Procrastination, uncertainty, and optimization

George Etheredge[1] and Don Ross[1,2]
1 – School of Economics, University of Cape Town
2 – Center for Economic Analysis of Risk, Georgia State University

1. Introduction

Ainslie (2010) and Ross (2010) argue that the basic logical and psychological structure of addiction, regretted binges, and other common forms of consumption that involve recurrent alternation of indulgence and regret is exhibited in its most transparent form by procrastination. Put another way, procrastination encapsulates the distinctively economic element of all of these common self-management syndromes.

By a distinctively economic element we refer to challenges the syndromes raise to modeling individual consumers as dynamic optimizers. As emphasized by the well known 'rational addiction' model of Becker & Murphy (1988), the economist need not struggle to model an agent merely because she chooses in a way that causes her welfare to steadily and predictably decline, provided we allow for one estimation error at the origin of the consumption path and acknowledge that correcting this error might require an exogenous augmentation of the capital stock that becomes increasingly far out of reach as the effects of the initial error are amplified by sequences of choices. However, as Ross (2011) argues, this is essentially a model of habit-formation rather than addiction. A rational addict who somehow does acquire the needed exogenous capital – that is, who successfully pays all of the costs of withdrawal and becomes physiologically 'clean' – would not be predicted by the Becker-Murphy to then repeat their initial estimation error and relapse. Yet most addicts successfully complete withdrawal 3-7 times before eventually finding a less tormented equilibrium (Heyman 2009). This pattern *is* a source of consternation to the economic modeler. The recurrently quitting and relapsing addict, who at all times is paying costs both to consume her addictive good *and* paying costs to try to stop or reduce her consumption, can only be thought to be optimizing if we suppose that she periodically quits in order to reduce her tolerance and thereby improve her cost-benefit ratio from addictive consumption. But this behaviour is not recognized by any currently accepted clinical model of addiction.

Once we recognize dynamic inconsistency as the distinctively economic problem raised by the typical life-cycle of an addiction, we can quickly grasp the motivation for framing it as a form of procrastination. The addict prefers to become a non-addict. She knows this will require aversive effort, which she would like to minimize. She also enjoys consuming her addictive target, particularly as her patterns of social interaction are likely to be oriented around this consumption. The marginal cost to her of postponing the commencement of effort by a small increment is significantly greater than the marginal benefit from starting slightly later rather than slightly sooner. Thus the addict postpones payment of some or all of the costs associated with an effort to quit that has good prospects of success. Like most procrastinators, she does not entirely give up in the face of this dilemma, but tries to shave costs. For example, a smoker might throw away her cigarettes and tell her friends she's quitting, but then go the pub in which her friends are smoking while drinking and socializing, thus dramatically reducing the likelihood that her resolution will persist[1].

Addiction is a relatively complex form of procrastination. This is partly due to its special psycho-neural properties, which interact with its economic ones. It is also due to the fact that the addict typically finds out how much room for procrastination she had before bringing catastrophe down upon herself only well after the period in which her choice could make a significant difference. Most people who stop smoking probably could 'get away with' smoking – in the sense of paying a small health and monetary cost they would regard as acceptable – for somewhat longer than they do, but are risk averse in the face of uncertainty; and this risk aversion implies that their uncertainty will never be resolved. Very few people's utility functions are consistent with a plan to smoke without concern for life, so planning to quit at some point is the easy part of the smoker's economic problem. The *very* hard part is selecting the optimal such point from the set of all of the days ahead of her.

To the extent that addiction is a kind of procrastination, then, we might best make progress by beginning with a simpler form of procrastination problem that idealizes away the peculiar dynamics of the timing of information receipt that bedevil the typical addict. We consider, then, a person who must complete a project by an exogenously fixed deadline, for example a student or an employee with an assignment and a due date. But it is crucial to our motivation to stress that the deep informational uncertainty about deadlines is the only aspect of the addict's economic problem we claim to be abstracting away. Otherwise our common-and-garden procrastinator models the typical addict from an economic point of view. Just as the addict wants, in the end of the day, not to be addicted, so the everyday procrastinator wants, in the end of the day, to complete all of the projects she is assigned (and self-assigned) by their deadlines. But – perhaps because she knows that she works most efficiently under pressure – she wants to enjoy as long a period of leisure as she can before she starts to work. Like the addict, she confronts uncertainty: she knows she must start to work at some time-point within a set, but, since she has not fully scoped the demands of the project by seriously starting to do it, she is unsure which point in the set is the optimal one. Like the addict, she has probably acquired some information relevant to her estimation problem through previous episodes of procrastination[2]. The only core *economic* difference between the common-and-garden procrastinator and the addict – or so we claim – is that the former

[1] The psychological and neurodynamic basis for the cue-sensitivity of addictive cravings is broadly understood; see Ross et al (2008), among many possible sources.

