# Self-control, discounting and reward: Why picoeconomics is economics

**Don Ross**
University of Cape Town

**George Ainslie**
Veterans Affairs Medical Center, Coatesville PA

**Andre Hofmeyr**
University of Cape Town

## *1 Introduction*

Many economists are skeptical, or at least uneasy, about the agenda of behavioural economics (BE) because it appears to aim at the transformation of the microeconomics of human agency into a branch of psychology. Gul and Pesendorfer (2008), for example, have produced a widely-circulated anti-behavioural polemic that is sometimes thought to go too far, but to nevertheless articulate appropriate concern for defense of the distinctive way in which economists conceive and model choice (Schotter 2008, Harrison 2008, Harrison & Ross forthcoming). The crux of the issue is as follows. Economists' core phenomena of study are responses to changes in incentives. Such responses, whether effected by individual agents or by aggregates of them, are all conceived as *choices*. This language might suggest commitment to an 'inboard' process of psychological (so, ultimately, neuropsychological) computation of selection of options from sets of alternatives. However, as various authors including Clark (1997), Ross (2005) and Smith (2008) have emphasized, this suggestion is misleading; it is not of direct importance to economists whether the influence of incentives on choices is channeled mainly through such internal processes or mainly through factors supporting what Smith calls 'ecological rationality', that is, structures in interactive environments that sculpt behavioural adaptations so as to track implications of changes in incentives even when they are not explicitly represented. (To cite in example developed by Ross, pp. 262-264, starfish make 'rational choices' in this sense even though they have no brains in which alternatives *could* be represented.) By contrast, psychologists are interested in modeling, and ultimately in integrating, *all* factors that enter into the production of behaviour, including purely causal influences that are independent of incentives, such as reflexes, or 'automatic' or conditioned behaviors. Economists who are skeptical or uneasy about BE worry that the so-called irrational motivations emphasized in that literature simply represent the displacement of the economist's explanatory project by the psychologist's. It might be thought that such displacement should better be thought of as *widening* of the economist's focus, because incentives are a subset of the set of psychological factors. The very point of the 'ecological rationality' critique, however, is to challenge this interpretation, and urge attention to incentives that are channeled mainly through the social environment rather than through psychological mechanisms.

We sympathize with economists' reasons as canvassed above for holding BE at arm's length. We also agree, however, with Schotter and Harrison that Gul and Pesendorfer's blanket rejection of the relevance of any psychological dynamics to economic modeling and testing is theoretical dogmatism. In this paper we aim to persuade economists that modeling strategies should be pursued by which a particular strand of the psychology of choice, Ainslie's *picoeconomic* account of self-control (Ainslie 1992, 2001), should be incorporated into the microeconomist's toolkit. (This does not imply that it should be taken out for use on every occasion.) Many economists are likely to think that this has already been done in a small but steadily growing tradition of literature often thought to have begin with Laibson (1994, 1997). A second aim of this paper is to show why no existing economic model adequately captures the choice dynamics that picoeconomics addresses. At the same time, we are concerned, as our third aim, to convince the reader that these are indeed *choice* dynamics, in the economist's sense. If they were not, we would agree that economists could and should regard them as somebody else's – psychologists' – business. We will then seek to support our claims empirically, by giving an account of a recent choice experiment we performed with a group of addicts. The final section summarizes and urges implications for future research.

## *2. Impulsive choice and alternative discount functions*

One of the phenomena that has preoccupied the BE literature is intertemporal preference reversal associated with choices thought to be impulsive due to impatience and related 'hot preferences' for 'viscerally' attractive rewards (Loewenstein 1996; Loewenstein & O'Donoghue 2004) that operate on agents with disproportionate strength at short ranges. Such preference reversal is often regarded as 'irrational', in which case 'rationality' is correspondingly associated with successful exercises of self-control. Such exercises are in turn associated with personal and social policy options for managing the visceral attractors, especially targets of potential addiction.

Early neoclassical economists such as Fisher considered the discounting of utility as a mere function of temporal delay to be irrational. In the mid-twentieth century, however, such discounting became a fundamental element of welfare economics based on cost-benefit analysis. The case for the rationality of temporal discounting crucially rests on the idea that risk increases with delays to consumption and with the length of intervals between investments and returns. Risk factors arise with respect to two general issues: accuracy in estimating magnitudes and signs of future returns, and questions about which agents will still be around to enjoy benefits. Assuming an intertemporally consistent utility function and rational expectations, an agent will discount delay to returns and consumption as a linear function of the passage of time. In the simple linear case, the discount formula is given by

$$v_i = A_i e^{-kD_i} \qquad (1)$$

where $v_i$, $A_i$, and $D_i$ represent, respectively, the present value of a delayed reward, the amount of a delayed reward, and the delay of the reward; $e$ is the base of the natural logarithms; and the parameter $0 > k > 1$ is a constant that represents the influence of uncertainty and the agent's idiosyncratic attitude to risk.

The choice-theoretic model of addiction from which all currently leading ones among economists are conceptually descended is due to Becker and Murphy (1988). This is usually referred to as the 'rational addiction' model. The basis of the label is that according to the model, although the addict's repeated choices cause her welfare to steadily decline *over* time, *at* any given moment of choice she is best off choosing to consume her drug. Thus the addict can never be said on this model to make an irrational choice. This has struck many people as too intuitively bizarre to take seriously. However, the judgment that a scientific model or theory confounds everyday intuitions is not a substantive objection to it, since there is no reason to believe that everyday intuitions track truth. The *sound* reason for the eclipse of the rational addiction model is that it makes a general false prediction about addicts' attitudes to relapse following periods of withdrawal. According to the model, *if* an addict successfully overcomes withdrawal *and* intended at the time of entering withdrawal to quit for good, then either she should never relapse or she should never permanently quit. But the overwhelming majority of addicts repeatedly undergo the rigors of withdrawal, planning never to relapse, and then return to addictive consumption – before eventually quitting successfully. Thus the rational addiction model is empirically refuted by a basic aspect of the standard addictive pattern. (For the detailed logic underlying this argument, along with relevant empirical evidence on the typical addictive life-cycle, see Ross *et al* 2008, Chapter 3.)

The roots of this misprediction by the rational addiction model lie in the fact that it has no device for allowing either an individual's preferences or her objects of preference to change over time. Consequently, the most common general strategy for economically modeling inconsistent people has been to drop the assumption that a person implements one economic agent over her whole biography. Instead, the person is broken up into interacting sub-agents on synchronic and / or diachronic dimensions. An analogy can be drawn here between people and countries (Ross 2005). When countries behave irrationally – erecting self-harming barriers to imports, for example, or alternately running down national savings and enforcing unpopular belt-tightening – this is straightforwardly explicable by reference to the interactions of rational citizens acting according to their divergent their parochial preferences. Note that this has not deterred economists from often modeling countries as agents. This has regularly been adapted as a model for the dissolution of puzzles about preference reversals in individual people. If the person is divided into agents synchronically, she becomes a *community* of agents. If the person is divided diachronically, she becomes a *sequence* of agents. In either case, we can model the whole person using game theory: her so-called 'molar' behavior is treated as a dynamic equilibrium of bargaining games amongst synchronic or diachronic agents. The analogy between countries and people breaks down in at least one crucial aspect, however. Unless one implausibly models national governments as standing outside of the bargaining games among citizens and institutions, in the case of a country there is no over-arching agent to which to ascribe the national utility function; modeling a country as an agent is purely an idealization. By contrast, many commentators have suggested that people act strategically to manipulate the sub-interests into which they are partially but incompletely decomposed. Economist Robert Strotz was the first behavioural scientist to suggest that people need in principle to deal strategically with their own predictable changes of preference (1956). He pointed out that a person cannot rely on her relative valuation of two future rewards staying the same as the rewards get closer – that a curve depicting the current value of a perspective reward as a function of its delay is apt to stray from the exponential function that describes consistent preference. He did not specify a particular alternative function, but noted that any function that was not exponential would lead to inconsistent preferences over time. Previous writers had observed that people sometimes pay for committing devices to forestall their own impulses. Marshall, for instance, noticed that the poor preferred to buy small amounts of coal or alcohol at retail rather than buying a large stock at a lower unit price to keep at home (1921, p. 814). However, Strotz was the first theorist to describe the need for such commitment as dynamic inconsistency, a product of the nonexponential discounting of delayed rewards. Later, Schelling (1978, 1980, 1984) wrote persuasively about the subtle self-management of motivations and opportunities in the implementation of such personal policies as initiating and maintaining abstinence from smoking.

In a seemingly unrelated line of experiments, psychologist Richard Herrnstein reported that subjects tend to sample two concurrent streams of reward in proportion to the mean rates, amounts, and immediacies of those rewards (the "matching law;" 1961, 1997). Building on both ideas, one of us (GA) pointed out that application of the matching law to single (discrete, i.e. not streamed) choices between smaller sooner (SS) and larger later (LL) rewards would imply four specific predictions (Ainslie, 1975, 1992):

1. The decline in rewarding effect with delay is described better by a value function that is inversely proportional to delay (hyperbolic discount curve) than by a function that declines by a constant proportion of remaining value for each unit of delay (exponential discount curve). That is, data on the evident value of a single prospective reward at varying delays will be described better by a hyperbolic than by an exponential function of those delays. The most commonly used form of the hyperbolic curve has been is given not by the exponential function (1) but by a hyperbolic function as described by Mazur's formula (2):

$$v_i = \frac{A_i}{1 + kD_i} \tag{2}$$

   A second exponent parameter on the right-hand denominator is sometimes adopted to produce closer fits to empirical discounting data in people (Myerson & Green 1997).
2. Preferences between some pairs of an SS reward and an LL alternative at varying delays but with a constant lag between SS and LL rewards will favour LL rewards when both are distant, but shift to SS alternatives when they become closer that is, will imply dynamic inconsistency of the sort noted by Strotz.
3. During the period when an LL reward is preferred, subjects may have motivation to choose options whose only function is to prevent their own subsequent choice of the SS alternative that is, they have incentives to seek to precommit their future behaviour.

4. Subjects choosing between whole series of SS/LL pairs at once will have a greater tendency to choose the LL rewards than will subjects choosing between the same pairs in discrete succession. A subsidiary prediction is that people who perceive current choices between SS and LL rewards as test cases that is, as examples that predict their own future preferences between similar pairs of rewards – will be more likely to prefer LL rewards than when they frame a pair of alternatives as an isolated choice. That is, people may create *bundles* of interdependent expectations by predicting their future choices *recursively* on the basis of each current choice.

Predictions (3) and (4) define subject matter for the sub-field of *picoeconomics*, the study of the strategic interaction of successive motivational states within the person.

Subsequent work by both economists and psychologists has often confirmed the first three predictions, both in human subjects whose choices are spontaneous (uninstructed) and in nonhuman animals.[1] Highlights of this experimental literature may be summarized in terms of three main phenomena:

- *Curve fitting.* Plots of indifference points between pairs of SS and LL rewards as delay is varied often show less least-squares deviation from hyperbolic than from exponential curves (reviewed in Green & Myerson 2004; Kirby 1997). (A limitation of this literature is that attitudes to risk and convexity of the utility function are not controlled for; see Andersen *et al* 2008. Such controls reduce the proportion of choices in aggregate data that are best described by hyperbolic functions, but do not eliminate them.)The steepness of discounting has been found to vary greatly, both among individuals and among tasks within individuals (Frederick *et al* 2003), often in ways that reflect differences in psychological health –steeper in addicts than in non-addicts, for instance (Bickel et.al., 1998; Bickel & Marsch, 2001). However, the hyperbolic shape has been observed in all time ranges, even among outcomes that will not be realized for decades. People expressing preferences for public programs that would pay off at delays up to 100 years produced a similar curve (Cropper *et al* 1992), as did an analysis of discount rates implied by actual long term public works investments (Harvey 1994).
- *Preference reversal.* When subjects are given choices between a given pair of SS and LL rewards with a fixed lag between them but a varying delay until the SS reward would be available, their preferences often change from LL to SS as this delay gets shorter (reviewed in Green & Myerson, 2004). This temporary preference for SS rewards is found in both people (Kirby & Herrnstein, 1995) and pigeons (Ainslie & Herrnstein, 1981; Green *et al* 1981), and can occur even when there is a delay before the earliest SS reward is available (Ainslie & Haendel 1983; Green *et al* 2005). That is, a subject who preferred $50 in one year to $100 in three years often did not prefer $50 in four years to $100 in six years, choosing the SS reward even when it would be delayed by one year but not four years.

