

Economic models of procrastination

Don Ross

University of Cape Town

University of Alabama at Birmingham

don.ross@uct.ac.za

dross1@uab.edu

ABSTRACT

The paper compares a leading psychological account of procrastination and related phenomena, the *picoeconomic* framework due to George Ainslie, with economic models of self-control lapses generally, and procrastination specifically. By ‘economic’ I refer to models that admit of solution by maximizing utility functions, and which can be tested econometrically against behavioral or neural processing data. I suggest that bargaining games among picoeconomic interests as Ainslie depicts them can take two forms. One of these forms corresponds to dual self models of self control that have unique Markov Perfect Nash Equilibria as solutions. The other form has yet to be modeled, but promising tools for the job are in place.

1. Introduction

Almost all people procrastinate, some to the point of suffering severe welfare losses.¹ This fact has recently become a prominent topic at the interface between academic and popular discussion of economics, as attested by a wave of briskly selling books². Much of this discussion carries the explicit or implicit message that economics should be so embarrassed by the ubiquity of procrastination, as well as of addiction, loss aversion, salience effects, and other so-called ‘anomalies’ of rational choice, that it should collapse into psychology. In this paper, my modest aim is to survey recent work in the microeconomic modeling of procrastination so as to indicate resources available for dispelling such embarrassment.

Economic rationality implicates two aspects of agency: (i) consistency of choice from one consumption or investment opportunity to another, and (ii) use of full information in arriving at beliefs about the relative expected payoffs from

¹ Ted O’Donoghue, and Matthew Rabin, “Procrastination in preparing for retirement.”

² E.g. Dan Ariely, *Predictably Irrational*; Richard Thaler and Cass Sunstein, *Nudge: Improving Decisions About Health, Wealth and Happiness*.

possible choices (so called 'rational expectations'). The case for the displacement of economics by psychology rests in part on the claim that people are systematically inconsistent in their choices and prone to act on scanty information even when more could be obtained with relatively little effort.

The economist trying to model procrastination can trade off these aspects of rational agency against one another. On the one hand, she can suppose that people procrastinate in full awareness that they are doing so, but exhibit intertemporal preference inconsistency. That is, they choose courses of behavior that are rationalized by the payoff that would accrue to completion of various specific investments, then choose not to complete the investments in question, then subsequently reverse preferences a second time and suffer regret over their irresoluteness. Alternatively, procrastinators can be modeled as having intertemporally consistent utility functions but as lacking accurate information about their probable future behavior when they choose schedules of activities. Since what is mainly of interest is *recurrent* procrastination, the poverty of sound expectations here must be comparatively radical: the procrastinator fails to learn to predict future procrastination from her own history of past procrastination.

If effort is assumed to be aversive, then procrastination expresses at least temporary preference for smaller immediate rewards over later larger ones. This more general disposition has a range of other common manifestations, including addiction, impulsive spending, and impulsive risk-taking. It can be argued that procrastination is the basic form of this family of behaviors. For example, the addict who wants to discontinue use at some point, but not today, fails to quit because she continually procrastinates; and the difficulty of overcoming such procrastination is the main *problem* with addiction. Non-economists are often suspicious of economists' standard assumption that effort is aversive; artists, athletes, craftspeople and even businesspeople engrossed in their work tend to report when probed that such experiences of so-called 'flow' are highly rewarding. However, this simply makes the point that later rewards from avoidance of impulsive consumption, including impulsive consumption of leisure as in procrastination, *do* tend to be larger. Procrastination is in fact evidence for the aversiveness of effort.

The first extended account of procrastination in the economic literature was offered in a 1991 paper by Akerlof.³ He drops without apology the assumption that people are economically rational, emphasizing instead the extent to which they tend in nearly all their decisions to over-weight the inferential relevance of salient perceptions. A necessary condition for procrastination, according to Akerlof, is that people perceive immediately present costs of action more vividly than they perceive future benefits, or (especially) than they perceive the marginal values of increments of benefits that accumulate gradually. Akerlof doesn't explicitly address the question of why experiences of regret don't teach chronic procrastinators to discount vividly perceived immediate costs (or, alternatively, to augment dimly perceived future or incremental benefits). However, he indicates the conceptual resources available for

³ Akerlof, 'Procrastination and obedience.'

addressing this problem when he partly attributes procrastination to the fact that “individuals possess cognitive structures of which they are less than fully aware” (p. 10). He adds immediately that “[t]he assumption that such structures influence behavior is unfamiliar in economics, but central to other social sciences.” Thus Akerlof implicitly drops commitment to both aspects of economic rationality as relevant to explaining procrastination: procrastination is assumed to result from factors exogenous to rational choice.

In Akerlof’s model of the procrastinator’s revealed value function, costs of action made salient by temporal immediacy are multiplied by a parameter $\delta > 1$ that represents entrapment by a brute force of perceptual vividness. The parameter, which does all of the work in the model, amounts to a black box in which resides everything that personal psychology might identify as relevant to procrastination. The explicit model – the only part of Akerlof’s otherwise verbal account that has the flavor of economic analysis – thus sheds no extra light on the phenomenon; it simply encodes the fact that some psychological forces, and not merely relationships between valuations and prices, drive patterns of intertemporal choice. Much of the so-called ‘behavioral’ literature in economics continues in this vein, while adding richer psychological detail. As we will see, however, this is not the only available option.

I will organize the discussion as follows. Section 2 reviews a leading psychological approach to intertemporal self-control problems and regulation, including procrastination, the ‘picoeconomic’ framework descended from Herrnstein and most extensively developed by Ainslie. Section 3 identifies affinities between Ainslie’s picoeconomic interpretation of effective self-control (or ‘wilpower’) and the late Michael Bacharach’s model of team reasoning. Despite its name, picoeconomics does not provide a genuinely *economic* modeling framework. By an ‘economic’ model I mean a model that admits of solution by maximizing one or more utility functions (possibly interactively, using game theory), and which can be tested econometrically against behavioral or neural processing data. Section 4 considers ways in which the elements of the picoeconomic framework can be modeled economically. Section 5 concludes, arguing that economic analysis helps to clarify distinct modeling ideas within the picoeconomic framework that otherwise remain obscure.

2. Picoeconomics

There is a straightforward general strategy for economically modeling inconsistent people. This is 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.⁴ An analogy

⁴ Robert Strotz, “Myopia and inconsistency in dynamic utility maximization”; Thomas Schelling, “Economics, or the art of self-management”; Thomas Schelling, “The intimate contest for self-command”; Thomas Schelling, “Self-command in practice, in policy, and in a theory of rational choice”; George Ainslie, *Picoeconomics*;

can be drawn here between people and countries. It is obvious that countries often behave irrationally – erecting self-harming barriers to imports, for example, or alternately running down national savings and enforcing unpopular belt-tightening – due to the interactions of rational citizens acting in pursuit of their parochial interests. Nevertheless, for many purposes, especially macroeconomic policy analysis, economists model countries as agents. It is quite natural to talk of countries procrastinating when their institutions repeatedly defer hard policy choices. In these cases there is no mystery as to what is going on; we explain national procrastination by reference to strategic interactions of politicians, bureaucrats and lobbyists with divergent preferences. This can be adapted as a model for the dissolution of puzzles about weakness of will 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.