[2] For most addicts, learning from unsuccessful attempts to quit is likely a necessary condition for eventual success. The friends and families of addicts typically find the relevant learning rates to be despairingly slow.

but not the latter expects and receives clear feedback on whether she missed the optimum, in cases when procrastination is indulged for too long, in time for this information to be applied to the next similarly structured episode. She also receives less decisive evidence of mis-estimation whenever she finishes her task on time without experiencing significant stress. We suggest that this difference involves the kind of subtlety that is best postponed for a refinement of the basic model.

The paper proceeds as follows. In Section 2 we compare our approach to the dynamic inconsistency involved in procrastination and addiction to prevailing approaches in the literature. In particular, we attend to tensions between psychological and economic models that partly motivate our approach. In Section 3 we explain the structure and assumptions built into our model, and indicate what it does and doesn't aim to accomplish. Section 4 states the model and proves that it establishes the intended relationships. Section 5 briefly concludes.

2. Picoeconomic and economic models of procrastination and addiction

Economists who have considered procrastination and addiction have of course noticed that dynamic inconsistency with respect to welfare optimization is the crux of the challenge. The most influential models (O'Donoghue & Rabin 1999, 2001; Bénabou & Tirole 2004; Fudenberg & Levine 2006) have responded to this by the straightforward device of dividing the agent into two or more sub-agents, based on differences in time-preference. It is evident that if a person's behaviour results from the outcome of competitive bargaining or warfare amongst sub-agents, the result may be apparently ambivalent choice patterns that can be modeled in terms of multiple equilibria. In adopting this approach, economists have aimed to reconcile a dominant conceptual approach promoted by Ainslie (1992, 2001) with the need to arrive at testable predictions through identification of a value function that monotonically converges as observations of choices are taken into account (Ross 2010; Ainslie 2011a). Ainslie's own 'hyperbolic discounting' model does not satisfy, and is not intended to satisfy, this desideratum.

Before we turn to some critical comments on this modeling approach, we take note of a main alternative found in the economic literature. Gul & Pesendorfer (2001, 2004), Benhabib & Bisin (2004) and Bernheim & Rangel (2004) model addiction as resulting from a unified agent's failure to muster resources sufficient for overcoming exogenous temptations to enjoy counter-optimal visceral sensations, the lure of which was inadequately estimated or accounted for in ex ante planning. These models, relying as they do on information uncertainty of a particular kind, have an abstract family resemblance to the one we develop in the next section. However, they have significant limitations. First, they ignore the economic structure that addiction shares with procrastination, as emphasized above. Most everyday procrastinators are not plausibly modeled as confronting vivid temptations to storms of temporary pleasure that could not have been anticipated. Second, these models downplay instead of helping to rationalize and explain the cunning, highly foresighted exercises in re-framing values and alternatives that, as Ainslie (1992) illustrates, characterize much of the self-conscious life of modern individuals, including addicts (Ainslie 2011a, 2011b). Exogenous temptation models emphasize one strategy that foresighted people sometimes adopt when faced with threats to counter-impulsive consumption scheduling, namely avoidance. However, as Ainslie (2011a) discusses in detail, it is at least as common for people to concentrate on potential triggers of impulsive choice in order to carefully construct superior alternatives. A device Ainslie has long emphasized is reward *bundling*, the re-framing of current choices as implicating the stakes in expected similar future choice situations in which outcomes are predicted on the basis of the present decision. We do not mean to suggest that exogenous temptation models fail to capture the dynamics of some, or even very many, cases in which people are at risk of impulsively undermining their welfare. Our point is only that they reach for a lower level of generality than interests us here.

Let us therefore return to models based on bargaining or other forms of competition among sub-agents. We will concentrate on the details of Ainslie's (1992, 2001) account, despite the fact that it is not an economic model, for reasons we will indicate. However, it is the descriptively richest such work in the literature, and in that respect provides the fullest identification of the ultimate modeling target; furthermore, our motivations for modifying one of its core representational features in our own modeling apply in one step to all of the other economic models that Ainslie's work has inspired, as will be clear later.

Ainslie's 'picoeconomics' derives from a tradition in psychology that has its proximate intellectual roots in the laboratory of the late Richard Herrnstein (see Herrnstein 1997). Herrnstein and his students from the outset made central conceptual use of a basic economic concept, namely, discounting of future relative to present utility.

To the extent that uncertainty of expectations increases with the passage of time forward from the moment of choice, a rational agent should discount the value of future rewards relative to present ones. This does not undermine either static or dynamic consistency if the discount function is linear in the time argument, as in the standard exponential function given by

$$v_i = A_i e^{-kD_i} \tag{1}$$

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.

Ainslie accounts for preference ambivalence on the basis of a drastically modified representation of the personal discount function. Based on experimental work, Ainslie (1992) 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 (among other variants found in the literature) Mazur's (1987) formula (2):

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

Hyperbolic intertemporal discounting is compatible with intertemporal preference reversals when agents choose between smaller, sooner rewards (SSRs) and larger, later ones (LLRs). A pair of temporally spaced rewards $a[t1]$, $b[t2]$ for which the person's utility function gives $b \succ a$ at a point in the future from the current reference point, where the slope of the discount function is relatively gentle, may swivel into the relation $a \succ b$ as the time of $a\prime s$ possible consumption comes closer to the reference point, where discounting is steeper. Here b is a LLR – say, completing one's assignment by the deadline – and a is a SSR – for example, sharing images and remarks with one's Facebook friends.