[1] The nonhuman experiments demonstrate that the findings cannot simply be attributed to experimenter suggestion, cultural norms, or executive functions unique to humans.

Such findings argue against the proposal, to be discussed presently, that SS rewards are temporarily preferred only because immediacy creates an arousal effect.

- *Precommitment*. The simplest evidence of making a choice purely to avoid a later choice comes from experiments in which nonhuman animals will push a key to avoid a later choice of immediate SS versus delayed LL food (Ainslie 1974), or press a lever that commits them to a later SS electric shock to avoid having to choose at that later time between SS and LL shocks (when the SS would be immediate; Deluty *et al* 1983). It is hard to strain humans' self-control ability with respect to important impulses in the laboratory, but surveys of purchasing behaviour in combination with subject self-reports have found many examples of apparent deliberate precommitment of the kind pointed out by Alfred Marshall, such as buying tempting consumables in small packages despite strong price incentives to buy in bulk (Wertenbroch 1998), and voluntarily arranging deadlines to avoid procrastinating unpleasant medical tests (Trope & Fishbach, 2000) or term papers (Ariely & Wertenbroch 2002). Likewise in much larger choices people are said to make illiquid investments that would be otherwise inefficient, in order to commit their future spending (Laibson, 1997).

For reasons that we will discuss shortly, economic research and theory have tended to veer away from the fourth prediction. Although it has the most significance for practical self-control, it has been studied the least.

Economists have been hesitant to adopt true hyperbolic functions in models of intertemporal choice. A key reason for this is that utility functions based on Mazur's formula are not guaranteed to converge. Concern about this is not merely a matter of methodological dogma. In practical terms, it blocks use of standard econometric tests of specific models against data, which rely on asymptotic properties of functional forms: as we collect more preference data, fit to an estimated function should get no worse. It should not be intuitively surprising that this fails for hyperbolic discounting, since the slope of a function estimated from two indifference points will be sensitive to their relative positions in time and to the temporal reference point from which they are evaluated. This practical problem soon found a popular solution among economists who require workable functional forms for model testing but aim to capture intertemporal preference inconsistency that appear to be real phenomena. Laibson (1994, 1997) observed that predictions 2 and 3, and much existing evidence about human impulsiveness, were consistent with a function that was much closer to formula 1. This was a step function, borrowed from an existing model of the motives in intergenerational wealth transfers (Phelps & Pollack, 1968), and usually referred to as 'quasi-hyperbolic' or 'β-δ'. This class of functions is expressed by

$$v_i = A_i\beta\delta^D \qquad (3)$$

where $v_i$ represents the present value of a delayed reward, $A_i$ the amount of that reward, β a constant discount factor for all delayed rewards, δ a per-period exponential discount factor, and $D$ the delay of the reward. Where $\beta = 1$ the equation reduces to standard exponential discounting. Where $\beta < 1$ discounting is initially steeper up to some inflection point, then flattens. β-δ discounting predicts that value drops precipitously from no delay to a one-period delay, but then declines more gradually (and exponentially) over all periods thereafter.

Accordingly, a *hyperboloid* (or *quasi-hyperbolic*) discounter chooses consistently until the SS reward is imminent, when its value suddenly spikes upwards by $(1 - \beta)/\beta$, potentially reversing an existing preference for an LL alternative.

Laibson's original purpose in suggesting β-δ discounting was to create an approximation to 'true' hyperbolic discounting, adopted to achieve convergence of functions that sum over flows of values (personal communication). The hyperbolic function was taken to give the correct molar-scale account of discounting, with β-δ discounting then treated simply as an approximation for the sake of making closed-form representation and conventional econometric testing analysis tractable. Impulsiveness can readily be modeled as such a problem for a person decomposed diachronically into sub-agents. The present agent, like the present generation, discounts the utility of future agents (generations) relative to its own at the β rate, while setting policy for later agents (generations) who discount one another's utility at the δ rate. Models of this kind, applied to the more general class of myopic-choice phenomena, can be found in a number of influential papers (Laibson 1998; Laibson, Repetto, & Tobacman 1998, O'Donoghue & Rabin 1999a; O'Donoghue & Rabin 2001).

The two discrete components of motivation in the β-δ model fit seductively into the long tradition of two-faculty models that date back to Plato's chariot of the soul, pulled by the well-behaved horse of reason and the unruly horse of passion (*Phaedrus*, 253e). When Laibson proposed his hyperboloid formula Thaler and Shefrin (1981) had already described their two agent model, in which a farsighted (or slowly discounting) *planner* sets rational policies but must rely on a shortsighted (or rapidly discounting) *doer* to execute them. The planner was the familiar agent from conventional utility theory, but the doer did not have a clear basis – and nor did Laibson's β factor.

Loewenstein (1996; Lowenstein & O'Donoghue, 2004) duly proposed a rationale for the doer-- that it could arise from the "viscerality" of some motives (1996; Lowenstein & O'Donoghue, 2004).[2] The viscerality hypothesis came to be widely accepted in the BE literature, and was soon married to the βspike of imminent reward in formula 3 (McClure *et al* 2004), producing a model in which a person discounts the future consistently until

---

[2] The defining property of viscerality has been vague, having come to include such appetites as curiosity and the prospect of money. Viscerality has been attributed to any motive associated with appetite or emotion, indeed any motive that generates temporary preference for an SS reward, a usage that begs the question of how this preference reversal occurs.

an appetite or emotion is aroused, and then becomes markedly less patient. In an additional refinement the all-or-none spike has since been softened into a second, steep exponential discount curve from the moment of expected reward, which sums with the shallower standard discount curve (McClure *et al* 2007):

$$v_i = v_0\{(\omega\beta^D) + [(1-\omega)\delta^D]\} \qquad (4)$$

where $v_0$ = value if reward is immediate and ω is a weighting factor. We will refer to upward departure from the exponential curve in either (3) or (4) as the *beta spike.*

The β-δ function has become the standard discounting model in the BE literature. Benabou and Tirole (2004) re-founded much of the recursive self-prediction model as it appeared in Ainslie (1992) on the assumption that the basic discount curve is β-δ. Some empirical evidence has been consistent with it. Loewenstein and various collaborators have built on the work of psychologist Walter Mischel (e.g. Mischel & Moore 1980) to show that decisions made in a state of emotional arousal have a greater tendency to favour SS alternatives. Recently Laibson and colleagues have exploited β-δ discounting to interpret neuroimaging data in support of a molecular-scale account of intertemporal preference reversal. McClure *et al* (including Laibson) obtained fMRI evidence they interpret as suggesting that evolutionarily older orbito-frontal brain areas discount more steeply than later-evolving prefrontal areas. They then propose that hyperbolic discounting at the molar scale be understood as an aggregation of the tug of war between neurally localized β-discounting orbito-frontal and δ-discounting frontal sub-agents.

Thus the idea that viscerality is the source of the βspike that explains impulsiveness has roots in psychological research as well as common experience. However, there are problems both with viscerality as an ultimate cause of impulsiveness and with the β-δ model itself. The dual brain center interpretation of the McClure *et al* findings has recently encountered empirical difficulties. Glimcher *et al* (2007) report fMRI data they interpret as indicating that neurons in both midbrain and prefrontal areas in fact implement similar discount functions as the molar subject. They do find, as have others (Hariri *et al* 2006), variability in activation levels in striatum between subjects, which correlates with variability in molar steepness of discounting. However, they find no areas in which activation levels are correlated with steeper or shallower discount functions than those inferred from molar behavior. Glimcher *et al* interpret these findings as directly challenging the McClure *et al* hypothesis, and indeed generalize this doubt: not only does it undermine the idea of distinct δ and β discounting areas in the brain, but also, they say, the very existence of "separable neural agents that could account for multiple selves that are used to explain hyperbolic-like discounting behavior" (p. 143). Similarly, Delaney *et al* (2008) report neuroimaging evidence for activation of separate neural areas for discounting in different behavioral domains (health maintenance and financial planning), but no evidence for separable processing across time horizons. Glimcher *et al* (2007) also find no evidence of brain valuation sites that do not respond to delayed rewards. All of the sites reported by London *et al* 2006) likewise respond to to both immediate and delayed rewards, a finding that is not surprising in light of the fact that animals lacking a prefrontal cortex have long been known to evaluate SS/LL reward tradeoffs (Ainslie &

Herrnstein 1981) and sometimes to commit themselves against future impulses (Ainslie 1974; Deluty *et al* 1983).

However the neuropsychological picture ultimately settles, the hypothesis that arousal of appetite or emotion even when reward is not imminently available creates a β spike complicates the interpretation of formulas 3 or 4, since we can no longer assume that spikes occur just at near-zero delays. The pre-reward spike of Laibson's hyperboloid discounting becomes an upward departure that can occur any time that appetite is stimulated (McClure *et al* 2004) –or, presumably, when appetite arises without a stimulus.[3] The point from which delay should be computed is then ambiguous. The key factor in valuation for the imminent future becomes the arousal of emotion, which has as much to do with the person's choices, for instance to entertain "hot" or "cool" thoughts (Metcalfe & Mischel, 1999), as it does with the properties of the reward itself. Abstract rewards such as money, Amazon coupons, and the satisfaction of curiosity have been claimed to behave as visceral rewards (Loewenstein, 1999; McClure et.al., 2004), as Ainslie and Monterosso (2004) pointed out; while tempting foods can be experienced as cool options (Metcalfe & Mischel, 1999). There is also the problem that some impulsive choices involve mundane activities that are not emotionally arousing – procrastination is a prominent example (Ainslie, 2010c; O'Donoghue & Rabin, 1999a) – while others may take longer than arousal can literally be sustained. Examples might include buying a house against one's better judgment, taking an unskilled job instead of going to college, or not contributing to a retirement plan.

Recently, behavioural economists have begun to view the putative line between visceral and ordinary rewards as problematic (Fudenberg 2006, pp. 699-700). Fudenberg suggests that a partial solution may be that token rewards such as money act as conditioned cues that predict later rewards. However, conditioning is just contingency learning, so *all* predictive cues for *all* rewards can equally be said to be conditioned (Rescorla, 1988). Many authors are now agnostic as to the origins of the βspike, but they still tacitly or explicitly assume the operation of β-δ discounting in models of self-control. In those models, a problem with β-δ functions themselves becomes apparent: although β-δ curves simulate a pure hyperbola's especially high discount rates when a reward is imminent, they do not preserve its especially low discount rates at long delays, and thus do not support prediction 4. We will argue that the lack of this property has hobbled the many attempts that have been made in the mainstream economics and BE literature over the past decade to model internal self-control, that is, willpower.

*3. Willpower with pure hyperbolae*

---

[3] Arousal in response to unpredictive stimuli, or no stimuli, is also problematic for pure hyperbolic functions. Ainslie (2010a) has suggested a positive feedback process based on recursive self-prediction, that is converse to the mechanism for self-control to be described presently.