An early proponent of this approach to modeling procrastination and related phenomena was the economist Thomas Schelling.⁵ However, it has been most extensively developed by a group of psychologists whose intellectual roots lie in the laboratory of the late Richard Herrnstein. The group’s most notable contributors aside from Herrnstein are Drazen Prelec, George Ainslie and Howard Rachlin.⁶ Herrnstein’s group from the beginning made central conceptual use of a basic *economic* concept which Akerlof mentions but puts to no formal work, namely, discounting of future relative to present utility.

As Akerlof⁷ notes, 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

George Ainslie, *Breakdown of Will*; Don Ross, *Economic Theory and Cognitive Science: Microexplanation*; Don Ross, “Integrating the dynamics of multi-scale economic agency.”

⁵ See Schelling references in previous footnote.

⁶ Key highlights of this literature are collected as Richard Herrnstein, *The Matching Law*. See also George Ainslie, “Specious reward: A psychological theory of impulsiveness and impulse control”; and also Ainslie references in footnote 4 above.

⁷ Procrastination and obedience, p. 6

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

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

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.

Picoeconomic accounts of molar-scale intertemporal preference inconsistency model it around properties of the personal discount function. Based on experimental work, Herrnstein and Prelec⁸ and Ainslie⁹ argued that the default intertemporal discount function for animals, including people, is given not by the exponential function (1) but by a hyperbolic function as described by Mazur's formula¹⁰ (2):

$$v_i = \frac{A_i}{1 + kD_i} \quad (2)$$

A second exponent parameter on the right-hand denominator typically produces closer fits for empirical discounting data in people,¹¹ but I will henceforth ignore this technical complication. Hyperbolic intertemporal discounting allows for (though it does not entail) intertemporal preference reversals when agents choose between smaller, sooner rewards (SSRs) and larger, later ones (LLRs). A pair of temporally spaced rewards a [t_1], b [t_2] for which the person's utility function gives $b > a$ at a point well out into the future from the current reference point, where the slope of the discount function is relatively gentle, may swivel into the relation $a > b$ as the time of a 's possible consumption comes closer to the reference point, where discounting is steeper. Here b is a LLR – say, beginning a major home renovation – and a is a SSR – for example, spending the weekend at the cottage.

Figure 1 shows a standard diagram of the sort used by Ainslie to illustrate preference reversal. The short bar shows the value of an SSR, perhaps a bit of goofing off when there is work to be done. The long bar gives the value to be realized if procrastination is avoided. The crossing of the two hyperbolae illustrates preference reversal: at choice point Y perception of relative values is such that the agent will choose to avoid procrastination; at choice point X temptation looms large,

⁸ Richard Herrnstein and Drazen Prelec, "A theory of addiction."

⁹ Ainslie, *Picoeconomics*

¹⁰ James Mazur, "An adjusting procedure for studying delayed reinforcement."

¹¹ Joel Myerson and Leonard Green, "Discounting of delayed rewards: Models of individual choice."

and the agent postpones the task. At the point where the curves cross, we would expect a probe of the agent to find her wavering.

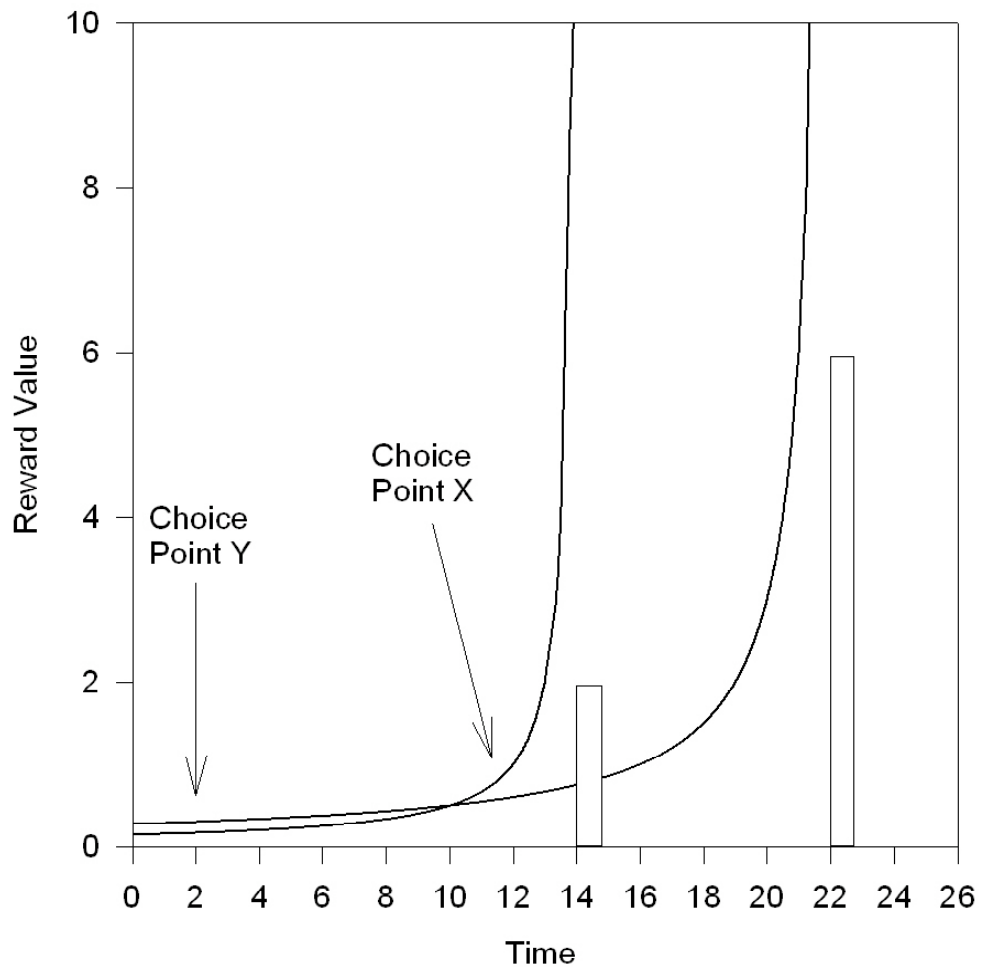


Figure 1: Preference reversal as crossing hyperbolae

Among the earliest empirical findings of the Herrnstein group¹² was that animals allocate responses to two choice alternatives proportionally to the relative frequencies of reinforcement received from them. The relation given in (3):

$$B_1/B_2 = R_1/R_2 \tag{3}$$

where the Bs and Rs refer to scalar measures of, respectively, responses allocated to and reinforcements received from the two alternatives, became known as the matching law because it claims that animals match relative behavioral allocation to

¹² Richard Herrnstein, "Relative and absolute strength of response as a function of frequency of reinforcement."

relative reinforcement. Herrnstein¹³ later showed that a choice rule that directly compares the local rates of reinforcement on the two alternatives and selects the alternative with the higher local rate describes the process that produces matching, which he termed *melioration*. Subsequent work building on these basic relations showed that we get accurate predictions of human and animal choice behavior if we model organisms as maximizing local rather than overall expected utility. The family of models that represent melioration specify that an animal will distribute its resources to a range of activities during a time interval in exact proportion to the value of the rewards it experiences for each activity *in that time interval*. Importantly, the time intervals over which local rates of reinforcement are computed and compared may vary in response to exogenous influences both across and within individuals. If variation in the temporal interval over which rates are computed changes the estimate of those rates, this alters the valuation of the choice options and leads to a change in behavioral allocation. Thus if the organism is modeled as a single agent then the economist's assumption of intertemporal preference consistency as a criterion of rationality may be violated.