Figure 1 shows a standard diagram of the idiosyncratic style that Ainslie uses for depicting the procrastinator's or addict's preference ambivalence. The short bar shows the value of an SSR, such as the time on Facebook. The long bar represents, in the present example, the value associated with completion of the assignment on time. The crossing of the two hyperbolae allows for intertemporal preference reversal: at choice point **Y** perception of relative values reflects the agent's preference for shutting down her browser while she still has adequate time for completing the assignment. However, the relative values at choice point **X** indicate that the agent will, at that reference point, attach higher marginal value to further chatting and posting, so that she may put her LLR at risk by further procrastination. At the point where the curves cross, **X** on the graph, we would expect a probe of the agent to find subjective experience of internal conflict.

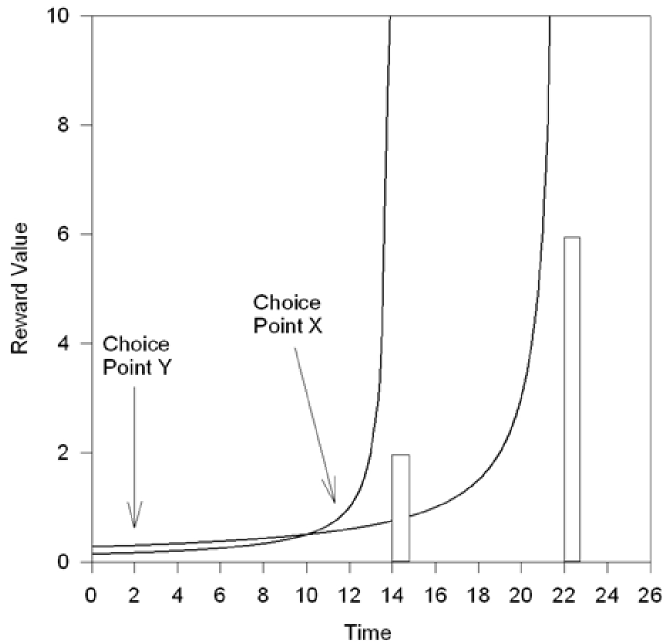

Figure 1: Preference ambivalence as crossing hyperbolae

Ainslie's picoeconomic 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. (In our example, a person who neglects her work for constant immersion in Facebook may find herself unable to pay for her internet subscription or electricity.) 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, which constitutes 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 the short-range interest will not be indulged if it is patient reinforces the rationality of its defection.

This analysis predicts, among other things, that everyone should *always* procrastinate. Explaining how its that most people complete most of their assigned and self-assigned tasks – including breaking free of addictions, which most addicts eventually do without clinical assistance (Heyman 2009) – is Ainslie's primary explanatory objective. He attributes this to coalitions of long-range and short-range sub-personal interests that form around *personal rules*. A personal rule might be a ban on consumption of a particular kind of SSR (e.g., drop one's Facebook subscription), or a restriction on the circumstances under which the SSR may be indulged (e.g., Facebooking is only allowed between 10 and 11 p.m. each evening). For most people, an attempt at complete prohibition on procrastination would not be credible, so a workable personal rule should allow for loopholes. This gives rise to a higher-order self-management problem, namely rationalizations of expansions of loopholes (e.g., an escape clause such as , "Of course I can stay online beyond the hour if the person I'm hoping to romantically attract starts posting messages" may encourage the person to liberally define her set of romantic targets). Recognition of that risk might in turn produce an opposite problem, pathologically rigid self-repression (e.g., "My closest friend is evidently in desperate need of comforting sentiments, but it's 10:59 p.m.").

Ainslie argues that maintenance of stable personal rules, and hence control of procrastination, does not imply that underlying discounting ceases to be hyperbolic. Rather, the hyperbolically discounting agent recognizes that present behavior *predicts* future behavior in similar situations. In consequence, she derives *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. According to Ainslie (2001), the traditional idea of 'the will' finds scientific vindication as a nominal 'device' for generating personal rules. The device in question is not to be reified as an imagined neural processing module or other 'organ'. Rather, the idea is that willpower is shown to have a virtual reality if activities of personal-rule production and maintenance are found to emerge endogenously when people make intertemporal choices.

This is evidently not an economic model, for two main reasons. First, the value function implied by hyperbolic discounting is not guaranteed to converge in the limit as new observations of behaviour are entered as data. Thus the model cannot be used to restrict econometric estimation. Second, the hyperbolic discount function is latent, since we are to suppose that it governs underlying valuation even when choice is governed by a personal rule that enforces economic consistency. Modeling the intertemporal utility function as incorporating hyperbolic discounting under these circumstances violates both the weak and the generalized axioms of revealed preference.