Consider a β-δ function $F$ with reference point $R$ and interval $I$ between $R$ and the inflection point where δ discounting begins. Suppose that $k = x$ for the function that describes discounting over $I$ (i.e. $\beta = (v_i = A_i e^{-xD_i})$). Finally suppose that $F$ converges asymptotically to $v_i = L$. Note that $I$ is a period of time. Now consider a pure hyperbolic discount function $F^*$, conforming to formula (2), defined over the same temporal interval as $F$, with reference point $R$ and $k = x$ over $I$. If $F^*$ also converges to $v_i = L$, then it must do so more slowly than $F$. Note also that although there are hyperbolic functions where $k = x$ in which the value of $k$ over the first sub-interval of $I$ is less than the value of $k$ over the corresponding sub-interval of $I$ for $F$, there are no hyperbolic functions $F^*$ where $k = x$ and the value of $k'$ in any discount function over the first sub-interval of $I$ consistent with $F^*$ is *greater* than the value of $k''$ over the corresponding sub-interval of $I$ for any exponential discount function consistent with $F$. This precisely expresses the correspondence relationship underlying the intuitive idea that pure hyperbolic discount curves are not only steeper than relevantly corresponding β-δ functions when delays are short, but are also flatter than β-δ curves when delays are long. Still more intuitively, the tails of hyperbolic curves are higher than the tails of β-δ curves 'all else being equal'.

This property predicts people's disproportionately high valuation of long range investments, as described above for prediction #1, and a more robust tendency to make long range commitments against impulses than β-δ functions predict (prediction #3). However, the most important consequence of these high tails is that for a given long-range temporal horizon, the discounted value of a *sequence* of rewards adds up to a greater total for hyperbolic than for the corresponding β-δ function. β-δ functions share with hyperbolic ones the property that the aggregate discounted value of a series of LL rewards will be greater than that of a single LL reward, and for some empirically relevant temporal interval could be greater than the value of βspikes from SS alternatives. But this aggregation effect attenuates faster for a β-δ function than for its corresponding hyperbolic function.

A constructed numerical example illustrates the point. If four LL rewards of non-discounted present value 10 units of reward are available 10 time intervals apart, and a delay of 10 time intervals is beyond the range of the β component of formulas 3 and 4, then by formulas 1, 3, or 4 a slope that would yield a discounted present value of 1.0 units at a reference point 10 intervals before the first reward (δ = 0.794) would yield 1.0 + .1 + .01 + .001 units, or just 1.111 units from that same reference point for the summed prospect of all four rewards together. For the corresponding hyperbolic function (formula 2) a slope that gave a value of 1.0 units at the reference point 10 time intervals before the first LL reward (k = 0.9) would yield 1.0 + .526 + .357 + .270 units, or 2.15 units for the discounted sum of the values of all four ((10/(1+(.9 x 10))+10/(1+(.9 x 20))…). For a series of ten such rewards the β-δ sum is 1.111111111 units. The present value never reaches 1.2 units, whereas the hyperbolic function gives 3.08 units for ten rewards and keeps climbing, albeit more slowly, for longer series. If we accept the empirical evidence that hyperbolic discount curves from multiple delayed rewards are additive (Kirby 2006; Mazur 1997), then a person's perception of her choices as being between whole series of

LL versus SS rewards could significantly increase the attractiveness of the LL rewards, without calling upon any supplementary source of motivation.

In consequence, hyperbolic functions over combinations of delayed rewards are more likely than corresponding β-δ functions to yield more aggregate value than their combined SS alternatives, for each of which the value is taken on the margin as it becomes imminent. In effect, a planner and a doer could exist simultaneously with regard to the same set of alternatives, depending on how the person predicts contingent reward, rather than having to take turns at sole control as in standard planner / doer models descended from Thaler and Shefrin (1981). Even a single LL reward in the previous example would have a value of .001 units of reward at delay = 40 if discounted in β-δ or exponential fashion, but .27 units if discounted according to the corresponding hyperbolic function. The hyperbolic function thus predicts not only a greater aggregate motive but a stronger motive for pre-commitment. Therefore, as we will argue in some detail, the power of summed hyperbolically discounted LL rewards suggests a mechanism *based directly on comparative incentives* for the otherwise vaguely intuited willpower, without resort to any specific hyptheses about brain resources.

The higher tails of hyperbolic functions relative to corresponding β-δ functions have been generally ignored in the economic literature, and misunderstood in the rare article dealing with the summation of multiple rewards. Musau (2009) uses exponential, quasi-hyperbolic, and hyperbolic functions to calculate the incentive for cooperation in an iterated prisoner's dilemma with an unknown endpoint, which is the sum of the discounted values of each outcome. He assumes that the hyperbolic discount curve becomes increasingly steep relative to an exponential curve as delays increase, leading to more stringent restrictions on conditions under whichcooperation is sustainable. However, although hyperbolic curves are steeper than corresponding exponential ones at short delays, they are *decreasingly* steep with longer delays and eventually become *less* steep than their exponential counterparts, leading to increased incentive to cooperate as more delayed rewards are taken into account. The importance of this property in repeated prisoner's dilemmas specifically becomes obvious as we consider how people might actually choose reward in whole categories.

When it comes to implementation, there is a hitch to the strategy of choosing options in whole categories instead of discretely. Under the pressure of temptation, a doer could simply propose to take the imminent SS reward and the series of LL rewards in subsequent choices. This is indeed the behaviour that O'Donoghue and Rabin (1999a, 1999b, 2000, 2001) expect, and which limits the choices of procrastinators in their model. The intuitive solution is not obscure. Mere awareness of the relevance of the question, "If not now, when?" may suffice for an agent that isn't strictly myopic – and was in fact proposed by Ainslie (1975) many years ago (1975). But without the strong additive effect based on the comparatively high tails of hyperbolic curves, this solution is hard to model. With a strong additive property a hyperbolic function can stake a whole category of choices on a current choice, but both parts of prediction #4 are necessary:

4a. Subjects will choose LL rewards more when choosing between series of LL versus SS rewards than when choosing singly. This has been shown for both college students and for non-human animals. Students chose LL rewards (pizza in one experiment, money in another) over SS alternatives more often when they expected to get five weekly deliveries of the same kind than when they expected to choose week by week (Kirby & Guastello 2001). And rats chose LL amounts of sugar water more often when making choices for the next three trials at once than when choosing on a trial-by-trial basis.

4b. Self-aware subjects, presumably only humans, who frame their current choices as predictive of a category of similar choices in the future may thereby tie together their expectations of these choices into just such a series (Ainslie, 1975, 1992, 2001). The hypothesis is that a combination of imperfect self-prediction and a tendency to temporarily prefer SS rewards sets up a limited warfare relationship among successive motivational states (or 'selves', if taken to mean only 'oneself in different motivational states'), which can be resolved by discerning, or defining, a variant of repeated prisoner's dilemma among these selves: defection in the present case makes defection in future cases more likely, not from a motive of retaliation but by making cooperation seem likely to be wasted.[4] The aggregate LL reward may be from the evident consequences of each choice, as in binges with hangovers, or may be a continuous state such as good health, adequate savings, etc.. The necessary element is that the person's expectation of getting the aggregate LL reward is put at stake when each opportunity for a definable set of SS rewards occurs.

### *4. Hyperbolic discounting and economic modeling*

Hyperbolic discounting has featured much less prominently in economic models than β-δ discounting, for a variety of reasons. As noted previously, hyperbolic functions are less tractable than exponential ones because the recursive self-prediction that they imply means that competitions between SS and LL rewards cannot be solved by standard equilibrium analysis.. The bargaining among sub-personal interests motivated by picoeconomics invites Nash equilibrium analysis and solution, but this seems to rely on the challenging idea that a person incorporates simultaneously active sub-agents while continuing to exercise economic agency – i.e., respond appropriately to incentives – herself. As we have pointed out, dividing agents into sub-agents has a long history in Western philosophy, from which economists have frequently borrowed. However, when the agent is strictly *decomposed* into these sub-agents, and then especially if the sub-agents are identified with neuropsychological and other biological causal forces, the outcome is abandonment of the development of a general account of incentive-response, that is, a collapse of economics into psychology and neuroscience. Economists wisely resist such disinvestment in a research enterprise that has been highly productive. We aim to show, however, that picoeconomics, unlike much of BE, does not motivate such

[4] Thus Musau's (2009) article on aggregate incentives to cooperate in iterated prisoner's dilemmas with unknown endpoints is highly relevant.

disinvestment, even though it raises modeling challenges that economists have not yet met.

It is characteristic of most economists' approach that they treat the *objects* of choice – e.g., drug consumption or alternative activities such as working or enjoying family time – as fixed through all stages of a given run of analysis. For reasons that have partly to do with mathematical and statistical tractability, but are also related to the demands of policy domains in which we cannot, for practical or ethical reasons, discriminate amongst individuals, the economist's modeling strategy is most useful for generalizing the costs and benefits of forms of impulsiveness, such as substance abuse, across *populations* of people. The approach collapses all of the specific ways in which different people idiosyncratically experience dilemmas of choice into a variable parameter that represents a cost of resisting temptation to choose SS rewards over LL rewards. This allows us to then represent, for purposes of empirical measurement in a given range of cases, the comparative effects of both contextual factors in the environment and factors 'intrinsic' to the tempting rewards, on the cost variable. In the case of abused and addictive substances, examples of contextual factors include monetary prices of drugs relative to budgets, expectations of social disapproval of drug consumption and intoxication, risks and magnitudes of legal sanctions, and cultural expectations that normalize, to varying degrees in different environments, the social role of the substance abuser or addict. Examples of intrinsic factors include expected health consequences, expected intensity and magnitude of highs, and anticipations of hangover and withdrawal.

The economic literature features a range of models in this family. Since they are descended from predictions 2 and 3 above, one might be tempted to call these models picoeconomic. However, they have not made use of what we argue is the crucial factor for understanding the dynamics of impulse control, prediction 4. We will reserve the term 'picoeconomic' for models that incorporate recursive self-prediction.

In depicting how someone can control a fundamental tendency to temporarily prefer imminent SS rewards, the theorist employing $\beta$-$\delta$ discount functions has only two kinds of model at her disposal: a person can behave strategically toward a future self upon which her present self will be unable to exert continuing influence, or she can bring to bear an additional source of motivation, inside her biological self but outside the her agency as a responder to incentives. Authors have pursued both possibilities, which could be labeled *prior strategy* models and *sovereign will* models, respectively (Ross 2010). Leading instances of prior strategy models are Akerlof (1991), Laibson (1997, 1998), O'Donoghue and Rabin (1999b, 2001), and Brocas and Carrillo (2008). Willpower is introduced by Gul and Pesendorfer (2001), Bénabou and Tirole (2004), Benhabib and Bisin (2004), Bernheim and Rangel (2004), Fudenberg and Levine (2006), and Heidhues and Koszegi (2009). Beginning with O'Donoghue and Rabin these authors have suggested strategies that can work without recourse to external commitment devices. The first class of models tend to envision competing choices made by different sub-parts of people's brains, while those in the second group often treat temptations as costs that are imposed on a unitary chooser by forces from the external environment – which, crucially, can include parts of a person's own brain over which she lacks voluntary control. In other words, models of the first type treat different parts of the brain as distinct choosers, while the second type treat parts of the brain as external to the agent who makes choices.