Ainslie's piceoeconomic framework represents this situation as a bargaining game between two synchronous sub-personal interests, one of which (the 'short-range interest') has a utility function such that $SSR > LLR$ while the other (the 'long-range interest') has the opposite preference ordering over these alternatives. The short-range and long-range interests are motivated to bargain because the former controls behavior, but typically depends on resources harvested by the latter. However, their bargaining game has the structure of a Prisoner's Dilemma (PD): if the long-range interest will allow the short-range interest to obtain a payoff at *some* point, then the short-range interest prefers to obtain its payoff now, but that amounts to defection in the PD; if the long-range interest will never indulge the short-range interest in future then the short-range interest is also best off defecting at the first opportunity. Thus defection on any bargain with the long-range interest is a dominant strategy for the short-range interest. This in turn implies that the long-range interest never maximizes by indulging the short-range one; it also defects. The fact that it will not be indulged if it is patient reinforces the rationality of defection by the short-range interest.

This analysis predicts, among other things, that everyone should *always* procrastinate. Ainslie must therefore explain how it is that people frequently get things done (including breaking free of addictions). He attributes this to coalitions of long-range and short-range interests that form around *personal rules*. A personal rule may take two forms: a ban on consumption of a particular kind of SSR (e.g., cigarettes), or restriction on the circumstances under which the SSR may be indulged (e.g., drinking is only permitted on social occasions). For most people, an attempt at complete prohibition on procrastination would not be credible; therefore, in focusing on this most general expression of hyperbolic discounting, we will be especially interested in personal rules that make concessions to indulgence.

¹³ Richard Herrnstein, "Melioration as behavioral dynamism."

Ainslie argues that governance by personal rules can be made consistent with hyperbolic discounting if the agent recognizes that present behavior *predicts* future behavior. In that case, the person can derive *present* satisfaction from evidence that the personal rule is in place, making it a currently valuable asset. If the rule is broken, then this asset is damaged or destroyed near the reference point for discounting, where its value might thus dominate the value of the competing SSR. Ainslie¹⁴ analyzes the hoary philosophical and popular idea of ‘the will’ as a nominal ‘device’ for generating personal rules. He supposes that satisfactory explanation requires avoidance of the temptation to posit an *organ* to stand in for this device, and believes he achieves this by showing that personal-rule production and maintenance – ‘willpower’ – emerges endogenously as a virtual process from his picoeconomic bargaining dynamics.

3. Picoeconomics and team reasoning

The short account of picoeconomics given above does not remotely do justice to the empirical richness of Ainslie’s account, which is its most attractive and persuasive property. However, considering it from the perspective of economic analysis reveals some problems in need of solution.

Consider the kind of personal rule I suggested is most relevant to procrastination, in which bargaining occurs amongst interests. In a one-shot PD, there is no room for such bargaining; mutual defection is the unique Nash equilibrium. In personal correspondence, Ainslie replies to this by suggesting that games between interests in SSRs and interests in LLRs have the structure of a *repeated* PD, in which equilibrium vectors include cooperative strategies. However, this raises two further problems. First, it conflicts with Ainslie’s suggestion that interests, in having access to all of a person’s cognitive and behavioral resources, coincide with complete time-slices of a person. A present time-slice of a person does not recur, and so cannot play repeated games; interests can play such games only if they persist through time. Second, the repeated PD has an infinite set of Nash equilibria. Thus, limiting the analysis to identification of such games risks draining it of explanatory power.

A helpful idea for dealing with both of these problems is to combine the picoeconomic model with the ‘team reasoning’ model of Bacharach.¹⁵ Consider a person who establishes a policy to watch sports on television in the evening only after completing three hours of work, and to watch until no later than midnight. Suppose, contrary to Ainslie’s suggestion that interests coincide with personal time-slices, that the interest in watching sports has some *standing* interest in a present expectation that it will get to go on watching sports *generally* (though it always prefers watching sports *now*, in part because it cares *only* about watching sports). Then this interest faces the problem that its standing expectation is threatened

¹⁴ *Breakdown of Will*

¹⁵ Michael Bacharach, *Beyond Individual Choice*.

unless it allows long-range interests some scope to harvest resources. Given these assumptions, the person can emerge from the bargaining dynamics as a corporate institution with which the distinct short- and long-range interests partly identify. Formally, as Bacharach shows, this must amount (if it is to be coherent as non-cooperative game theory) to the claim that where personal rules are effective the bargaining game amongst the implicated interests has the structure of an Assurance Game instead of a PD.

Here is how this alchemy works. Begin with the familiar one-shot PD as in Figure 2, and imagine it as played between a long-range interest, Player I, and a short-range interest, Player II.

		II	
		C	D
I	C	2, 2	0, 3
	D	3, 0	<u>1</u> , <u>1</u>

Fig. 2: One-shot Prisoner’s Dilemma, strategic form

As usual, C denotes the strategy of ‘cooperating’ and D denotes ‘defecting’. The unique NE is (D, D). In equilibrium, to continue the example above, the short-range interest will watch sports all evening, no evening work will get done, revenue will dry up, and the TV will be re-possessed, putting an end to all sports-watching. If this in turn facilitates the resumption of work and the eventual recovery of financial well being, the long-range interest, playing D in the game, will not support investment in another TV.

Bacharach argues that under many circumstances people re-frame games by reasoning as *teams*, that is, asking ‘What is best for *us*?’ instead of ‘What is best for *me*?’. More precisely, players view enhanced welfare – and, perhaps, solidarity – of the collective as an additional basis for utility ranking. Such reasoning does not produce cooperation in one-shot PDs, which violates the axioms of game theory. Rather, by increasing payoffs associated with circumstances in which the team coordinates, it has the effect of transforming the one-shot PD of Figure 2 into the Assurance game of Figure 3. This has two NE, (C, C) and (D, D), where the former is Pareto efficient.