The models due to O'Donoghue & Rabin (1999, 2001), Bénabou & Tirole (2004) and Fudenberg & Levine (2006) all follow Ainslie in modeling procrastination and its related intertemporal preference ambivalence phenomena as expressions of complex discounting and sub-personal disunity. The models brought into conformity with economic modeling principles resort to an alternative form of discount function called 'quasi-hyperbolic' or '$\beta - \delta$', borrowed from Phelps and Pollack (1968) and applied individual agent modeling by Laibson (1997). This class of functions is expressed by

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

where v_i, A_i and D retain the interpretations as from (1) above, β is a constant discount factor for all delayed rewards, and δ is a per-period exponential discount factor. Where $\beta = 1$ the equation reduces to standard exponential discounting. Where $\beta < 1$ discounting is initially steeper up to some inflection point, then flattens. $\beta - \delta$ 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. It has recently been suggested that, unlike Ainslie's virtual sub-agents, the β and δ discounters are distinct neural areas (McClure et al 2004). However, other neuroimaging studies specifically aimed at testing this hypothesis do not find support for it (Glimcher et al 2007).

We suggest that these models, in understanding procrastination more literally as an outcome of discounting conflict, move further away from fealty to the phenomenon rather than closer to it. Ainslie (2011a) agrees, arguing in particular that re-framing choices so as to bundle streams of future rewards, the importance of which he intends as the primary insight of picoeconomics, ceases to be a natural prediction when the hypothesis of a single discount function is abandoned. We suggest that the lesson goes deeper: the best approach to modeling the distinctively economic aspects of the dynamics Ainslie identifies may be not involve forcing them into the procrustean bed of formal discounting. But of course an economic model must be formal.

3. Procrastination and choice frames

In this section we proceed as follows. We take Ainslie to have provided a psychiatrically expert description of the phenomena of procrastination and related self-management problems. However, emphasizing the fact that picoeconomic 'discounting' and economic discounting are fundamentally different concepts, with divergent histories and commitments, we drop the $\beta - \delta$ modelers' attempt to formally translate picoeconomics with minimal semantic adjustment. Like the exogenous temptation modelers, we will begin our economic reconstruction, as it were, from scratch. However, unlike the developers of that approach, we will aim to capture all of the central phenomena emphasized by Ainslie, in particular irreducible behavioural ambivalence and potential for open-ended re-framing that should not be expected to necessarily resolve, even given rationality on the part of the agent, due to informational uncertainties. Like Ainslie, but unlike exogenous temptation modelers, we will aim for generality, that is, a model of all procrastination-like phenomena rather than only a limited sub-variety of addiction.

In criticizing models of intertemporal preference ambivalence based on discounting, we should henceforth be understood as referring to 'discounting' in the economist's strict sense. Our attitude toward Ainslie's use of a different concept of 'discounting', as in Figure 1, is that this is a useful metaphor that should not restrict economists' use of alternative concepts from their own toolbox.

The most important way in which the model we will develop departs from the existing economic literature concerns the representation of the payoffs between which the agent chooses. Discounting models assume too casually that the rewards that accrue to effort on a task and to procrastination, respectively, are comparable unit for unit. We also abandon the assumption that these rewards are mutually exclusive. It is because they are not that people express often irresolvable ambivalence in their intertemporal choice behaviour, just as Ainslie argues.

In the standard economic characterization of procrastination, the phenomenon is described as resulting from a temporary preference for an SSR over an LLR. The fact that this formulation involves a preference relation implies that the two arguments on either side of the relation can be directly compared. The picoeconomic perspective casts doubt on this. Consider the example of a nicotine addict. The standard characterization implies that when procrastinating (i.e. smoking another cigarette despite planning to quit 'soon') the agent chooses between the utility of smoking the cigarette (the small, short term payoff) and that which would be gained by stopping smoking altogether (the large, long term payoff). It is doubtful that these payoffs are in fact directly comparable. In choosing not to have the cigarette, the agent does not necessarily, nor indeed in any likelihood, receive the large payoff of improved health prospects. This will be obtained only if she chooses to invest significant effort on many subsequent occasions when she experiences nicotine cravings. The duration of this time frame and the proportion of times that the agent must choose not smoking over smoking are set by agent-specific factors about which the chooser is likely to be highly uncertain, at least until she has become a veteran of several sustained but ultimately failed attempts. The cost that corresponds with the agent stopping smoking is not, as implied by the standard characterization of procrastination, the cost of abstinence over one period; rather, it is the cost of repeated choices to abstain. If the agent chooses to pay these costs, her choice is that of a long-run policy, which is formally consistent with any particular choice over the comparatively trivial stakes applying to the choice of smoking or not smoking the next cigarette. Ainslie (2011b, and elsewhere) emphasizes that if she does not link these choices, then smoking the next cigarette is a dominant strategy. This is of course true, and important; but it should not obscure the point that a *formal* model should not treat the two sides of the preference relation as drawn from the same choice set.