In general, these differences do not reflect divergent ideas about how the brain works; they rather represent alternative abstract *descriptions* of the relationships between relevant neural mechanisms, whatever exactly these turn out to be, and incentivizing pressures and opportunities. Economics is a kind of engineering, and like other engineers economists choose descriptions that are best fit to practical policy alternatives. A property of all of the models cited above is that they admit of *equilibrium solutions*. That is, they allow one to look for simultaneous quantitative values of costs and benefits such that, to the extent that the model specification is approximately accurate, proportions of choices of addictive behaviour are predicted to remain *stable*. In principle this allows a policy maker to try to influence levels and prevalence rates of impulsive choice through manipulations of variables that are under her control, such as taxes on alcohol and cigarettes, penalties for possession of illegal drugs, convenience of access to vendors of fatty foods, and so on. Of course, such control is always imperfect; but economists can factor such imperfection into their models through use of probabilistic response functions.

Having noted the central role of equilibrium in models of self-control, we can compare the picoeconomic model with those that have recently become prominent in economic modeling.. The most basic model of internal self-control involves no control at all, only navigation. A relatively powerless but farsighted agent is what Strotz proposed (1956); this dates back to Ulysses and the Sirens. Beginning in 1999 O'Donoghue and Rabin published a series of game-theoretic analyses of impulsiveness (1999a, 1999b, 2000, 2001). They pointed out that people who are not aware of their tendencies to form temporary preferences for SSs ("naifs") would always fail to complete long term plans. People who are fully aware of these tendencies ("sophisticates") may take evasive action, but the authors proposed no way in which people could increase their motivation in opposition to βspikes. Sophisticates have to take the dominance of β spikes at certain choice points as given, and find prior actions that prevent themselves from arriving at those choice points, just as Strotz had originally concluded. In Thaler's terms the doer of one moment is the planner of future moments, with the functions distinguished only by the βspike that influences the doer but not the planner.

Fudenberg and Levine model the general family of impulse control problems using what they call a 'dual self' approach. In this model, long-run selves choose actions that influence the utility functions of short-run selves. Short-run selves then choose behaviors for the organism. Short-run selves are completely myopic, so are (in the first place) meliorators rather than expected-utility maximizers. Costs to long-run selves of manipulating short-run selves are functions of the differences between what non-manipulated and manipulated short-run selves would do, and thus are endogenous. When decision problems concern only future options, short-run selves are indifferent, so long-run selves can determine their actions at minimal cost. Some manipulations of current short-run selves constrain options available to future short-run selves; these manipulations are sometimes efficient with respect to expected utility optimization. The solution concept imposed on the games is what the authors call 'SR-perfection', which requires that short-run selves maximize in every subgame, and long-run selves anticipate that they do. This is equivalent to what is elsewhere called 'Markov Perfect Nash Equilibrium', and is sufficient for unique equilibria in the games. Every SR-perfect NE profile is provably equivalent to a reduced-form optimization problem. Equilibria exist in which self-control costs are too high to be rationally paid; thus the model predicts manifestations of procrastination and molar-scale preference reversal.

There are clear differences between the dual self model and picoeconomics. Picoeconomics supposes that long-range and short-range interests have different, often conflicting, preferences. By contrast, Fudenberg and Levine's selves differ only in the extent of their information about the future. But the crucial difference is that Fudenberg and Levine's short-run selves are essentially forces outside the agent that behave in such a way as to impose management costs on her, much as in sovereign will models.

Once one has driven a wedge between the 'fully-fledged' agent, the planner, and internal, manipulable but ultimately exogenous forces, doers, then degrees of freedom in modeling open up – often, perhaps, too widely for empirical testing to readily distinguish between them. Brocas and Carrillo (2008) make a running with the extreme assumption that intellectualized planner functions ("principals") can only get information about the value of alternatives by calling upon ("delegating to") short-sighted but sensitive doer functions ("agents"), at a risk of being seduced by the options that tempt the doer functions. A person's balance between impulsiveness and control then depends on the extent to which the planners invite this sampling, gathering information but sometimes finding out thereby that they face a dominant SS reward.[5] This fable-like division of faculties, effectively into one that can only see and one that can only feel, contradicts evidence that brain reward centers all see into the future. There seem to be no reward-evaluating parts of the human brain that are insensitive to long range consequences (Kable & Glimcher, 2007), and even pigeons can learn a committing device to forestall impulsive choices (Ainslie, 1974).

The more common solution to the self-control problem in β-δ models is to restore a sovereign will. Gul and Pesendorfer criticize solutions that rely on prior commitment (2001, 2004) by pointing out that self-control can often be exercised in the presence of tempting choices. They also criticize intertemporal bargaining models for not generating unique equilibria, and thus disturbing established economic problem-solving procedures as much as pure hyperbolic curves do. They propose that self-control has a measurable cost, defined as the difference between how much a person values achieving a given long

---

[5] This model does recognize some little-appreciated phenomena, such as the role of vicarious trial and error (VTE) in most or all choice-making (see Tolman,1939; Johnson & Redish, 2007) and the long-appreciated danger that VTE will arouse appetite. (Thus the Catholic Church makes the contemplation of a mortal sin a venial sin in its own right.)

range goal when she can make a binding external commitment to it, versus when she has to exert self-control. The person is not formally divided into planner and doer in their model, but implicitly there is an agent who values the future rationally, competes with another agent who is prone to impulsiveness, and chooses among three options: to forestall the impulsive agent in advance by external precommitment; to give in to the temptation; or to oppose a currently dominant impulse with a special budget of motivation called willpower. The first and second options reprise modeling ideas developed previously by O'Donoghue and Rabin (1999a, 1999b 2000, 2001). The third option lets the rational agent weigh the cost of both impulse and self-control at each choice point, and sometimes wrestle down the impulse while it is current, as people often say they do. In Gul and Pesendorfer's model, then, choice is distinct from the action or failure of willpower, which furthermore has nothing directly to do with discounting. Of all existing accounts, this is the only model of intertemporal vacillation that requires no revisions to standard consumer theory: the agent discounts exponentially, and may manifest preference reversal and/or failure of self-control even with full information.

The problem with sovereign will models is that they do not model an ultimate choice mechanism, but only move it back into a self within the self. The planner is a homunculus that gets its authority from its power to summon a special and rather mysterious force, willpower. In itself, this is not an objection that should trouble economists, if the details of the processes hidden in the box do not make a difference to reduced-form descriptions of agent-level responses to incentives. In any case, the picoeconomic account is also highly abstract so far as its account of mechanisms is concerned. The problem is that sovereign will models put the force that threatens willpower *outside* the box. They thereby remove from the choice dynamics factors that, according to the picoeconomist, are properly endogenous to it. They are consequently likely to mis-specify data and deliver unreliable predictions about choices – just what the economist aims to specify correctly and predict successfully, at least at aggregate levels.

The sovereign will theorists' conception built around an unstructured black box is a convenient folder in which to file some observed psychological phenomena. Experiments in which subjects' willingness to avoid SS rewards fatigues with massed practice, but increases with repeated spaced practices, have led to the suggestion that will resembles a muscle (Baumeister *et al* 1994, 2006). However, an analogy is not yet a model. If the will is a discrete faculty at the disposal of the planner, how can it be strong against one kind of target, e.g. wasting money, at the same time as it is weak against another, e.g. smoking? And if the will has its own kind of motivation, dimensioned as "strength," how does this motivation interact with the kind of motivation that presumably governs the planner's decisions? The existence of separate motivational centers for long and short term rewards would not answer these questions as long as there is a single marketplace in the brain; and in any case the separation of long and short term motivational centers is in doubt.

Benhabib and Bisin (2005) produce a model that implies but does not specify a β-δ value function ("Agents trade off 'excessive' and 'impulsive' immediate consumption with a consumption-saving rule requiring the exercise of self-control" [p. 462]). Here the usual rational agent or "control process" competes with an "automatic process," as described in psychology to underlie "classical conditioning and Pavlovian responses" (pp. 464-464). They mean their model not to depend on the "visceral / rational dichotomy per se" (p. 464n), but since a Pavlovian response is unmotivated by definition (Rescorla 1988) and their automatic processes compete on the basis of motivation, the best interpretation of their model is that they do mean automatic processes to be viscerally rewarded. These two kinds of "processing pathways" compete strictly on the basis of prospective reward: "An *executive function,* or *supervisory attention system*, modulates the activation levels of the different processing pathways, based on the learned representation of expected future rewards" (p. 464). The authors propose that self-control requires "active maintenance of a [long term] goal-like representation" whose cost, unlike that of Gul's and Pesendorfer's willpower, "is independent of the temptation" (p. 470). The executive function calls on one or the other (necessarily before the automatic pathway is active) depending on (1) what reward the automatic pathway is expected to bring (when?), and (2) the cost of inhibiting the automatic pathway, which is determined not by (1) – as in Gul & Pesendorfer – but by "the costs of maintaining a representation [of the long term goal] in active memory" (p. 466). Temptations estimated to be not worth this effort are given in to. This model does not seem to depict true willpower, in which, as William James said, the temptation is "held steadily in view" (1890); rather it uses the direction of attention as a commitment device. The sustainability of this tactic is necessarily limited, as in the Stroop-like examples the authors cite from experimental psychology (pp. 464-465).

Benhabib and Bisin's model relies on no explicit appeal to discounting at all. Once again, as in the general family of models we have been considering, the temptation of impulsive consumption (including impulsive consumption of leisure) is made exogenous to the agent. They assume that the agent can pay a cost to inhibit responses to temptations. Whether payment of this cost is justified in a given case is a function of the exponentially discounted relative values of the consumption plans with and without purchase of self-control measures. Conceptually, there is no gain here over the Gul and Pesendorfer model. However, the greater generality of Benhabib and Bisin's formulation allows them to reduce to variations in a single parameter *b* the difference between their model and, respectively, (i) the model of a perfectly consistent optimizer of permanent income, with or without smoothing over the life cycle, and (ii) the model of a $\beta$-$\delta$ discounter restricted to Markov Perfect Nash Equilibria.This allows them to generate and test relative empirical predictions of aggregate saving behavior for each of the three models. Of course the agent who must pay a cost to avoid succumbing to temptation saves less than the agent who is not tempted in the first place. The more interesting result is that the agent who can buy self-control saves more than the agent whose successive selves play Markov games, even in circumstances where the latter is never actually impulsive. The reason is that the former agent, but not the latter, expects to be least likely to be impulsive when it matters most; therefore the marginal expected value of present savings are higher for her.

The model of Benabou & Tirole (2004) captures many of the features of picoeconomics, including self-signaling,[6] self-enforcing personal rules, the magnified danger of lapses, and compulsive overcontrol. However, since they base their model on β-δ value functions they are, like other authors, led to invoke "imperfect recall" to account for the role of self-signaling. "Whenever people look back on their past actions to infer what they are likely to do in the future, it must be that the motives that led to these choices at the time are no longer accessible with complete accuracy or reliability" (p. 850). The reason that β-δ modelers consistently appeal to faulty recall is that they otherwise cannot explain why the prospect of visceral reward is not fully valued in advance.

Self-signaling is motivated to the extent that an agent routinely underestimates the effect that past and future visceral states will have on her future behavior, and thus neglects to incorporate their influence in making present choices. For example, an abstinent substance abuser may reinitiate consumption in part because he fails to remember the effect of past craving on drug use, and thus underestimates how much future craving will make it difficult to quit. This psychological hypothesis can be modeled either in a unified or multiple agency setting. In the multiple (diachronic) agent context, present agents with some degree of special concern for their descendent agents choose prudently because prudent choice now *predicts* prudent choice later, and the *present* agent gets lightly discounted *present* utility from this reassurance. As in picoeconomics, this explains a personal rule to frame rewards as series as a presently valuable asset, and self-signaling depends neither on hyperbolic discounting nor on dividing the person into intertemporally multiple agents. Bénabou and Tirole demonstrate rigorously that self-signaling can be justified by any imperfection in an agent's knowledge about her own capacity to cope with temptations.