		II	
		C	D
C	<u>4</u> , <u>4</u>	0, 3	

I		
	D	3, 0
		<u>1</u> , <u>1</u>

Fig. 3: One-shot Assurance game, strategic form

Bacharach supposes that teams – corporate entities with which individual players identify their utility – arise from natural psychological framing by people. Thus if we try to appeal to Bacharach’s model in the piceoeconomic context, we must hypothesize that a person as a unified economic agent – a utility function that survives over multiple periods – arises from psychological framing by sub-personal interests. Many readers are likely to find this proposal incredible. In fact, I read Ainslie, in *Breakdown of Will*, as suggesting some affinity for the idea. It is consistent with the philosophical thesis of Dennett,¹⁶ substantially elaborated upon by Ross,¹⁷ that people are narrated post-hoc rationalizations of sequences of behavior, where the rationalizations in question feed back to influence behavioral choice because personal consistency is valued and incentivized by other people for the sake of social coordination (including division of labor), and most of the projects of loosely associated sub-personal interests rely on degrees of interpersonal collaboration. The sense in which ‘interests’ are the sorts of entities that could *frame* choices – that is to say, could make perceptual judgments – requires embedding in a certain kind of theory of consciousness, for which the reader is again referred to Dennett and Ross.¹⁸ This road evidently leads us well away from economic modeling considerations and deep into the foundations of psychology. Note, however, that it threatens no *collapse* of economics into psychology. The appeal to a theory of the relationship between consciousness and value is used to explain the genesis of basic, exogenous, utility functions. This reflects the traditional relationship of psychology to economics, in which tastes are treated by the economist as given. What is novel is that the utility functions in question are utility functions of sub-personal interests, with personal utility functions then emerging *endogenously* (i.e., within the *economic* model) through games amongst the interests.

Recourse to recursive processes is crucial if this style of modeling is to address the problem we started with. That problem has a chicken-and-egg logic: a bargaining game that doesn’t yield procrastination in equilibrium must include as a player a short-term interest in the maintenance of a personal rule; but according to the modeling framework characterized above, the person for whose welfare these rules are generated is herself a product of the dynamics of self-control. The answer to the puzzle lies in the fact that optimization in ranges of games that humans play give them standing long-range interests in consistency *and* short-range interests in

¹⁶ Daniel Dennett, *Consciousness Explained*.

¹⁷ Ross, *Economic Theory and Cognitive Science*.

¹⁸ See immediately preceding notes.

not risking embarrassment *now*. These general interests beget more specific interests in specific circumstances – e.g., an interest in preserving the rule that requires evening work before indulgence.

How might this picture be captured in formalism suitable for specifying games? Suppose to begin with that we allowed a long-range interest facing a game with short-range ones to summon into being by fiat a new agent with a new utility function, namely, a short-range interest in maintenance of a personal rule. In that case long-range interests would simply create whatever new short-range interests they needed as circumstances arose, and procrastination would never occur. (In Ainslie’s terms, all personal rules would be complete bans on indulgence in consumption of SSRs.) The analyst would still construct and solve games that generated procrastination in equilibrium; but these would be counterfactual games¹⁹ elaborated for the sake of inferring the actual assurance games, featuring personal-rule police, engineered by long-range interests to avoid internal PDs. More precisely, we would have the structure depicted in Figure 3, with the long-range interest receiving a higher payoff in the left-hand Assurance subgame than in the right-hand PD subgame. The long-range interest would always play left given subgame perfection, so procrastination would occur only off the equilibrium path and would not be observed. Let us call games with this generic structure ‘Ainslie-bargaining games’.

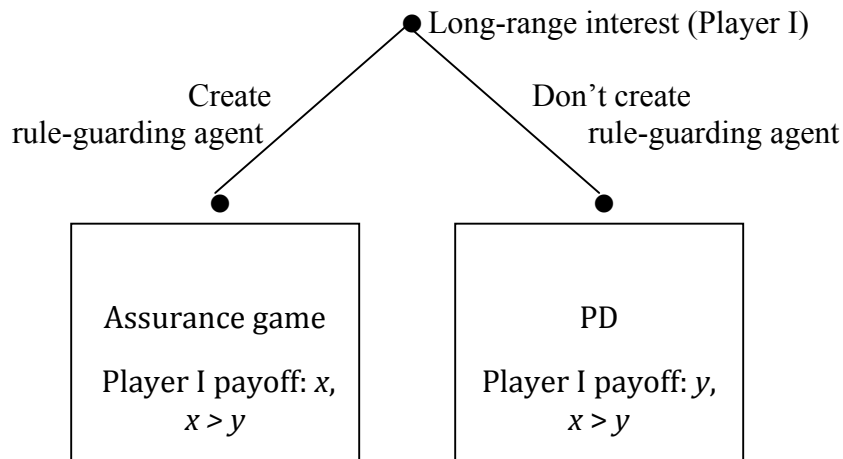


Figure 4: Ainslie-bargaining game

¹⁹ That is, they would be sub-games lying off the equilibrium paths of the larger games of which they formed parts.

It is easy to get from this limiting case of the model to a more satisfactory one in which procrastination sometimes or often appears: simply introduce a production function for the long-range interest such that creating short-range interests has non-zero cost. Depending on the relationship of these costs to other costs and benefits in the game, we can then get instances where long-range interests play to the right-hand side of Figure 3 and procrastination is observed.

4. Picoeconomics as economics

We have now arrived at a model of self-control we can compare with those that have recently become prominent in economic modeling of the problem. Let us begin with Fudenberg and Levine²⁰, who model the general family of impulse control problems, with procrastination as the simplest instance, 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’.) Since this is sufficient for unique equilibria in the games, there is no motivation for going all the way to imposing subgame perfection. Every SR-perfect NE profile is provably equivalent to a reduced-form optimization problem for the short-run self. 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 as I have reconstructed the latter. Ainslie 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 second difference is sufficient to generate different strategic choices by them. Fudenberg and Levine’s short-run selves are completely myopic. There is therefore no room in the dual control model for short-run selves that value personal rules, as in picoeconomics. Thus short-run selves are not equivalent to short-range picoeconomic interests. These differences in the models are reflected in divergent possibilities for philosophical interpretation. Fudenberg and Levine say that “while

²⁰ Drew Fudenberg and David Levine, “A dual self model of impulse control.”

we find the language of multiple ‘selves’ to be suggestive, the model can equally be interpreted as describing the behavior of a single ‘self’ whose overall behavior is determined by the interaction of two subsystems” (p. 3). As will be discussed below, the dual self model thus allows for a reductionist (molecular) interpretation that the piceconomic model does not.

Let us for now emphasize what the models have in common. Both explain the frequency and distribution of procrastination and related phenomena as a function of the *cost* of natural meliorative processes emulating expected utility maximization. Asked to explain why such emulation occurs, defenders of both models could appeal to the same evolutionary and institutional pressures operating on people. Both models, then, rely fundamentally on the idea that transcendence of matching is not free. Thus they are compatible with the dominant hydraulic perspective in the clinically focused psychological literature on self-control, according to which future-directed prudence requires expenditure of an energy resource that renews itself at a fixed rate and may thus be depleted if consumed at a faster rate (Baumeister *et al* 1994; Baumeister and Vohs 2007).²¹ (Fudenberg and Levine build this link explicitly, extending their model to make cost of self-control sensitive to current cognitive load.) Neither the dual self nor the piceconomic model require basic units of action that are irrational in the sense of having cycling preferences, or access to information that they willfully ignore or distort.