A second barrier to direct formal comparability of the rewards to effort and leisure involves the varying time frames over which rewards are accounted. Consider again the example of smoking. The reward yield from one cigarette is very short lived while the rewards from quitting are experienced over a lifetime. When only a short time period is considered, the utility from smoking, at least for the addict, exceeds the utility gained from not smoking over the same period. This serves to remove one possible solution to the incomparability of payoffs, namely looking at the average reward per period. Doing this would require using either a long period length over which the reward from marginal smoking is negligible, or a short period over which the reward from quitting is negligible. It takes about five minutes to smoke one cigarette and the benefits could accrue for about twenty five minutes afterwards. Thus the duration of reward from marginal smoking is about thirty minutes. Contrast this with the period of reward delivery from quitting. The benefits begin with decreased blood pressure about twenty four hours after the last cigarette, and then accrue for at least ten years through decreased risk of heart disease. We cannot divide the rewards from quitting by an appropriate factor to arrive at the reward per half hour of quitting, as these rewards are not meaningful over such a short time frame. Neither can we multiply the reward to marginal smoking to arrive at a time period of ten years as the benefits to marginal smoking are meaningless over a time period this long. Thus the rewards to effort and procrastination are incomparable because of the time frames over which these rewards become meaningful. Again, this is not a criticism of Ainslie; the incomparability of the alternative rewards in simple accounting terms is part of the explanation for both the efficacy of bundling, and for the persistent ambivalence that justifies the active cognitive involvement in such framing. The problem, again, arises for literal modeling of procrastination as expressing a preference for an SSR over an LLR.

Following our general expository strategy, let us check these intuitions against the case of the common-and-garden procrastinator, the student choosing between studying and hanging out on Facebook. According to the standard characterization she faces a choice between the utility from a few extra hours of social networking (the SSR) and the utility that would accrue from the completion of her degree (the LLR). This is obviously wrong: in all likelihood the student can procrastinate up to a point, both now and in the future, and also graduate. Nor does studying now ensure that the student will receive the intended reward from the effort; rather, she must make similar choices repeatedly for several years. The student, like the smoker, is involved in two choices, on different time scales, which may be psychologically linked in the way that Ainslie describes, but which are *formally* separate. On the one hand, she chooses whether to study or network now. On the other hand, she chooses between a policy of studying at least enough to be confident of success and a policy of trying to maximize her leisure, with a risk that she will fall short of the goal that is the point of studying at all.

One advantage of developing the basic model by reference to everyday procrastination on a short-run task, rather than addiction, is that the former permits a more intuitive idealization of a 'point of no return'. Cognitive and emotional management of the implications of a proposition such as "if I spend one more hour on Facebook now I'll probably fail my deadline" is arguably much less challenging than in the case of the formally similar young addict's knowledge that "if I don't quit within five years I'll do long-run damage to my life prospects." In our model we will treat this as a matter for point estimation, recognizing that what must be estimated in real cases is an interval. The idealization gives us one decision point at which the agent faces a decision between a mutually exclusive SSR and LLR. At any points temporally prior to this, she is able to choose another SSR that only slightly reduces the probability of obtaining the LLR.

Following the logic of standard economics, we will regard the agent's preference for maximum utility as a matter of definition. Given that procrastination has an immediate positive payoff and effort has an eventual positive payoff, the agent's utility is increasing in both of these choices. As the payoffs from effort and leisure are not generally mutually exclusive and both are positive, the agent maximizes her utility by procrastinating up to some point that is a function of her aversion to the risk of foregoing the LLR. The agent thus maximizes utility with a certain mix of effort and procrastination. The location of this mix depends on the ability of the agent and the difficulty of the project at hand. A more competent agent will be able to choose to procrastinate more than a less competent one. Similarly, a difficult project will require more effort to complete than an easy one.

We do not model a switch from effort to leisure or, for that matter from leisure to effort, as involving a change of preference between SSRs and the LLR; the agent consistently prefers an optimal mix of leisure and effort. Failure to actually optimise is diagnosed as estimation error with respect to the point at which the risk of foregoing the LLR becomes too high relative to her tolerance for risk. The agent can make too liberal an estimate of this point, leading to super-optimal procrastination (Tice & Baumeister 1997); alternatively, she might make too conservative an estimate and procrastinate less than she should.

The idea of sub-optimal procrastination may seem counter-intuitive. However, this goes to the heart of the irresolvable ambivalence that Ainslie identifies and that the standard characterization of procrastination sweeps under the rug. Normal people *want* to procrastinate, and this does not necessarily reflect either hyperbolic or $\beta - \delta$ discounting. A person knows that she must make policy choices with respect to sustained projects. If these choices were all made ex ante, then a well-informed agent would execute all of her projects briskly as soon as her policy preferences were settled, and happily consume SSRs after her investments in LLRs were securely lodged. In that framework, welfare-destroying procrastination could only possibly reflect ex ante ignorance of costs and benefits, as in rational addiction models. However, this idealization loses all touch with any real person's conception of optimality. A well-off person must rely on exogenous influences to play a major role in providing and scheduling her opportunities for policy choices. As a first approximation, she then makes choices over harvesting of SSRs within the gaps that her policy choices leave open.