However, when self-signaling models preserve full agent rationality in not requiring intertemporally inconsistent preferences, they must invoke compromised inferential rationality. Suppose we grant Loewenstein's (1996) hypothesis that people are poor at remembering and imagining the qualitative intensity of states resulting from visceral consumption. It would nevertheless be puzzling if the procrastinator or the addict could not *infer* from her past failures at maximization that the visceral temptations must have been quite powerful. Indeed, some such inference seems crucial to motivate the anxiety that is alleviated by self-signaling. But in that case one wonders why for most or many people the tendency to fall into impulsiveness is chronic and recurring. By contrast, the self-signaling involved in recursive self-prediction is not motivated by any lack of recall. Uncertainty is intrinsic to this intertemporal bargaining because the self-prediction is recursive, so that predictions not only can change because of factors not yet apparent, they are themselves agents of this change. What is uncertain is not the agent's episodic memory, but her likely interpretation of experiences clearly perceived and recalled.

[6] This is Prelec and Bodner's (2003) term for a person's interpretation of a current choice as a test case for similar future choices.

The proposals that have descended from Laibson's β-δ formula have not fully modeled self-control, not only with respect to its psychological dynamics, *but also with respect to its influence on choice*. They have not specified the nature and source of the motivation that can sometimes dominate current temptations. We have suggested here that a return to the pure hyperbolic discount function as originally proposed (formula 2; Ainslie 1975, 1992) can effectively integrate the various models that have used β-δ value functions by repairing the deficiencies of each. Like prior strategy models, the picoeconomic account predicts sub-personal bargaining. Like sovereign will models, picoeconomics addresses the recruitment of additional motivation that can overcome what would otherwise be dominant temptations, but does so without appeal to a special and mysterious source of psychic energy. The secret of this integration is to return to the neglected prediction #4, which does not follow from functions of the β-δ sort.

## *5. Evidence for the reward bundling effect*

Recursive self-prediction is not intuitively difficult. The factor that keeps a dieter from eating a small piece of chocolate now is obviously not the tiny weight gain it would entail but the risk to the credibility of her diet. The hypothesis that willpower depends on framing current choices as test cases predicts many familiar phenomena that do not follow from the β-δ models just examined. This hypothesis would seem to be necessary to explain some of them, such as rationalization, heterogeneous discounting, chronic failures of will, compulsiveness, and the experience of free will, since β-δ models do not suggest a need for them. All of these phenomena except the last should be of direct import to economists because they as recurrent influences on choice. And the nature of the experience of free will may explain why people are such competent economic agents while being such poor natural economists.

*Rationalization.* The perception that a much larger reward, or bundle of rewards, depends on a current choice puts pressure on a person to either forego a tempting satisfaction or to find some reason that the current choice is not a valid test case for the larger prospect. What is commonly called rationalization is the activity of distinguishing the case at hand from a larger category – e.g. "New Year's Eve doesn't count," "I'll never have a chance to eat that dish again," "I'm on vacation," and many other forms of "just this once." Without recursive self-prediction, this internal legalism seems pointless. But are we to deny that people regularly engage in it, and that it influences behavioural choice?

*Rational planning.* Within limits a person should be able to adopt various norms for discounting under varying circumstances (Ainslie 1991), including exponential discounting and no discounting at all (as earlier economists and philosophers recommended). These limits would be determined by the aggregate discounted value of the expected fruits of such a norm, relative to the discounted value of individual temptations as they drew near. If a person defines paying more than the prime interest rate as a defection from a prudent rule she has adopted, she will be more likely to lapse than if she defines paying greater than her current credit card rate as such a defection.

*Chronic failures of will and mental accounts.* Serious lapses are likely to result in the person's discrimination of the relevant circumstances from the larger bundle that is the basis of willpower, leading to what are in effect involuntary mental accounts--circumscribed areas of dyscontrol in which the person gives up on willpower ("I *have to* smoke after a meal," "I *can't* get myself to speak in public," etc.). People's experience of what would seem to be choices as involuntary symptoms has otherwise puzzled theorists. In fact, we need not deny that they are based on choice – but the choice is of a policy rather than of a specific action, and is not typically proximate in time to the action.

*Compulsive character.* A person's uncertainty about how she will interpret a current choice when looking back from the moment of a future choice will create an incentive to give temptations a wider berth than sheer calculation would require. If she is especially afraid of impulses and becomes aware of the rationale for this extra caution, she may come to perceive a failure to give such a wide berth as itself something of a lapse (Ainslie, 1992, p. 188). Bodner and Prelec (1997, 2001) showed that the additional diagnostic utility of such self-signalling would motivate drift into poles of scrupulous self-control or irresistible impulsiveness, a process that has been observed in a two-person analog of the intertemporal prisoner's dilemma (Monterosso *et al* 2002). If the direction of drift is away from scrupulous self-control, the result may be a chronic circumscribed failure of will, as above.

*"Free will."* Small changes of perceived symbolism – and hence category membership – may radically change the implications of a self-signal, so that a person cannot be absolutely sure of what her motivation and hence her choice will be even in the near future. This will create both the introspective opacity and sense of participation in one's own decisions that have been held to be key to the near-universal subjective experience of free will (Ainslie 2010b).

*Great variability in revealed patience.* The degree of patience (represented by the *k* parameter in Formula 2) estimated both from self-reports and actual choices varies widely among people, and among kinds of reward for a given person (Frederick *et al*, 2002). This heterogeneity – for systematic measurement of which we require the expertise of economists – contrasts with the single-digit consistency observed in nonhuman animals. The contrast makes sense if we understand the human findings not to represent spontaneous preferences, but preferences based to a greater or lesser extent on mental accounting practices, which can be exercised with more or less skill.

*6. Experimental evidence for reward bundling as a factor in choice*

Controlled experiments on intrapsychic bargaining processes are difficult, because they entail manipulating subjects' perceptions of what their current choices imply for their own future choices. However, the only published example of such manipulation reported promising results**.** Psychologists Kris Kirby and Barbarose Guastello studied choices between SS and LL amounts of money in undergraduate volunteers (2001). They used an auction procedure to find SS rewards that subjects slightly preferred to LL alternatives, then offered series of choices between these SS and LL rewards in one of three patterns:

(1) five separate choices at fixed intervals of several days; (2) delivery of the same five SS or LL rewards at the same intervals, but chosen all at once at the same moment as the first choice in (1); or (3) the same free choices as in (1), but with the additional instruction that "the choice you make now is the best indication of how you will choose every time" (p. 159). They found that a much larger proportion of subjects chose the LL rewards when choosing all at once (2) than when they expected separate choices (1), and an intermediate number choosing LL when it was suggested to be a test case (3). The authors repeated the procedure using slices of pizza (to introduce the celebrated viscerality effect from BE), and found the same pattern. They interpreted the results as showing that the students were lesspatient toward the first choice in a series than they were toward the series as a whole, and that this difference was reduced by the suggestion that the first choice set a precedent. Unfortunately, the subsequent choices that the subjects were told to expect in conditions (1) and (3) were not actually elicited.

Following Kirby and Guastello, we report the results of a study to investigate South African undergraduates' dispositions to frame rewards as bundles in the course of temporally sequenced monetary choices. We compared three conditions in separate groups of subjects: a 'free' condition, a 'suggested' condition and a 'forced' condition. Besides wishing to replicate and improve upon Kirby and Guastello's protocol in a moderately different culture, we wondered whether subjects in the free and suggested conditions would change their preferences with experience if they expected trials to continue into the future. Unlike Kirby and Guastello, our procedure involved no deception or merely hypothetical choice. In addition, we were interested in whether subjects who were more addiction-prone would respond differently to manipulations of the salience of current choices as predictors of future choices. In this normal sample we used smoking as an indicator of addiction-proneness.

*Methods*

Subjects were recruited through sign-up sheets circulated in undergraduate economics classes at the University of Cape Town (UCT). Each potential subject indicated whether he or she was a 'regular smoker', and was screened using the World Health Organisation's (WHO) Alcohol, Smoking and Substance Involvement Screening Test (ASSIST) (WHO ASSIST Working Group 2002). This allowed us to check the reliability of subjects' self-identifications of smoking status. In addition, subjects that met the criteria for alcohol or substance abuse were excluded to eliminate potential comorbid confounds. The Problem Gambling Severity Index (PGSI), the scored module of the Canadian Problem Gambling Index (CPGI) (Ferris & Wynne 2001), was also administered to all potential subjects, and those with gambling problems were excluded.

60 students – 30 smokers and 30 non-smokers – met the criteria for inclusion in our study. 1 subject was excluded after self-identifying as a regular smoker, because his WHO ASSIST Smoking score was lower than that of several self-reported non-smokers. No other cases of contradictory self-report and WHO ASSIST measures occurred. Table I presents summary statistics for the self-identified regular smokers and non-smokers in the sample. Smokers' mean score of 26.03 implies that they are at moderate risk for health

and other problems associated with their current pattern of substance use. 53 percent of smokers have scores on the tobacco module in excess of 27, which places them in the high risk category and implies they are likely to be dependent. Nine of the non-smokers reported having never tried cigarettes, and only 3 reported having had more than 2 cigarettes over the preceding 3 months. WHO ASSIST results thus confirm that self-judged smoking status in the sample reliably tracked status as measured by the screen.

TABLE I
SUMMARY STATISTICS AND SMOKER - NON-SMOKER COMPARISONS

| | Mean (Standard Deviation) | | Significant Difference? | |
|---|---|---|---|---|
| Variable | Smokers | Non-smokers | z-statistic or χ2 | p value |
| ASSIST tobacco score | 26.03 (6.44) | 3.03 (2.48) | z = -6.69 | **0.00** |
| Income | 2658.33 (2047.54) | 2458.33 (2891.55) | z = -1.40 | 0.16 |
| Strictly preferred SS | 29.13 (12.53) | 33.78 (15.14) | z = 1.63 | 0.10 |
| Age | 20.97 (1.52) | 21.23 (3.18) | z = -0.31 | 0.75 |
| Fraction selecting LL | 0.57 (0.50) | 0.60 (0.50) | χ2 = 0.02 | 0.79 |
| English first language | 0.83 (0.38) | 0.83 (0.38) | χ2 = 0.00 | 1.00 |
| White | 0.47 (0.51) | 0.37 (0.49) | χ2 = 0.61 | 0.43 |
| Black | 0.27 (0.45) | 0.57 (0.50) | χ2 = 5.55 | **0.02** |
| Male | 0.53 (0.51) | 0.53 (0.51) | χ2 = 0.00 | 1.00 |

Notes: Summary statistics computed from a sample of 60 subjects
Only wave 1 data is presented

Table I shows a significant difference between smokers and non-smokers in the proportion of Black subjects, with more Blacks among the non-smokers. We therefore control for race in the estimation framework to guard against misattribution of results.