In constructing the Ainslie-bargaining game model, I suggested that short-range interests need to attach some present value to the future maintenance of their consumption patterns. Part of Ainslie’s motivation for resisting this suggestion emerges when we see how it is in tension with a favorite example of his, an annoying interest in scratching an itch. This interest, Ainslie emphasizes, will fade entirely if even briefly ignored; unless the itch is caused by a foreign irritant, as most itches are not, the interest in scratching *is* the itch. Ainslie characterizes itches as very short-range addictions. Failing to stop scratching at the current instant involves the same sort of procrastination, at a more micro scale, as failure to stop smoking today. If short-range interests are purely *instantaneous* interests then, as I argued in the general case of such interests, they can’t be bargained with because there is no bribe they will be around to collect. Now, in the case of itches it seems that bargaining indeed does *not* occur. An interest in scratching now isn’t bought off by an offer to scratch more at a less inconvenient time. Itches must be overcome by being ignored. One might of course say the same thing about many longer-range addictions; though allowing limited drinking works for some former alcoholics, others must suppress the behavior entirely. The dual self model handles this kind of case straightforwardly: the long-run self knows, but the short-run self does not, that scratching will give rise to further demand for scratching, and there is no self that

²¹ Roy Baumeister, Todd Heatherton and Diane Tice, *Losing Control*; Roy Baumeister and Kathleen Vohs, eds., *Handbook of Self-Regulation: Research, Theory and Applications*.

now wants to scratch later. Scratching will thus occur when and only when the long-run self declines to pay to arrest it.

What this suggests is that Ainslie-bargaining games aren't the only scenarios Ainslie has in mind as models of self-control. In some instances, instead of choosing to play Assurance games with a short-run interest B_1 , a long-range interest recruits a purely rival short-range interest B_2 to play a game that is zero-sum with respect to B_1 . Such recruitment cannot be costless on pain of vacuity. Thus one kind of self-control model Ainslie seems to have in mind, for application to some cases, is equivalent to the Markov games of Fudenberg and Levine. This point will emerge as important in later discussion.

Another problem that must be addressed if picoeconomics is to be rendered as economics 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 must 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 has a popular solution among economists who require workable functional forms for model testing but recognize the reality of the intertemporal preference inconsistency that procrastination implies. Laibson²² showed that the intertemporal preference reversals between SSRs and LLRs modeled by hyperbolic discount functions be represented by an alternative form of discount function called 'quasi-hyperbolic' or ' β - δ ', borrowed from Phelps and Pollack²³. This class of functions is expressed by

$$v_i = A_i \beta \delta^D$$

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.

One possible way of understanding β - δ discounting is as an approximation to 'true' hyperbolic discounting, adopted simply to achieve convergence of functions that sum over flows of values. On this interpretation, the hyperbolic function is

²² David Laibson, "Golden eggs and hyperbolic discounting."

²³ Edmund Phelps and Robert Pollack, "On second-best national saving and game equilibrium growth."

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. Phelps and Pollock developed the β - δ model in the first place to study intergenerational wealth transfers; and procrastination 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.²⁴

More recently, however, 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)²⁵ obtained fMRI evidence they interpret as suggesting that pre-frontal brain areas discount more steeply than frontal 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 pre-frontal and δ -discounting frontal sub-agents. This idea can in turn be given either of two generic kinds of interpretation in economic modeling. On the first interpretation, one *reduces* the piceoeconomic interests involved in self-control to the brain's rival discounting systems. This amounts to *eliminating* piceoeconomic models in favor of neuroeconomic ones. It entails a commitment to seeking a neural site for the kind of personal agency involved in the formation of personal rules, since people quite evidently *do* formulate, express and use such rules. Theoretical and neuroimaging study of working memory by psychologists has furnished some evidence that people can make conscious decisions that buttress their δ -discounting system's relative strength in internal conflicts.²⁶

The reductionist interpretation of the McClure *et al* findings has recently encountered empirical difficulties. Glimcher *et al*²⁷ report fMRI data they interpret as indicating that neurons in both midbrain and prefrontal areas in fact implement

²⁴ David Laibson, "Life-cycle consumption and hyperbolic discount functions"; David Laibson, Andrea Repetto, and Jeremy Tobacman, "Self-control and saving for retirement"; Ted O'Donoghue and Matthew Rabin, "Incentives for procrastinators"; Ted O'Donoghue and Matthew Rabin, "Choice and procrastination."

²⁵ Samuel McClure, David Laibson, George Loewenstein and Jonathan Cohen, "Separate neural systems value immediate and delayed monetary rewards."

²⁶ Randall Engle, "What is working memory capacity?"

²⁷ Paul Glimcher, Joe Kable and Kenway Louie, "Neuroeconomic studies of impulsivity: Now or just as soon as possible?"

similar discount functions as the molar subject. They find, as have others²⁸, 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). Delaney *et al*²⁹ 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.

An alternative interpretation of the McClure *et al* proposal, and one that is more interesting from the perspective of economics, associates it with a family of models that has recently come to dominate the economics of self-control problems. These treat pre-frontal processing as outside the boundary of the economic agent: β -discounting is understood as an *exogenous* challenge to personal utility maximization. Loewenstein³⁰ (along with Read³¹) points out that only certain sorts of goods are hyperbolically discounted at the molar scale: for example, desserts³² and hamburgers³³, but not petrol³⁴ or computer paper³⁵. The former sorts of reward are referred to as ‘visceral’: they are hypothesized to be perceptible *as* rewards with minimal cognitive processing. The idea is then that people may procrastinate in their saving due to being lured by visceral consumption (which includes

²⁸ Ahmed Hariri, Sarah Brown, Douglas Williamson, Janine Flory, Harriet de Wit and Stephen Manuck, “Preference for immediate over delayed rewards is associated with magnitude of ventral striatal activity.”

²⁹ Liam Delaney, Carolyn Rawdon, Kevin Denny, Wen Zhang and Richard Roche, “Event-related potentials reveal differential brain regions implicated in discounting in two tasks.”

³⁰ George Loewenstein, “Out of control: Visceral influences on behavior”; George Loewenstein, “A visceral account of addiction.”

³¹ Daniel Read, “Is time-discounting hyperbolic or subadditive?”

³² Stephen Hoch and George Loewenstein, “Time-inconsistent preferences and consumer self-control.”

³³ Read, ‘Is time-discounting hyperbolic or subadditive?’

³⁴ Hoch and Loewenstein, ‘Time-inconsistent preferences’.

³⁵ Read, ‘Is time-discounting hyperbolic or subadditive?’

consumption of idleness). If personal utility is identified with valuation resulting from relatively cognitive or 'cold' processing, then the siren call of visceral satisfactions can be treated as a form of *disutility* to the person, a threat to successful maximization lying in ambush in the 'external' environment within the brain.

There are several economic models that unpack this general idea in different specific ways. Gul and Pesendorfer³⁶ provide the most direct such approach. Their model defines a 'temptation' as a choice option with the property that its presence in the choice set makes an agent worse off, either because this results in her making a worse choice than she would have made in the option's absence, or because to cope with the option the agent must incur a cost of self-control. Thus the agent is incentivized to take steps to avoid encountering temptations. These may include personal-rule formation, along with other generic devices itemized by Ainslie: pre-commitment, control of attention, and preparation of emotion.³⁷ Gul and Pesendorfer point out that, despite predicting some of the same behavioral phenomena as they do, the piceoeconomic model admits no role for self-control *in their sense*, that is, as resistance 'at the point of choice' to a storm of visceral temptation. In the piceoeconomic model, overcoming 'weakness of will' is identified with reconfiguration of commodity spaces, with personal rules being added to possible consumption bundles. On Gul and Pesendorfer's model, by contrast, 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 can use as much information as the modeler thinks the data warrant.