This is only a first approximation to an adequate conception of the situation because it presents policy choices and whimsical choices as being independent of one another to an unrealistic degree. This is the opposite mistake to that incorporated in models based on (economic) discounting, which staple policy and whimsical choices too tightly to one another, putting SSRs and LLRs into general direct competition. A well-governed person moves back and forth between two frames of choice, recognizing that they dynamically interact *up to a point*. On the one hand she chooses and schedules projects on longer time scales that she hopes will deliver various LLRs. On the other hand she makes whimsical choices that primarily deliver – when all goes well – a diet of SSRs. These choice sequences are ideally not in continuous tension with one another, but of course choices in one sequence frequently risk undermining choices in the other sequence. Due to the non-comparability of most LLRs with most SSRs, rationality mandates no meta-policy about how careful she should be about guarding her policy choices against accidental interference from her whimsical choices. A very risk-loving agent will tolerate the sacrifice of some LLRs for the sake of a more spontaneous life, while a very risk-averse agent may expect to suffer fewer plan-destroying failures but have less fun.

Finally, agents will make self-management errors as they go along, due to common-and-garden uncertainty about contingencies, about the demands of projects, and about their own capacities. Problematic procrastination, including addiction, is the anxiety that results from the combination of uncertainty and the absence of an ideal meta-policy for balancing policy choices and whimsical ones. Of course the addict is an extreme case, who recurrently abandons policy choices, perhaps to the point where she despairs of being able to make any at all. Fortunately, only a minority of addicts – fewer than 30% in non-clinical populations (Heyman 2009) – remain mired in this condition. Even people who have squandered most of their capital resources are resilient policy adopters.

Normal, and not just pathological, switching between policy and whimsical choice frames will appear as preference reversal if the duality of the frames, and the subtleties of their interaction, are squeezed out of models. And insistence on ranking all preferences in a single monotonic field then loads all of the deep anxiety inherent in activity scheduling onto the discount function. Thus the tension between psychological and economic models of choice must emerge somewhere, though analysts can play endless shell games with it. The modeling preference expressed here is to give up on the idea of modeling a normal person as a single consistent economic agent over the scale of her whole life, in order that we can build good economic models of choices made over more limited ranges of options among which agents *do* intend and try – and very often succeed – in exercising competent consistency management.

We emphasize that we do *not* favour the general abandonment of independence or completeness axioms in economic models; where these axioms should be relaxed in particular cases, there should always be a sound case-specific motivation. Our philosophical perspective is instead that no isolated economic model can reasonably aspire to describe all of an individual biological person's behaviour at all scales of analysis simultaneously. Biological individual people are merely correlated with complex sets of temporally indexed economic agents (Ross 2005), and any given economic model of an individual's incentive-response profile should thus be expected to be accurate only within empirically determined boundary conditions. Very short-run models of individual response might be based entirely on preference orderings relevant to whimsical choices with policy choice constraints held fixed. On the other hand, models of aggregate preference, which we regard as the core business of the economist as a social scientist, will generally treat variations in whimsical choices by individuals as noise. Then sometimes, for example in modeling addicts and others who suffer from disruptive procrastination anxieties, we must introduce elements of risk and uncertainty that arise from the potential for whimsical and policy choices to interfere with one another.

The model we produce in Section 4 below should be understood in these terms. It depicts the agent as statically rather than dynamically optimizing within the interval in which procrastination anxiety arises. The intuition here is that the agent does not acquire significant new information within that interval. Her policy choices and preferences are implicitly treated as fixed. The agent procrastinates until she sets to work on a project, and then she discovers ex post whether she has left herself enough time to complete it. She procrastinates in the first place because dynamic optimization on the longer-run scale of policy preferences precludes zero procrastination, i.e., a life plan in which every desired project is carried out one after the other without breaks between them. The following intuition is relevant as an example: it is not irrational for even very busy people to punctuate bursts of work with drinking and other stereotypically 'wasteful' activity; yet all drinking is procrastination[3]. Since dynamic optimization *does* characterize the scale of policy choice, we should expect learning to occur on this scale, which will inform static optimization on shorter scales. Thus we extend the model to allow that agents become better-informed estimators of optimal procrastination-effort mixes as they recurrently encounter choices with similar elements and structures. In the picoeconomic framework this emerges as reward bundling.

4. A model of procrastination

Based on her policy preferences, an agent i has chosen to complete a project ψ that is estimated to require a quantity of effort E. She considers how to allocate E within a time interval X from the point of choice that is composed of two sub-intervals. During the second sub-interval $\bar{X}$ she will expend E. During the first sub-interval $X - \bar{X}$ she will make whimsical choices. We will refer to $X - \bar{X}$ as involving consumption of 'leisure', L, but this should not be interpreted as excluding productive activity devoted to something other than the project produced by E. The agent chooses a bid value $B\left(\hat{V}_E\right)$ to maximize

$$EPO = \left(\hat{P}_C | AB\left(\hat{V}_E\right)\right)\hat{V}_E - B\left(\hat{V}_E\right) + V_P \tag{4}$$

where

$\hat{V}_E$ = Estimated value of effort payoff

$B\left(\hat{V}_E\right)$ = Portion of time budget allocated to effort $(\bar{X})$

[3] Professional wine tasters are exceptions. Philosophers will think of others.