Each group of smokers and non-smokers was randomly sorted into the three experimental conditions described below. Subjects then took part in a temporal discounting binary titration procedure to elicit a strict preference for a SS reward to be delivered in 1 day over a LL reward to be delivered in 10 days. We used a 1 day front-end delay (FED) to hold the subject's transaction costs constant for SS rewards and LL rewards, following Coller and Williams (1999). This FED also removes the possible influence of emotional arousal provoked by imminent reward (the putative visceral or "β" incentive as extensively discussed above). The LL reward was fixed at 50 South African Rand (R50)[7] and the SS reward varied according to the subject's choices. The titration implements a binary search algorithm that halves the difference between a subject's choices. Thus if a subject facing the choice between R25 in 1 day (SS) and R50 in 10 days (LL) chooses the LL reward, then the next choice the subject faces is between R37.50 in 1 day and R50 in 10 days. If instead the subject in the example chose the SS reward, the next choice would be between R12.50 in 1 day and R50 in 10 days. By this procedure one narrows in on the subject's indifference point between an SS reward in 1 day and R50 in 10 days. Lest the algorithm continue indefinitely, the titration terminated when it was forced to halve R0.50. Subjects were not told how their choices would be used, so as to avoid presenting

[7] 1 Rand had the domestic purchasing power of approximately $0.22 at the time of the experiment.

them with an incentive to misreport their preferences in order to be subsequently offered higher SS rewards.

Once the titration had reached its limit for a subject, that subject was again offered the smallest SS reward he or she had previously preferred to the LL reward, to ensure that the preference was stable. If the subject selected the LL reward, then the procedure was repeated until the subject selected the same SS reward over the LL reward twice in a row when the titration had reached its limiting value of R0.50. This selection established his or her baseline strict preference.

Past studies have shown that smokers tend to choose smaller SS rewards relative to LL reward magnitudes than non-smokers [22]. In our sample, however, the difference in strictly preferred SS rewards as between smokers and non-smokers was not found to be significant by a non-parametric Mann-Whitney test ($\chi^2 = 1.64$; $p = 0.10$).

The experimenter noted the subject's strict preference amount and then posed one of three condition-specific questions, depending on the condition into which the subject had been randomly placed:

FREE CONDITION:
"[As you know][8] I will be calling you every two weeks for six to ten weeks. Every second week I will be asking you to choose between *SS reward*[9] in one day and R50 in ten days. Which would you like me to give you this week: *SS reward* in one day or R50 in ten days?"

SUGGESTED CONDITION:
"[As you know] I'll be calling you every two weeks for six to ten weeks. Each week I will be asking you to choose between *SS reward* in one day and R50 in ten days. Each time you are offered this choice you will be in the same situation that you are now, facing a choice between *SS reward* in one day and R50 in ten days. Therefore, the choice you make now is the best indication of how you will choose every time. What somebody chooses one week is often what they go on choosing in later weeks, but you'll be completely free to choose between these two options every two weeks. Which would you like me to give you this week: *SS reward* in one day or R50 in ten days?"

FORCED CONDITION:
"You will now make a choice for a set of rewards. If you choose *SS reward* in one day then you will receive *SS reward* in one day and *SS reward* every two weeks after that for six to ten weeks. If you choose R50 in ten days then you will receive R50 in ten days and every two weeks after that for six to ten weeks. Which would you prefer: *SS reward* in

---

[8] Phrases in brackets were only used in follow-up phone calls.

[9] SS reward was subject-specific and is determined by the discounting titration procedure as described.

one day and every two weeks after that or R50 in ten days and every two weeks after that?"

All subjects were told they would make further choices or receive further prizes for the indefinite ‘six to ten weeks’ in order that all repeated choices in the Free and Suggested conditions were made in possible expectation of a similar future choice.

The three conditions allowed us to test hypotheses about the propensity of individuals to increase their preference for LL rewards when we varied the salience of the subjects’ current choices for their expectation of future rewards. The Free condition provided a comparison group without manipulations: individuals were simply told that we would contact them again to offer them the same choice. The Suggested condition primed subjects to think about their current choice as a predictor of future choices and thus framed an individual choice as occurring in a context of similar choices to be made in the future. The Forced condition determined all the future rewards by the subject’s current choice, forcing him or her to choose the rewards as a series.

SS rewards used were the smallest SS values reliably chosen (strictly preferred) over the LL reward in the binary titration session. From here on subjects actually received the money they chose, deposited directly in their bank accounts. After subjects in the Free and Suggested conditions made their initial choices in the presence of the experimenter they were contacted by telephone three more times and offered the same choice in intervals of two weeks. The two week interval was used to reduce the possible impact on preferences of accumulation of money over the course of the experiment. The data were gathered in four waves, at weeks 0, 2, 4, and 6. There was no attrition, so the panel of 60 individuals was balanced across the waves.

To test for any potential experimental confounds we investigated whether the variables from Table I differed significantly across the three experimental conditions into which subjects were sorted. The only significant differences that emerged were for income between the Free and Suggested conditions ($z = 2.59$; $p = 0.01$) and between the Suggested and Forced conditions ($z = -2.39$; $p = 0.02$) and for age between the Suggested and Forced conditions ($z = -2.02$; $p = 0.04$). We control for these factors in our statistical models, and find that they do not influence the experimental results of interest.

*Results*

Figure 1 displays the fraction of LL choices made during Wave 1 for smokers and non-smokers in the three experimental conditions. Among smokers, there is a significant increase in the fraction of subjects selecting LL between the Free and Forced conditions ($\chi^2 = 9.98$; $p = 0.00$) and between the Suggested and Forced conditions ($\chi^2 = 7.01$; $p = 0.01$). Among non-smokers there are no significant differences in the fraction selecting LL in Wave 1 between any of the experimental conditions. These results imply that, in Wave 1, smokers’ choice of LL was affected by the experimental condition whereas non-smokers’ choices were not.

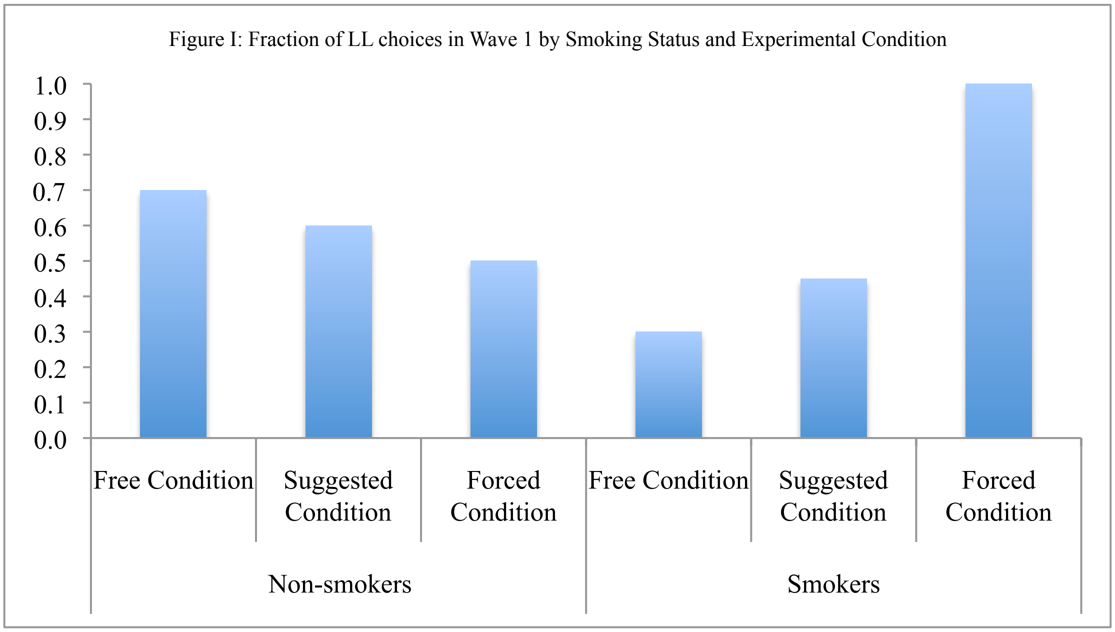


Directly comparing smokers and non-smokers across the experimental conditions at Wave 1, we find that, in the Forced condition, a significantly higher fraction of smokers selected LL than did non-smokers ($\chi^2 = 6.11$; $p = 0.01$). That is, smokers, when able to bindingly pre-commit to later choices, abandoned their strict (baseline) preferences for SS rewards at a markedly higher rate than non-smokers.

To clarify the effect of experimental condition and to test for the possibility of learning across waves in the experiment we present binary choice estimation models for smokers and non-smokers. We estimate a random effects probit model because of the dichotomous nature of our dependent variable, which follows the Bernoulli distribution, and to explicitly incorporate the longitudinal structure of the data. This model also allows us to control for, or partial out, the differences reported earlier in income and age across experimental conditions and in race between smokers and non-smokers.

TABLE II

BINARY CHOICE ESTIMATES FOR SMOKERS AND NON-SMOKERS

| Variable | Smokers | Non-Smokers |
|---|---|---|
| | Coefficient Estimates | |
| Suggested Condition | 0.65* | -2.73 |
| | (0.35) | (2.18) |
| Forced Condition | ∞*** | -0.39 |
| | | (1.31) |
| Black | -1.50*** | -1.36 |
| | (0.53) | (1.66) |
| Male | 0.44 | -1.51 |
| | (0.38) | (1.20) |
| English first language | 0.63 | -0.41 |
| | (0.55) | (1.68) |
| Age | 0.24** | -0.23 |
| | (0.11) | (0.18) |
| Income | 0.30 | 2.72* |
| | (0.30) | (1.41) |
| Strictly preferred SS | -0.049*** | -0.11** |
| | (0.018) | (0.047) |
| Wave 2 | 1.03** | 0.33 |
| | (0.47) | (0.57) |
| Wave 3 | 1.07** | 0.035 |
| | (0.47) | (0.57) |
| Wave 4 | 1.09** | 0.66 |
| | (0.48) | (0.58) |
| Constant | -7.06** | -8.06 |
| | (3.55) | (7.04) |
| Observations | 84 | 116 |
| Log Likelihood | -37,73 | -40,6 |
| | Marginal Effects | |
| Suggested Condition | 0.65* | -2.73 |
| Forced Condition | ∞*** | -0.39 |
| Black | -1.50*** | -1.36 |
| Male | 0.44 | -1.51 |
| English first language | 0.63 | -0.41 |
| Age | 0.24** | -0.23 |
| Income | 0.30 | 2.72* |
| Strictly preferred SS | -0.049*** | -0.11** |
| Wave 2 | 1.03** | 0.33 |
| Wave 3 | 1.07** | 0.035 |
| Wave 4 | 1.09** | 0.66 |

Notes:

Standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

∞ Forced condition is a perfect predictor and its coefficient tends toward infinity

Table II shows the results of our modeling, estimated on the smoker and non-smoker subsamples. In the model for smokers, the coefficient for participation in the Forced condition tends toward infinity because it is a perfect predictor of the response probability (i.e. all smokers in the Forced condition selected LL). We note that there is likewise a significant increase in the probability of selecting LL in the Suggested condition relative to the Free condition, our omitted category.

The results also show that the likelihood of selecting LL is significantly higher ($p < 0.05$) in Waves 2, 3 and 4 relative to the omitted category, Wave 1. This result confirms the pre-modeling observation of marginally significant pairwise comparisons in the Suggested condition, in smokers only, showing increased preference for LL rewards between Waves 1 and 2 and between Waves 1 and 3 ($\chi^2 = 3.14$; $p = 0.08$ for each comparison). This is depicted graphically in Figures II and III below, where the upward trend in the fraction of smokers selecting LL across waves is apparent.

