occasions for what seems empathically as cheer. They need only physical comfort and what we might call *texture* in their surroundings—patterns that can serve as occasions for emotion in play. Sometimes children are driven from this Eden by grown-ups' expectations regarding school, but if they are sent to a permissive school the Eden falls away anyway, not from the harsh demands of reality but from the inadequacy of self-generated play. Children get too good at anticipating their own fantasies, and need greater complexity with unpredictable elements. They learn increasingly evocative and durable fantasies that incorporate these elements, so that the level of competition for their attention rises.

This competition increasingly involves future prospects, but only as these prospects serve as occasions for current emotion. To illustrate how prospects do this: near prospects support feelings of anticipation, but more distant scenarios can be used to pace current feelings in much the same way that someone else's story can—"vicarious experience" (see Ainslie, 2001, pp. 179-186) Even a general appreciation that you have the necessary elements for a good future invites an emotion, which could be called a sense of wealth. Wealth is not a passive state; it entails the risk of losing it and an incentive to defend it if necessary. Loss demands that you avoid appetites for which the object is no longer available--fantasies that used to be "real"--a painful task apparently eased by generating the emotions of grief and/or anger. Wealth can be defined as all circumstances that are at risk for loss, and should thus include the prospect of personal relationships. Mourning is the process of learning not to generate appetites for the lost objects, a process which, if successful, ends this incentive for grief and anger.

Any value that the past can have must likewise be realized in current emotion. People differ widely about what this value is. Some regard the past as refuse, without current value beyond the constraints

or opportunities it may have left. Others regard their histories as their greatest wealth. Neither can be called wrong. The use of the past to occasion present feeling is an optional skill, although its components have been little explored. When “they can’t take that away from me,” what is it that they can’t take? The mere occurrence of pleasure may leave nothing usable. The extreme case is the crack cocaine high, which must be repeated indefinitely as though previous ones had never happened. Other entertainments may leave a sense of satisfaction or accomplishment without increasing your actual wealth; people sometimes avoid breaking up a completed jigsaw puzzle in an attempt to prolong this sense. There seems to be a continuum of pleasures from those that leave no ongoing feeling to those that have great momentum.

More distant events, at least toward the latter end of the continuum, can be useful as memories. But what is it about memories that makes them more useful than mere fictions? This question applies to vicarious experiences, too, and, I have argued, to future prospects. By the time a child has learned the main properties of occasions for emotion, two factors determine which emotional reward processes (or fantasies, or imaginings) win the competition to be entertained: a. the uniqueness of the occasions for the emotion and b. the goodness of the gamble that these occasions will occur, that is, how close it lies to an optimal point on a continuum between “a sure thing” and hopelessness. Together these two factors are what constitute texture.

- a. The commonest test for uniqueness is factuality, the qualities that support belief and thus distinguish this particular imagining from make-believe. The requirements of factuality for occasioning emotion are broader than the requirements for scientific or other instrumental purposes, which are shaped by success or failure in the outside world. However, instrumental usefulness is a good criterion for the broader kind of factuality, ironically, since occasioning emotion is a fundamentally non-instrumental use for the same imagining. Other criteria are consensual tradition among a uniquely defined population (your country, your company, your family...), your own past belief for a large fraction of your life, a rare coincidence, and doubtless others. The value of memories for

occasioning feeling is apt not to be related to the instrumental value of what you learned, but rather to their uniqueness, that they stand out from imagination in general.

- b. The goodness of a gamble as a generator of occasions depends on the combination of wins, which consume appetite, and losses, which refresh appetite. This is the factor that makes human relationships valuable beyond the instrumentality with which conventional utility theories explain it. The gamble itself must be adequately unique—not, for instance, a game of solitaire in which you change the rules to accommodate near-wins, or a personal relationship that you can wholly dominate.