$A = A(T)$

The function $A = A(T)$ weights the bid value in the agent's estimation that ψ will be completed on time. This weighting function captures an increasing efficiency effect of concentration as the agent becomes more certain that she is best off expending E and so experiences less distraction from the possibility of consuming L. $A(T)$ is an increasing function of time with an initial value of 1 at the point of choice (the beginning of X). As X progresses, the value of $A(T)$ increases non-linearly, reflecting the fact that the agent becomes more productive as the deadline approaches and anxiety due to the possibility of further procrastination vanishes. The function is characterised by the following properties:

$A'(T) > 0$

$A''(T) > 0$ up to an inflection point thereafter $A''(T) < 0$.

The agent's rate of productivity increases up to an inflection point (**F** in Figure 2) and thereafter declines, reflecting decreasing marginal benefits of concentration.

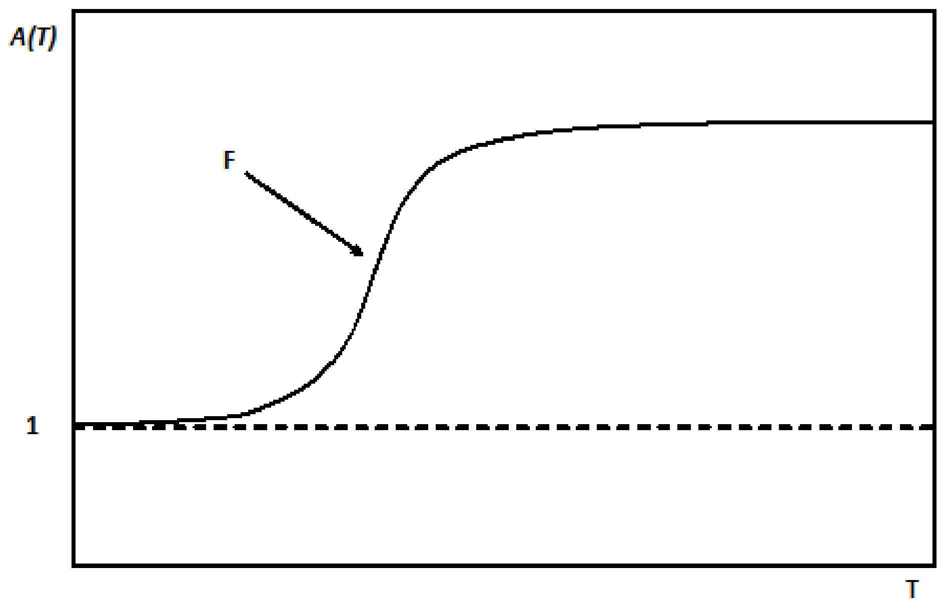

Figure 2: The A(T) function

$\left(\hat{P}_C | AB\left(\hat{V}_E\right)\right)$ =Estimated probability of completing ψ conditional on $AB\left(\hat{V}_E\right)$

$V_P = \left(\bar{T} - B\left(\hat{V}_E\right)\right) MU_P$

with

$\bar{T}$ = Time budget (X)

MU_P = Constant marginal utility of procrastination

$$B\left(\hat{V}_E\right) = b\left(\hat{V}_E\right) + \gamma$$

Here $b(\hat{V}_E)$ is a fixed project-type specific component and γ is a variable component that the agent adjusts to refine $B\left(\hat{V}_E\right)$ when a project, $\psi_{n\neq 1}$, belonging to the time-ordered set of informationally relevant projects, $\Psi : \psi_1, ..., \psi_n$, is attempted. Considering the progression from ψ_1 to ψ_n the agent's estimation accuracy would be expected to increase due to refinements made through γ. $b(\hat{V}_E)$ is thus the agent's estimate of how much of X to assign to E when attempting $\psi_1 \in \Psi$.

The form of $A(T)$ implies that $\left(\hat{P}_C|AB\left(\hat{V}_E\right)\right)$ accounts for both the quantity and quality of E employed by the agent. If the agent were to employ E from the begging of X (for a given project, ψ) then $\left(\hat{P}_C|AB\left(\hat{V}_E\right)\right) = \left(\hat{P}_C|B\left(\hat{V}_E\right)\right)$. However, as the agent becomes more productive throughout, an equal estimated probability of completion occurs further into X at the point where the positive effect of an increase in A outweighs the fall in $B\left(\hat{V}_E\right)$ implied by the agent delaying the earliest expenditure of E. The agent is thus better off procrastinating for some positive $X - \bar{X}$ because, due to $A(T)$, the agent is able to devote a lower proportion of X to E which implies a higher payoff through an increase in $\left(\bar{T} - B\left(\hat{V}_E\right)\right)$.