Figure II: Fraction of LL choices in Free Condition by Smoking Status

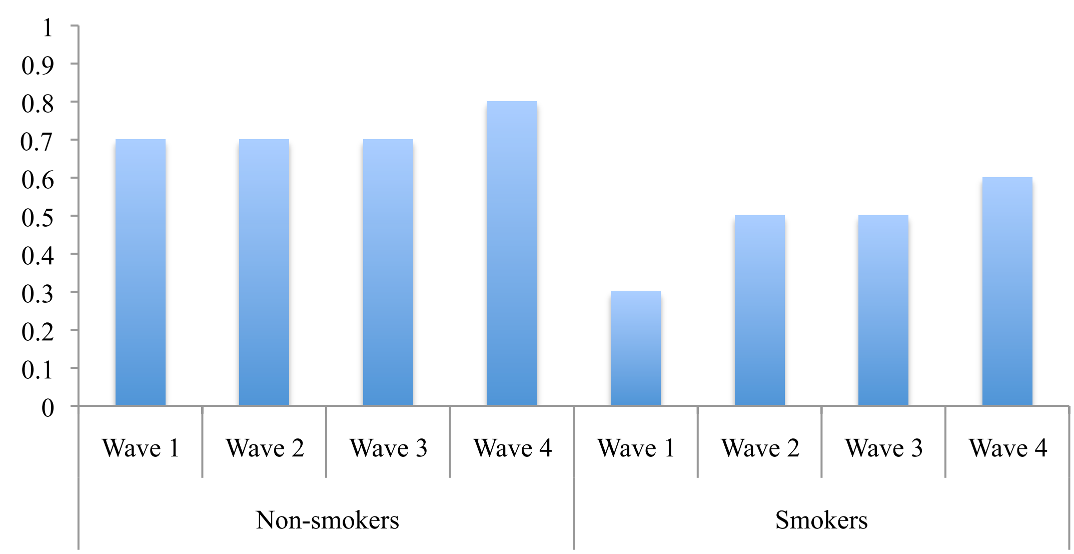

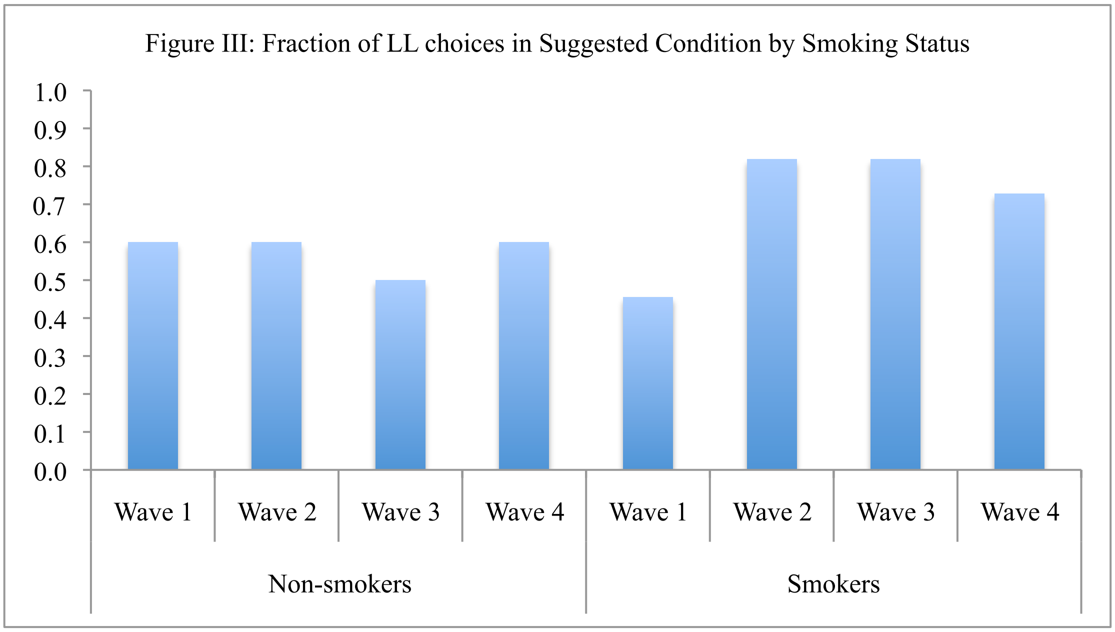


We also used the model to control for demographic differences between subjects. It is important to do this to ensure that the results of interest are not merely significant because of correlations with omitted demographic characteristics. In the model for smokers, being Black lowers the probability of selecting LL. This result highlights the importance of controlling for race in our model because the other variables would be biased in its absence. Being older increases the likelihood of selecting LL and a higher strictly preferred baseline SS lowers the probability of selecting LL. In the model for non-smokers, we find that only two variables influenced their choices in the experiment: having a higher income raises the probability of selecting LL whereas a higher strict preference at baseline has the opposite effect. By including these variables in our model, we partial out their effects and find that they are not driving the results.

In sum, the model indicates differential effects of experimental condition and repeated experience on smokers and non-smokers. Smokers are more likely to select LL in the Suggested and Forced conditions than they are in the Free condition. Furthermore, the probability that smokers select LL rises across repeated trials in the experiment. Non-smokers, on the other hand, are not influenced by the experimental conditions nor do they adjust their behavior significantly over time.

*Discussion*

Acknowledging the limitation of the small numbers in our study, our data suggest two important patterns. First, when the salience of a current choice to similar future choices is either suggested or offered as a binding commitment to subjects, smokers but not non-smokers make more future-oriented choices than they do in the absence of such external manipulation. This suggests that smokers may be less likely than non-smokers to spontaneously frame series of similar rewards as bundles.

Though other hypotheses about underlying choice-framing processes might be conjectured to explain this finding, its potential significance for treatment approaches to addiction and impulsivity is independent of specific interpretations. We suggest, however, that interpretation of these results in terms of assisted willpower is most appealing at this preliminary stage. It is not simply intuition that the perception of a current choice as a test case for a bundle of future choices should increase prudence. The hyperbolic discount functions that describe inverse proportionality to reward delay specifically predict this phenomenon.

*7. Conclusion*

One of the aspects of the BE literature that makes some economists uncomfortable with it – in our view, rightly – is its readiness to displace responses to incentives by exogenous hypothesized neurophsychological or other biological 'forces' (e.g., viscerality, β-spikes, limbic system discounting). This involves shifting phenomena that economists find difficult to model over into the province of psychology – and largely speculative psychology at that. This would be unobjectionable to economists if there were empirical evidence that the displaced processes are not elements of choice or are not influenced by incentives. However, as we have indicated, such evidence as has been cited for this claim is substantially weaker than evidence that points in the opposite direction.

By contrast, picoeconomics is about responses to incentives. If incentives are discounted in pure hyperbolae then there is no theoretical need for exogenous forces to account for inconsistent choice or dynamics of self-control. The sense in which picoeconomics complicates the picture of individual decision arises from its acknowledgment that choice sets can be recursively manipulated by choosers.

The seriousness of this complication should not be minimized. But there is *nothing* contrary to – or even in tension with – the subjectivist spirit animating the history of economics, from Walras through Pareto to Samuelson, in recognizing that people do not confront a world that comes uniquely parsed into choice sets by a metaphysical principle. It is a fact that a person can sometimes choose to jog now rather than watch television now. It is *also* a fact that she can choose to make it more likely that she'll be, in general, a jogger rather than a couch potato. It is a further fact that these real choices interact with one another in their implications for behaviour.

According to picoeconomics, framing and re-framing of choice sets isn't something that exogenously happens to people, as in most BE models. It is motivated by incentives and chosen.

Therefore, insofar as economists seek, at least in the limit, a general model of incentive-driven choice, they will need to find ways of coming to grips with picoeconomic phenomena, especially with reward bundling.

We squarely acknowledge that insofar as the picoeconomic framework demands representation by means of hyperbolic discount functions, economists will treat the framework as a problem description rather than a solution. In picoeconomics, the mathematical models that have been based on exponential functions are reduced to a special case (and not even the most forward-looking case at that – true hyperbolic functions assign relatively greater value to the distant future). As we have remarked at various points, hyperbolic functions resist easy integration into equilibrium analysis and econometric estimation and testing. At least as alarmingly, they do not promise the ability to predict an individual's choice; on the contrary, they depict individual choice as fundamentally unpredictable, including to the individual herself, and thereby explain people's apparent inability to shake off the conviction that they are endowed with free will. However, this should seem more ominous news to psychologists, whose business has in part been the prediction of individual choices, than to economists, who aim to model the general functional shape of the influence of incentives on choices.

We do not imagine that hyperbolic functions should sweep exponential models out of consumer theory. For one thing, such a suggestion would be counter to empirical evidence: as Vernon Smith (2008) has emphasized, people routinely collectively construct and maintain institutional and cultural scaffolding that helps them achieve "ecological rationality," which accounts for the efficacy of standard economic theory. In the language of picoeconomics, ecological rationality allows people to often exercise self-control, particularly under pressures of strategic interaction. We call upon economists to devote their modeling ingenuity to developing tractable formal machinery for representing this most subtle kind of choice that people constantly face, between framing opportunity costs in terms of discrete choices and framing them in terms of whole interconnected futures.

## ACKNOWLEDGMENTS

We gratefully acknowledge research support from the South African National Responsible Gambling Trust and from the National Research Foundation of South Africa.

## REFERENCES

Akerlof, G. (1991). Procrastination and obedience. *American Economic Review* 81: 1-19.

Ainslie, G. (1974). Impulse control in pigeons. *Journal of the Experimental Analysis of Behavior* 21: 485-489.

Ainslie, G. (1975). Specious reward: A behavioral theory of impulsiveness and impulse control. *Psychological Bulletin* 82: 463-496.

Ainslie, G. (1992). *Picoeconomics: The Strategic Interaction of Successive Motivational States within the Person*. Cambridge: Cambridge University Press.

Ainslie, G. (2001). Breakdown of Will. New York, Cambridge U.

Ainslie, G. (2010a). Hyperbolic discounting versus conditioning and framing as the core process in addictions and other impulses. In D. Ross, H. Kincaid, D. Spurrett & P. Collins, eds., *What Is Addiction?* Cambridge, MA: MIT Press.

Ainslie, G.(2010b). “Free will” as recursive self-prediction: Does a deterministic mechanism reduce responsibility? In G. Graham & J. Poland, eds., *Addiction and Responsibility.* Cambridge, MA: MIT Press.

Ainslie, G. (2010c). Procrastination, the basic impulse. In C. Andreou & M. White, eds., *The Thief of Time*. Oxford: Oxford University Press, pp. 11-27.

Ainslie, G., & Haendel, V. (1983). The motives of the will. In E. Gottheil, K. Druley, T. Skodola & H. Waxman, eds., *Etiological Aspects of Alcohol and Drug Abuse*. Springfield,IL: Charles C. Thomas, pp. 119-140.

Ainslie, G., & Herrnstein, R. (1981). Preference reversal and delayed reinforcement. *Animal Learning and Behavior* 9: 476-482.

Ainslie, G., & Monterosso, J. (2004). A marketplace in the brain? *Science* 306: 421-423.

Andersen, S., Harrison, G., Lau, M., & Rutström, E. (2008). Eliciting risk and time preferences. *Econometrica* 76: 583-619.

Ariely, D., & Wertenbroch, K. (2002). Procrastination, Deadlines, and Performance: Self-Control by Pre-Commitment. *Psychological Science* 13: 219-224.

Baumeister, R., Heatherton, T., & Tice, D. (1994). *Losing Control*. Burlington MA: Academic Press.

Baumeister, R., Gailliot, M., DeWall, C.N., & Oaten, M. (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.

Becker, G., and Murphy, K. (1988). A theory of rational addiction. *Journal of Political Economy* 96: 675-700.

Bénabou, R., & Triole, J. (2004). Willpower and personal rules. *Journal of Political Economy* 112: 848-886.

Benhabib, J. & Bisin, A. (2004). Modeling internal commitment mechanisms and self-control: A neuroeconomics approach to consumption-saving decisions. *Games and Economic Behavior* 52: 460-492.