I have described these hypotheses in more detail elsewhere (Ainslie, 2001, pp. 166-180 and in press). I use them here to illustrate how a person who has free access to emotional reward might come to depend on external occasions to generate it effectively. Goods that serve as occasions will appear in many respects to be the traditional stuff of commerce with which economics is familiar, but in other respects they will seem alien. They may be valued in proportion to their scarcity up to a point, for instance, but beyond that be devalued entirely. Or a gambler may arrange to lose while trying her best to win. In many ways the value of an outcome for an activity may diverge from the very value that is the ostensible rationale for the activity.

Appetite defeats prediction. The hypothesis that emotions and other appetites are self-generated—that they are reward-seeking behaviours with rewarding properties of their own—permits at least a sketchy explanation for why choice is not more dominated by the sheer proximity of rewards than it is. A large number of rewards can be had immediately if we so choose. Pleasant daydreams, rest from difficult tasks, entertainment, even food and alcohol can usually be had with little or no delay. The factor that keeps them from being overwhelmingly intrusive is probably appetite. Most rewards are not very effective cold, without preparation. Rewards do not reach their full effectiveness the moment you think of them, but only after an appetite for them develops. Conversely, appetites can be thought of as seeking rewards. An appetite based on a hunger, for food, say, needs the prospect of an object in order to compete for attention, and will arise insofar as it may succeed in motivating eating. In the population

of processes that comprise the self, these appetites have the same incentives as house pets that are sometimes fed under some circumstances and never under others—Begging is relatively cheap, but will not be worth it if they have no chance at all. An appetite based partly on emotion, say for sex, may be intrinsically rewarding enough to compete for attention without the prospect of an object, but will still reward only slightly until it has been entertained for a while, as will appetites without corresponding objects of consumption, *viz.* emotions. By limiting early reward, appetites provide a buffer against snap decisions, perhaps the only one available in the organisms where they first evolved. When combined with foresight, appetites make the choice process really complex.

The self-generated nature of reward defeats the neat utilitarian plots of value against time. A plot that is true given a current set of emotions and hungers may change radically if the person changes any element in the set. Reward and thus behaviour often depends on self-generated phenomena even for occasions whose rewarding effect is pegged rather strictly by physiology—food and recreational drugs, for instance, where false belief *i.e.*, placebo effect, has only limited capacity to occasion reward or prevent pangs/withdrawal symptoms. There may even be a recursive self-predictive process that generates new sets of appetites from second to second. I have argued that this happens in willpower, for instance, where a person's motive to succeed depends on her current estimate of her probability of success. It may happen in a simpler situation where a person's hunger may vary with her expectation that this hunger will motivate her to stop an ongoing activity and assuage it, or even where an ostensibly involuntary behaviour like vomiting depends on her belief about whether she will vomit (Russell, 1978, pp. 27-28). The urge to urinate is an appetite that clearly depends on the prospect of doing so. (This is not discussed much, but Elster concurs—1999, p. 227, note 2), and generating it in the absence of opportunity readily turns painful. Sudden urges conventionally ascribed to conditioned stimuli despite an absence of new information may be explainable as a sudden calling into doubt of self-expectations, when an appetite detects a circumstance that may give it more chance against a dominant intention.

Urination is a good paradigm of a drive-limited modality: The drive is the filling of the bladder, the appetite becomes increasingly likely as it fills, but until the level of drive is extreme the appetite is occasioned by opportunity. The conventional, conditioning explanation of this urge requires that thousands of opportunities become associated with urination, unless we say that “opportunity” itself, that is, the complex operant cue, is the conditioned stimulus—not very parsimonious. A paradigm for the larger category of appetite that is not drive-controlled is anger, which is intrinsically rewarding (Lerner, et.al., in press) but which can compete with other processes mainly when other sources of reward are poor. The appetite facilitates further appetite up to a point, is supported by good texture i.e., strong but not overwhelming opposition, and does not turn painful as hunger does if not somehow consummated. I submit that the combination of these types of appetite, when analyzed as behaviours rather than conditioned responses can describe the full range of motivated behaviour. However, this description will be far more complex than the smooth curves of conventional utility theory, even when these are bent into hyperbolae.