A second approach that implicitly makes pre-frontal processing exogenous to the economic agent is promoted by Loewenstein³⁸, and explains self-control problems as informational deficits. According to this account, cognitive memory for visceral intensity is systematically unreliable. In consequence, people routinely underestimate the effect that past and future visceral states will have on future behavior, thus neglecting 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. Applied to procrastination, someone might be unable to anticipate that she will feel a temptation to put off work tomorrow that will be every bit as strong as today's feeling of temptation. This psychological hypothesis admits readily of an economic interpretation based around *self-signaling*.³⁹ This can in turn be modeled either in a

³⁶ Farouk Gul and Wolfgang Pesendorfer, "Temptation and self control"; Farouk Gul and Wolfgang Pesendorfer, "The simple theory of temptation and self-control."

³⁷ Ainslie, *Picoeconomics*.

³⁸ Loewenstein, 'A visceral account of addiction.'

³⁹ Drazen Prelec and Ronit Bodner, "Self-signaling and self-control."

unified or multiple agency setting. In the multiple (diachronic) agent context, present agents with some degree of special concern for their descendant 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 piceoeconomics, this explains personal rules as presently valuable assets. However, the potential value of self-signaling depends neither on hyperbolic discounting nor on dividing the person into intertemporally multiple agents. As Bénabou and Tirole⁴⁰ demonstrate rigorously, self-signaling can be justified by any imperfection in an agent's knowledge about her own capacity to cope with temptations.

Though self-signaling models preserve full agent rationality in not requiring intertemporally inconsistent preferences, they must clearly involve an hypothesis of compromised inferential rationality. Suppose we grant Loewenstein's 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 procrastination is chronic and recurring.

So far as difficulties for economic modeling are concerned, there is an asymmetry between relaxing the two basic aspects of rationality. Allowing preference reversal without decomposing the agent blocks application of standard tools altogether: we lose assurance that we can formulate and solve an optimization problem. By contrast, the only objection to relaxing rational expectations is that in many circumstances, such as modeling financial markets, any *specific* such relaxation seems arbitrary, and invites questions about why other agents haven't learned to exploit the particular less-than-rational expectations that are hypothesized. But if psychological research independently shows that people tend to make certain systematic errors, then the worry about arbitrariness goes away. And although chronic procrastination can lead to unnecessarily high interest payments, it does not undermine itself in the way that less-than-rational expectations about future prices do; procrastination is not necessarily disadvantageous in markets in which asset prices are random walks, and in other kinds of markets procrastinators and non-procrastinators don't generally face zero-sum games in which non-procrastinators must drive out procrastinators in equilibrium. Thus economists have been much more attracted to models of self-control problems with exogenous temptations than they have been to models involving hyperbolic discounting by unified agents.

O'Donoghue and Rabin⁴¹ have produced a specific model of procrastination based on imperfect expectations about the strengths of future temptations. They use

⁴⁰ Roland Bénabou and Jean Tirole, "Willpower and personal rules."

⁴¹ Ted O'Donoghue and Matthew Rabin, "Incentives for procrastinators"; Ted O'Donoghue and Matthew Rabin, "Choice and procrastination."

a β - δ discounting framework, and construct as their crucial variable the difference between β and a person's perception of β , represented as $\hat{\beta}$. A person for whom $\hat{\beta} = 1$ is completely naïve about her future behavior; though she procrastinates now, she never expects to procrastinate in the future. As noted above, this would be a peculiar sort of agent. The more interesting and realistic case is the person for whom $\hat{\beta} \in (\beta, 1)$. This person is *partially* naïve; she expects to be tempted to delay, but underestimates to some extent the force that the temptation will have on her. O'Donoghue and Rabin show that a partially naïve person will delay doing a worthwhile task indefinitely whenever she believes that her future tolerance for delay will be at least one period less than her actual (including her current) tolerance for delay. This implies that any degree of partial naiveté is sufficient to produce procrastination in some circumstances.

O'Donoghue and Rabin also expand the scope of the model to address the phenomenon of procrastination in a richer context than has been treated by other economic literature. They introduce the new element as follows:

[Previous] models of procrastination assume that a potential procrastinator has only one task under consideration, and hence the only concern is when the person completes the task. In most situations, however, a person must decide not only *when* to complete a task, but also *which* task to complete, or how much effort to apply to a chosen task. If a person must revise a paper for resubmission, she can either respond minimally to the editor's suggestions or expend more effort to respond thoroughly. If she is choosing how to invest some money, she can either thoughtlessly follow the advice of a friend, or thoroughly investigate investment strategies. If she is putting together a montage of Johnny Depp photos, she can either haphazardly throw together a few press clippings or work devoutly to construct the shrine that he deserves.⁴²

O'Donoghue and Rabin's analysis of procrastination in the context of choice among tasks produces two interesting results in their model, which the reader can evaluate for empirical plausibility. First, they show that a partially naïve agent who would complete a particular worthwhile task now might do nothing at all if offered a more attractive alternative action that had a higher cost relative to its immediate benefit. One possible sense of 'more attractive' is 'more *important*', in the sense of yielding higher long-run benefits. This leads to the second striking result in their extended (2001) model: higher importance can lead agents to choose costlier tasks, on which, if partially naïve, they then procrastinate. For example "the person may severely procrastinate [on investment choice] when her principal is \$10,000, but not when her principal is only \$1,000 ... [S]he plans and executes a quick-and-easy investment strategy for the \$1,000, while she plans – but does *not* execute – a more

⁴² O'Donoghue and Rabin, 'Choice and procrastination,' pp. 121-122.

ambitious investment strategy for the \$10,000" (p. 149). However, one should not infer from this that adding options tends *in general* to induce procrastination. In the model, a person with *many* tasks to be spread over several periods will tend to procrastinate *less*, because costs of delaying any task rise by threatening to crowd out other tasks.

Thus it should be acknowledged that combining imperfect expectations with β - δ discounting has produced the richest existing economic model of procrastination, which offers distinctive empirical predictions. Note, however, that in this instance β - δ discounting can be justified purely on grounds of analytical tractability; all of the model's insights would carry over if discounting were 'really' hyperbolic. Indeed, no aspect of discounting *per se* merits being regarded as the cause of procrastination in O'Donoghue and Rabin's model; that honor goes to false beliefs about how future opportunity costs will be perceived.

From the perspective of piceoeconomics, a straightforward criticism of the O'Donoghue and Rabin model can be given. This is that, like Fudenberg and Levine, they implicitly restrict attention to Markov Perfect Nash Equilibria in the choice games among the β -discounting and the various δ -discounting intervals of the person. That is, no choosing agent takes account of anything other than its utility function and whatever expectations it has about choices in future intervals; there is no room for any agent to learn the value of strategic self-commitment. But this is the core of the piceoeconomic account, at least insofar as it is expressed in analyses of the sort I reconstructed as Ainslie-bargaining games. The criticism applies generally to models that rely on a β - δ discounting framework.