Bernheim, B. D., & Rangel, A. (2004). Addiction and cue-triggered decision processes. *American Economic Review* 94: 1558–1590.

Bickel, W., Madden, G., & Petry, N. (1998). The price of change: The behavioral economics of drug dependence. *Behavior Therapy* 29: 545-565.

Bickel, W., & Marsch, L. (2001). Toward a behavioral economic understanding of drug dependence: Delay discounting processes. *Addiction* 96: 73-86.

Bodner, R., & Prelec, D. (2001). The diagnostic value of actions in a self-signaling model. In I. Brocas & J. Carillo, eds., *Collected Essays in Psychology and Economics*. Oxford: Oxford University Press..

Brocas, I., & Carrillo, J. (2008). The brain as a hierarchical organization. *American Economic Review* 98: 1312-1346.

Clark, A. (1997). *Being There*. Cambridge, MA: MIT Press.

Coller, M., & Williams, M. (1999). Eliciting individual discount rates. *Experimental Economics* 2: 107-127.

Cropper, M., Aydede, S., & Portney, P. (1992) Rates of time preference for saving lives. *American Economic Review* 82: 469-472.

Delaney, L., Rawdon, C., Denny, K., Zhang, W., & Roche, Richard. 2008. Event-related potentials reveal differential brain regions implicated in discounting in two tasks.

Geary Institute, University College Dublin, Working Paper 200811. http://ideas.repec.org/p/ucd/wpaper/200811.html
Deluty, M., Whitehouse, W., Millitz, M. & Hineline, P. (1983). Self-control and commitment involving aversive events. *Behavioral Analysis Letters*, 3: 213-219.
Ferris, J., & Wynne, H. (2001). *The Canadian Problem Gambling Index: Final Report*. Ottawa: Canadian Centre on Substance Abuse.
Fisher, I.(1930). *The Theory of Interest.* New York: Macmillan.
Frederick, S., Loewenstein, G., & O'Donoghue, T. (2003). Time discounting and time preference: A critical review. In G. Loewenstein, D. Read & R. Baumeister, eds., *Time and Decision: Economic and Psychological Perspectives on Intertemporal Choice*, pp. 13-88. New York: Russell Sage Foundation.
Fudenberg, D. (2006). Advancing beyond 'Advances in Behavioral Economics'. *Journal of Economic Literature* 44: 694-711.
Fudenberg, D., & Levine, J. (2006). A dual-self model of impulse control. *American Economic Review* 96: 1449-1476.
Hariri, A., Brown, S., Williamson, D., Flory, J., de Wit, H., & Manuck, S. (2006). Preference for immediate over delayed rewards is associated with magnitude of ventral striatal activity. *Journal of Neuroscience* 26: 13213-13217.
Heidhues, P., & Koszegi, B. (2009). Futile attempts at self-control. *Journal of the European Economic Association* 7: 423-434.
Glimcher, P., Kable, J., and Louie, K. (2007). Neuroeconomic studies of impulsivity: Now or just as soon as possible? *American Economic Review* 97: 142-147.
Green, L., Fisher, E.B., Perlow, S., & Sherman, L. (1981). Preference reversal and self control: choice as a function of reward amount and delay. *Behaviour Analysis Letters* 1: 43-51.
Green, Leonard and Myerson, Joel (2004) A discounting framework for choice with delayed and probabilistic rewards. *Psychological Bulletin 130,* 769-792.
Green, L., Myerson, J., & Macaux, E. W. (2005). Temporal discounting when the choice is between two delayed rewards. *Journal of Experimental Psychology: Learning, Memory, & Cognition* 31, 1121-1133.
Gul, F., & Pesendorfer, W. (2001). Temptation and self-control. *Econometric,* 69: 1403-1435.
Gul, F., & Pesendorfer, W. (2004). Self-control, revealed preference and consumption choice. *Review of Economic Dynamics* 7: 243-264.
Gul, F., & Pesendorfer, W. (2008). The case for mindless economics. In A. Caplin & A. Schotter, eds., *The Foundations of Positive and Normative Economics: A Handbook*, pp. 3-39. Oxford: Oxford University Press.
Harrison, G. (2008). Neuroeconomics: A critical reconsideration. *Economics and Philosophy* 24: 303-344.
Harrison, G., & Ross, D. (forthcoming). The methodologies of neuroeconomics. *Journal of Economic Methodology.*
Harvey, C. (1994). The reasonableness of non-constant discounting. *Journal of Public Economics* 53: 31-51.
Herrnstein, R. (1961). Relative and absolute strengths of response as a function of frequency of reinforcement. *Journal of the Experimental Analysis of Behavior* 4: 267-272.

Herrnstein, R. (1997). *The Matching Law: Papers in Psychology and Economics.* Edited by H. Rachlin & D. Laibson. New York: Sage.

James, W. (1890 / 1950). *The Principles of Psychology*. New York: Dover.

Johnson, A., & Redish, A. D. (2007). Neural ensembles in CA3 transiently encode pathes forward of the animal at a decision point. *Journal of Neuroscience* 12: 483-488.

Kable, J., & Glimcher, P. (2007). The neural correlates of subjective value during intertemporal choice. *Nature Neuroscience* 10: 1625-1633.

Kirby, K. (1997). Bidding on the future: Evidence against normative discounting of delayed rewards. *Journal of Experimental Psychology: General* 126: 54-70.

Kirby, K. (2006). The present values of delayed rewards are approximately additive. *Behavioural Processes* 72: 273-282.

Kirby, K., & Guastello, B.(2001). Making choices in anticipation of similar future choices can increase self-control. *Journal of Experimental Psychology: Applied* 7: 154-164.

Kirby, K. & Herrnstein, R. (1995). Preference reversals due to myopic discounting of delayed reward. *Psychological Science* 6: 83-89.

Laibson, D. (1994). *Hyperbolic Discounting and Consumption*. Ph.D. Thesis, MIT.

Laibson, D. (1997). Golden eggs and hyperbolic discounting. *Quarterly Journal of Economics* 6:, 443-479.

Laibson, D. (1998). Life-cycle consumption and hyperbolic discount functions. *European Economic Review* 42: 861-871.

Laibson, D., Repetto, A., & Tobacman, J. (1998). Self-control and saving for retirement. *Brookings Papers on Economic Activity* 1: 91-196.

Loewenstein, G. (1996). Out of control: Visceral influences on behavior. *Organizational Behavior and Human Decision Processes* 35: 272-292.

Loewenstein, G. (1999). A visceral account of addiction. In J. Elster & O.-J. Skog, eds., *Getting Hooked: Rationality and Addiction*. Cambridge, UK: Cambridge University Press.

Loewenstein, G., & O'Donoghue, T. (2004). Animal spirits: Affective and deliberative processes in economic behavior. *CAE Working Paper #0414*. http://ssrn.com/abstract=539843

London, E., Monterosso, J., Mann, T., Ward, A., Ainslie, G., Xu, J., Brody, A., Engel, S. & Cohen, M. (2006). Neural activation during smoking self-control. Poster at the 68th annual meeting of the *College on Problems of Drug Dependence*, Scottsdale, AZ, June 20.

Marshall, A. (1921). *Industry and Trade.* London: Macmillan.

Mazur, J. (1997). Choice, delay, probability, and conditioned reinforcement. *Animal Learning and Behavior* 25: 131-147.

McClure, S., Ericson, K., Laibson, D., Loewenstein, G., & Cohen, J. (2007). Time discounting for primary rewards. *The Journal of Neuroscience* 27: 5796-5804.

McClure, S., Laibson, D., Loewenstein, G., & Cohen, J. (2004). The grasshopper and the ant: Separate neural systems value immediate and delayed monetary rewards. *Science* 306: 503-507.

Metcalfe, J., & Mischel, W. (1999). A hot/cool-system analysis of delay of gratification: Dynamics of willpower. *Psychological Review* 106: 3-19.

Mischel, W., & Moore, B. (1980). The role of ideation in voluntary delay for symbolically-presented rewards. *Cognitive Therapy and Research* 4: 211-221.

Monterosso, J., Ainslie, G., Toppi- Mullen, P., & Gault, B. (2002). The fragility of cooperation: A false feedback study of a sequential iterated prisoner's dilemma. *Journal of Economic Psychology* 23: 437-448.

Musau, A. (2009). Modeling alternatives to exponential discounting. MPRA Paper No. 16416. Online at http://mpra.ub.uni-muenchen.de/16416/

Myerson, J., & Green, L. (1995). Discounting of delayed rewards: Models of individual choice. *Journal of the Experimental Analysis of Behavior* 64: 263-276.

O'Donoghue, T., & Rabin, M. (1999a). Procrastination in preparing for retirement. In H. Aaron, ed., *Behavioral Dimensions of Retirement Economics*, pp. 125-156. Washington: Brookings Institute and Russell Sage Foundation.

O'Donoghue, T., & Rabin, M. (1999b). Doing it now or later. *The American Economic Review,* 89: 103-124.

O'Donoghue, T.,& Rabin, M. (2000). The economics of immediate gratification. *Journal of Behavioral Decision-Making* 13: 233-250.

O'Donoghue, T. & Rabin, M. (2001). Choice and procrastination. *The Quarterly Journal of Economics* 116: 121-160.

Phelps, E. & Pollack, R. (1968). On second-best national saving and game-equilibrium growth. *Review of Economic Studies* 35: 185-199.

Prelec, D., & Bodner, R. (2003). Self-signaling and self-control. In G. Loewenstein, D. Read & R. Baumeister, eds., *Time and Decision*, pp. 277-298. New York: Russell Sage Foundation.

Rescorla, R.A. (1988). Pavlovian conditioning: It's not what you think it is. *American Psychologist* 43: 151-160.

Ross, D. (2005). *Economic Theory and Cognitive Science: Microexplanation*. Cambridge, MA: MIT Press.

Ross, D. (2010). Economic models of procrastination. In C. Andreou and M. White, eds., *The Thief of Time.* Oxford: Oxford University Press, pp. 28-50.

Ross, D., Sharp, C., Vuchinich, R., and Spurrett, D. (2008). *Midbrain Mutiny: The Picoeconomics and Neuroeconomics of Disordered Gambling*. Cambridge, MA: MIT Press.

Schelling, T. (1978). Economics, or the art of self-management. *American Economic Review* 68: 290-294.

Schelling, T. (1980). The intimate contest for self-command. *Public Interest* 60: 94-118.

Schelling, T. (1984). Self-command in practice, in policy, and in a theory of rational choice. *American Economic Review* 74: 1-11.

Schotter, A. (2008). What's so informative about choice? In A. Caplin & A. Schotter, eds., *The Foundations of Positive and Normative Economics: A Handbook.* Oxford: Oxford University Press, pp. 70-94.

Smith, V. (2008). *Rationality in Economics*. Cambridge: Cambridge University Press.

Strotz, R. (1956). Myopia and inconsistency in dynamic utility maximization. *Review of Economic Studies* 2:166-180.

Thaler, R., & Shefrin, H. (1981). An economic theory of self-control. *Journal of Political Economy* 89: 392-406.

Tolman, E. C. (1939). Prediction of vicarious trial and error by means of the schematic sowbug. *Psychological Review* 46: 318-336.

Trope, Y., & Fishback, A. (2000). Counteractive self-control in overcoming temptation. *Journal of Personality and Social Psychology* 79: 493-506.

Wertenbroch, K. (1998). Consumption self-control by rationing purchase quantities of virtue and vice. *Marketing Science,* 17: 317-337.

WHO ASSIST Working Group (2002). The Alcohol, Smoking and Substance Involvement Screening Test (ASSIST): development, reliability and feasibility. *Addiction* 97: 1183-1194.

.