The complexities of a bottom-up theory will necessarily multiply until they can match the complexity of their subject. I imagine that I have made some wrong turns in building on the foundation of hyperbolic discounting, but I am sure that this foundation itself is solid. I have described these possibilities in gratitude to Ole-Jørgen, whose own work and comments on mine have repeatedly been useful in developing them.

References

- Ainslie, George (1991) Derivation of "rational" economic behavior from hyperbolic discount curves. *American Economic Review* 81, 334-340.
- Ainslie, George (1992) *Picoeconomics: The Strategic Interaction of Successive Motivational States within the Person*. Cambridge: Cambridge U.
- Ainslie, George (2001) *Breakdown of Will*. New York, Cambridge U.

- Ainslie, George (2005) Précis of *Breakdown of Will*. *Behavioral and Brain Sciences* 28(5), 635-673.
- Ainslie, George (2006) A selectionist model of the ego: Implications for self-control. In Natalie Sebanz and Wolfgang Prinz, Eds., *Disorders of Volition*. MIT.
- Ainslie, George (in press) What good are facts? the “drug” value of money as an exemplar of all non-instrumental value. *Behavioral and Brain Sciences*
- Ainslie, George and Monterosso, John (2004) A marketplace in the brain? *Science* 306, 421-423.
- Elster, Jon (1999b) *Strong Feelings: Emotion, Addiction, and Human Behavior*. Cambridge, MA, MIT.
- Flavell, J. (1976) Metacognitive aspects of problem solving. In B. Resnick (ed.), *The Nature of Intelligence*. Erlbaum.
- Green, Leonard and Myerson, Joel (2004) A discounting framework for choice with delayed and probabilistic rewards. *Psychological Bulletin* 130, 769-792.
- Kahneman, Daniel (2000) Evaluation by moments: Past and future In Kahneman, Daniel, and Tversky, Amos (eds) *Choices, values, and frames*. Cambridge University Press.
- Kirby, Kris N. (1997) Bidding on the future: Evidence against normative discounting of delayed rewards. *Journal of Experimental Psychology: General* 126, 54-70.
- Lea, Stephen E.G. and Webley, Paul (in press) Money as tool, money as drug: The biological psychology of a strong incentive. *Behavioral and Brain Sciences*.
- Lerner, J. S., Tiedens, L. Z. and Gonzalez, R. M. (in press [slated for 2006, issue #2]) Portrait of the angry decision maker: How appraisal tendencies shape anger’s influence on cognition. *Journal of Behavioral Decision Making*.
- McClure, Samuel M., Laibson, David I., Loewenstein, George, and Cohen, Jonathan D. (2004) The grasshopper and the ant:

- Separate neural systems value immediate and delayed monetary rewards. *Science* 306, 503-507.
- Mazur, J.E. (1986) Choice between single and multiple delayed reinforcers. *Journal of the Experimental Analysis of Behavior* 46, 67-77.
- Monterosso, John and Ainslie, George (1999) Beyond Discounting: Possible experimental models of impulse control. *Psychopharmacology* 146, 339-347.
- Ouellette, Judith A. and Wood, Wendy (1998) Habit and intention in everyday life: The multiple processes by which past behavior predicts future behavior. *Psychological Bulletin* 124, 54-74.
- Piaget, J. (1937/1954) *Construction of Reality in the Child*. M. Cook, Trans. New York: Basic.
- Pigou, A. C. (1920) *The Economics of Welfare*. London: Macmillan.
- Russell, J. Michael (1978) Saying, feeling, and self-deception. *Behaviorism* 6, 27-43.
- Skog, Ole-Jorgen (1999a) Rationality, irrationality, and addiction—Notes on Becker and Murphy's theory of addiction. In Jon Elster and Ole-Jorgen Skog (eds), *Getting Hooked: Rationality and Addiction* Cambridge University Press, pp. 173-207.
- Skog, Ole-Jorgen (1999b) Addiction, choice, and self-control In Jon Elster (ed), *Addiction: Entries and Exits* Russell Sage, pp. 151-168
- Tolman, E. C. (1939) Prediction of vicarious trial and error by means of the schematic sowbug. *Psychological Review* 46, 318-336.