Benhabib and Bisin⁴³ relax this restriction. Their highly general and elegant 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 β - δ 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

⁴³ Jess Benhabib and Alberto Bisin, "Modeling internal commitment mechanisms and self-control: A neuroeconomics approach to consumption-saving decisions."

games, even in circumstances where the latter never actually procrastinates. The reason is that the former agent, but not the latter, expects to be least likely to procrastinate when it matters most; therefore the marginal expected value of present savings are higher for her.

Benhabib and Bisin compare these generic differences in prediction with three bodies of empirical evidence. First, they note a systematic tendency for people to indulge in extra consumption, relative to the plans implied by previous consumption, that is greater for small windfalls than for larger ones. This is predicted by the Benhabib and Bisin model, and is opposite to the qualitative prediction of the O'Donoghue and Rabin model. Second, in the model restricted to Markov games, since β - δ discounters have no available mechanism for internal self-control, they can restrain their impulses only by buying illiquid assets. This implies that they should be willing to buy such assets even when they yield negative returns, up to the limit of the value of self-control. For agents who can adopt effective personal rules, illiquid assets are complementary self-control aids, so should be purchased only when they yield positive or slightly negative returns. Benhabib and Bisin cite data that are on their side in this respect. Finally, they note that in their model, over-stability of investment plans may be explained by a factor that is the opposite of procrastination: personal rules become more fragile as they become more complex, so agents who rely on them will tend to be inflexible and show over-commitment to prudent courses they have adopted. It is difficult to see how this last point can be genuinely said to favor the Benhabib-Bisin model, since it makes the same qualitative prediction as the rival model, but by a different mechanism. However, it seems we can say at least that the Benhabib and Bisin model of intertemporal resoluteness of choice fares just as well in the face of current evidence, both conventional economic and neuroeconomic, as the variants based on β - δ discounting.

4. Conclusion

Comparison of the precise analysis of Benhabib and Bisin with those of Fudenberg and Levine and O'Donoghue and Rabin allows us to appreciate an implicit underlying structure in Ainslie's account that is otherwise obscure. In reconstructing piceconomics game-theoretically, I argued that long-range interests might sometimes achieve self-control by inducing Assurance games, in which short-range interests are bought off, and sometimes achieve self-control by inducing games in which one short-range interest is set into a zero-sum game with another. The latter mode of self-control is essentially the case analyzed by imposing Markov Perfection as the solution concept, with the Fudenberg and Levine model as the general case, and the O'Donoghue and Rabin model as an application that captures further special features of the control of procrastination and other impulsive choice *when it is suppressed instead of bribed*.

However, we do not yet find in the literature an economic model of Ainslie-bargaining. Benhabib and Bisin, while sharing many of Ainslie's qualitative insights, do not imply any appeal to bargaining dynamics, since in their model temptations are simple exogenous threats to optimization rather than strategic agents in their

own right. Nor does their model suggest, like piceoeconomics, that procrastination and related phenomena are best explained by reference to properties of discounting.

I propose that one moral the piceoeconomist should take to heart from the attempt to provide an explicit economic account of his framework is that hyperbolic discounting *describes* procrastination but can't *explain* it. Certainly, Mazur's formula gives a better fit to most human and animal time preference data than the exponential discount function. And drawing hyperbolae is a very handy way of illustrating dynamic inconsistency of choice in behavior. However, hyperbolic discounting, unlike exponential discounting, involves more than *discounting*. Hyperbolic curves do not have the shapes and slopes that they do simply as a function of the contribution of the passage of time to estimates of relative risk. Their relationship mainly reflects contributions of psychological framing, which almost certainly is in turn influenced by properties of neural dynamics and by social learning. Because they summarize such a range of background influences, hyperbolic discount functions resemble Akerlof's δ parameter. To request an explanation of hyperbolic discounting and to request an explanation of procrastination is to ask for the same thing. And to explain how such discounting can be made approximately exponential is to explain how procrastination is avoided.

The point I am making can be put in a different way: the economist and the piceoeconomist do not disagree about how to represent discounting; rather, they mean slightly different things by the concept. The economist aims to refer to the effect of delay on present value *as mediated by risk*. The piceoeconomist refers indiscriminately to the entire set of influences of delay on predictions of future choice.

If we drop the idea that piceoeconomic and economic models differ from one another in empirically hypothesizing different styles of discounting, then it is possible to characterize the piceoeconomic account using the elements of economic models that have been reviewed here. Suppose that, by fiat, we follow most recent economists in identifying the person with the current long-range utility function. I suggest that no meaningful empirical question is begged by doing so as long as we don't foreclose the idea that (i) the short-range interests have different utility functions of their own, and (ii) the person is dynamically sculpted over time as a recursive output of bargaining among these interests. Then we can identify the piceoeconomic account as a variant of the Benhabib and Bisin model, in which the cost of self-control at any given point is a function of bargaining games, instead of simply an exogenously given price to be taken or left.

Are bargaining and non-bargaining models of self-control exclusive alternatives? The O'Donoghue and Rabin model predicts that a tendency to procrastinate should get worse if a person who isn't very busy formulates a plan with significantly higher long-run returns, but also higher-short-run costs, than the standing alternative activities. The Benhabib and Bisin model predicts the opposite. Forcing Ainslie's account into a stricter game-theoretic perspective helps us to see

that these may not be rival models of the same phenomenon, but models of different forms of self-control that may occur on different occasions.

What experiments does the piceoeconomic account suggest? Both Fudenberg and Levine and Benhabib and Bisin stress the significance of cognitive load as a major determinant of the cost of self-control. Both sets of authors recommend that neuroimaging studies be used to investigate this relationship. The relevant concept of cognitive load, however, is extremely vague. In typical experiments, it is proxied by asking subjects to do arithmetical puzzles or hold lists of items in memory. These aren't really the kinds of psychic stress that beset people faced with ecologically natural decisions. Picoeconomics suggests an alternative method: cue subjects with various short-range interests, some of which are in tension with long-range ones and some of which are complements. Does sheer complexity dominate in such cases, leading to impaired self-control, or is self-control capacity at least sometimes better predicted by the balance of forces? If the latter turned out to be the case, it would be natural to represent the contents of the 'cognitive load' black box as games among interests.

At the theoretical level, we might try to use the insights to which piceoeconomics has led us to produce a more general model of self-control than any we yet have. Suppose, as emphasized here, that self-control is sometimes implemented by suppressing short-range interests and sometimes by allowing them limited scope for expression. How might we try to construct a model that would allow us to analyze the determinants of the relative effectiveness, from case to case, of the two forms of self-control? We have already seen that Benhabib and Bisin's model can represent the Fudenberg and Levine model as a special case. An obvious way to try to formally unify the models is to nest Ainslie-bargaining within Benhabib and Bisin as another special case. Suppose that long-range interests, when they make offers to short-range interests instead of trying to overcome them, change the population of interests in the game, as suggested earlier. We might then set out to adapt an overlapping generations model, a standard part of the economist's toolkit, to represent such dynamics.

I suggest that the prospects for such modeling, along with the accomplishments already in hand that have been reviewed here, make the case that economics complements partner sciences in shedding light on procrastination and related phenomena, without simply abandoning its distinctive restrictions and collapsing into psychology.

References

Ainslie, George. 1975. "Specious reward: A psychological theory of impulsiveness and impulse control." *Psychological Bulletin* 82: 463-496.

Ainslie, George. 1992. *Picoeconomics*. Cambridge: Cambridge University Press.

Ainslie, George. 2001. *Breakdown of Will*. Cambridge: Cambridge University Press.

Akerlof, George. 1991. "Procrastination and obedience." *American Economic Review* 81 (2): 1-19.

Ariely, Dan. 2008. *Predictably Irrational*. New York: Harper Collins.

Bacharach, Michael. 2006. *Beyond Individual Choice*. Princeton: Princeton University Press.

Baumeister, Roy, Heatherton, Todd, and Tice, Diane. 1994. *Losing Control*. Burlington MA: Academic Press.

Baumeister, Roy, and Vohs, Kathleen, eds. 2007. *Handbook of Self-Regulation: Research, Theory and Applications*. New York: Guilford Press.

Bénabou, Roland, and Tirole, Jean. 2004. "Willpower and personal rules." *Journal of Political Economy* 112: 848-886.

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

Delaney, Liam, Rawdon, Carolyn, Denny, Kevin, Zhang, Wen, and 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>

Dennett, Daniel. (1991). *Consciousness Explained*. Boston: Little Brown.

Engle, Randall. 2001. "What is working memory capacity?" In Henry Roediger III and James Nairne, eds., *The Nature of Remembering: Essays in Honor of Robert G. Crowder*, pp. 297-314. Washington: American Psychological Association.

Fudenberg, Drew, and Levine, David. 2006. "A dual self model of impulse control." Harvard Institute of Economic Research Discussion Paper No. 2112. Available on Social Science Research Network http://papers.ssrn.com/sol3/papers.cfm?abstract_id=888752

Glimcher, Paul, Kable, Joe, and Louie, Kenway. 2007. "Neuroeconomic studies of impulsivity: Now or just as soon as possible?" *American Economic Review* 97: 142-147.

Gul, Farouk, and Pesendorfer, Wolfgang. 2001. "Temptation and self control." *Econometrica* 69: 1403-1436.

Gul, Farouk, and Pesendorfer, Wolfgang. 2005. "The simple theory of temptation and self-control." <http://www.princeton.edu/~pesendor/finite.pdf>

Hariri, Ahmed, Brown, Sarah, Williamson, Douglas, Flory, Janine, de Wit, Harriet, and Manuck, Stephen. 2006. "Preference for immediate over delayed rewards is

associated with magnitude of ventral striatal activity." *Journal of Neuroscience* 26: 13213-13217.

Herrnstein, Richard. 1961. "Relative and absolute strength of response as a function of frequency of reinforcement." *Journal of the Experimental Analysis of Behavior* 4: 267-272.

Herrnstein, Richard. 1982. "Melioration as behavioral dynamism." In Michael Commons, Richard Herrnstein and Howard Rachlin, eds., *Quantitative Analyses of Behavior, Volume II: Matching and Maximizing Accounts*, pp. 433-458. Cambridge, MA: Ballinger.

Herrnstein, Richard. 1997. *The Matching Law*. Cambridge, MA: Harvard University Press.

Herrnstein, Richard, and Prelec, Drazen. 1992. "A theory of addiction." In George Loewenstein and Jon Elster, eds., *Choice Over Time*, pp. 331-361. New York: Russell Sage Foundation.

Hoch, Stephen, and Loewenstein, George. 1991. "Time-inconsistent preferences and consumer self-control." *Journal of Consumer Research* 17: 492-507.

Laibson, David. 1997. "Golden eggs and hyperbolic discounting." *Quarterly Journal of Economics* 112: 443-477.

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

Laibson, David, Repetto, Andrea, and Tobacman, Jeremy. 1998. "Self-control and saving for retirement." *Brookings Papers on Economic Activity* 1: 91-196.

Loewenstein, George. 1996. "Out of control: Visceral influences on behavior." *Organizational Behavior and Human Decision Processes* 65: 272-292.

Loewenstein, George. 1999. "A visceral account of addiction." In Jon Elster and Ole-Jorgen Skog, eds., *Getting Hooked: Rationality and Addiction*, pp. 235-264. Cambridge, MA: Cambridge University Press.

Mazur, James. (1987). "An adjusting procedure for studying delayed reinforcement." In Michael Commons, James Mazur, John Nevin, and Howard Rachlin, eds., *Quantitative Analysis of Behavior Vol. 5: The Effect of Delay and of Intervening Events on Reinforcement Value*, pp. 55-73. Hillsdale, NJ: Lawrence Erlbaum Associates.

McClure, Samuel, Laibson, David, Loewenstein, George, and Cohen, Jonathan. 2004. "Separate neural systems value immediate and delayed monetary rewards." *Science* 306: 503-507.

Myerson, Joel, and Green, Leonard. 1995. "Discounting of delayed rewards: Models of individual choice." *Journal of the Experimental Analysis of Behavior* 64: 263-276.

O'Donoghue, Ted, and Rabin, Matthew. 1999a. "Procrastination in preparing for retirement." In Henry Aaron, ed., *Behavioral Dimensions of Retirement Economics*, pp. 125-156. Washington: Brookings Institute and Russell Sage Foundation.

O'Donoghue, Ted, and Rabin, Matthew. 1999b. "Incentives for procrastinators." *Quarterly Journal of Economics* 114: 769-816.

O'Donoghue, Ted, and Rabin, Matthew. 2001. "Choice and procrastination." *Quarterly Journal of Economics* 116: 121-160.

Phelps, Edmund, and Pollack, Robert. 1968. "On second-best national saving and game equilibrium growth." *Review of Economic Studies* 35: 201-208.

Prelec, Drazen, and Bodner, Ronit. 2003. "Self-signaling and self-control." In George Loewenstein, Daniel Read and Roy Baumeister, eds., *Time and Decision*, pp. 277-298. New York: Russell Sage Foundation.

Read, Daniel. 2001. "Is time-discounting hyperbolic or subadditive?" *Journal of Risk and Uncertainty* 23: 5-32.

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

Ross, Don. (forthcoming). "Integrating the dynamics of multi-scale economic agency." In Harold Kincaid and Don Ross, eds., *The Oxford Handbook of Philosophy of Economics*. Oxford: Oxford University Press.

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

Schelling, Thomas. 1980. "The intimate contest for self-command." *Public Interest* 60: 94-118.

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

Strotz, Rpbert. 1956. "Myopia and inconsistency in dynamic utility maximization." *Review of Economic Studies* 23: 165-180.

Thaler, Richard, and Sunstein, Cass. 2008. *Nudge: Improving Decisions About Health, Wealth and Happiness*. New Haven: Yale University Press.