The formulation of (4) reflects that more procrastination (implying a lower $B\left(\hat{V}_E\right)$) increases the agent's expected payoff through its effect on V_P throughout X. This effect is compounded by the fact that a lower $B\left(\hat{V}_E\right)$ implies ψ is started later in X and thus, by the logic above, a lower absolute value of $B\left(\hat{V}_E\right)$ is required in order to complete ψ due to increased productivity. This compounding effect will only hold as long as the positive productivity effect outweighs the negative bid value effect. The cost of L is, for a given value of $A(T)$, the fall in $\left(\hat{P}_C|AB\left(\hat{V}_E\right)\right)$ that results from a lower value of $B\left(\hat{V}_E\right)$. As mentioned above, this cost will initially be compensated for by $A(T)$.

We make the assumption that the agent's utility function increases with the expected payoff function (4) and thus that the agent maximises utility by finding the optimal allocation of time between $\bar{X}$ and $X - \bar{X}$.

As the agent has control only over $B\left(\hat{V}_E\right)$, the condition for a maximum is that

$$\frac{\partial\left[\left(\hat{P}_C|AB\left(\hat{V}_E\right)\right)\hat{V}_E - B\left(\hat{V}_E\right) + V_P\right]}{\partial B\left(\hat{V}_E\right)} = 0$$

Let $\left(\hat{P}_C|AB\left(\hat{V}_E\right)\right)$ be expressed by

$$\left(\hat{P}_C|AB\left(\hat{V}_E\right)\right) = \aleph AB\left(\hat{V}_E\right)$$

where $\aleph$ is a function that maps weighted values of $B\left(\hat{V}_E\right)$ onto expected probabilities such that $0 \leq \aleph AB\left(\hat{V}_E\right) \leq 1$. When $\left(\left(\bar{T} - B\left(\hat{V}_E\right)\right) MU_P\right)$ is substituted for V_P, then we derive

$$\aleph A\hat{V}_E - MU_P - 1 = 0$$

and thus show that the maximum of the expected payoff depends on the value of E and the value of L. This will hold for all $B\left(\hat{V}_E\right) < \bar{T}$.

Let $B\left(\hat{V}_E\right)^*$ be the value of $B\left(\hat{V}_E\right)$ that maximises (4). When the realised utility from ψ_{n-1} was smaller than expected utility[4] ($AU_{n-1} < EU_{n-1}$) the agent learns information relevant to ψ_n. The agent considers two broad hypotheses as to why $AU_{n-1} < EU_{n-1}$. The first hypothesis is that there was an error in estimation. This is treated as a single hypothesis as the agent is assumed to know the direction of this error. The second hypothesis is that the inequality was due to an independent exogenous factor with the set of all possible exogenous factors represented by Θ. This second hypothesis is a compound of all the other factors that could lead to failure to maximise. The agent learns according to which hypothesis better predicts the observed outcome. Formally, the Bayesian conditionals are:

$$h_1 : P\left(AU_{n-1} < EU_{n-1} | B\left(\hat{V}_E\right)_{n-1} \neq B\left(\hat{V}_E\right)^*_{n-1}\right)$$

$$h_2 : P\left(AU_{n-1} < EU_{n-1} | \Theta\right)$$

If h_1 is found to be a better explanation for failure to maximize than h_2, the learning process will proceed as outlined below. If h_2 is the preferred hypothesis, then the agent assumes that $B\left(\hat{V}_E\right)_{n-1} = B\left(\hat{V}_E\right)^*_{n-1}$ and thus no adjustment occurs. The assumption that new information is only available at the end of X excludes updating of estimation during X.

The learning process is mechanically simple. If $B\left(\hat{V}_E\right)_{n-1} \neq B\left(\hat{V}_E\right)^*_{n-1}$ then $B\left(\hat{V}_E\right)_n = B\left(\hat{V}_E\right)_{n-1} + \gamma$ where $|\gamma| > 0$. By assumption, the agent recalls whether $B\left(\hat{V}_E\right)_{n-1} < B\left(\hat{V}_E\right)^*_{n-1}$ or $B\left(\hat{V}_E\right)_{n-1} > B\left(\hat{V}_E\right)^*_{n-1}$ and thus knows whether to choose $\gamma_n = \gamma^+$ so that $B\left(\hat{V}_E\right)_n > B\left(\hat{V}_E\right)_{n-1}$ or $\gamma_n = \gamma^-$. so that $B\left(\hat{V}_E\right)_n < B\left(\hat{V}_E\right)_{n-1}$. The agent does not know the size of the estimation error and thus when attempting $\psi_{n+k} \in \Psi$, $B\left(\hat{V}_E\right)_{n+k}$ will approach

[4] And thus that the actual payoff (received by the agent at the end of the period) was smaller than the expected payoff.

$B\left(\hat{V}_E\right)^*_{n+k}$ under the assumption that the agent remembers the size of γ_{n+k-1}. What is important to note here is that this learning process only applies to elements of Ψ. It is plausible that the agent could notice a tendency to over-indulge in L and thus adjust E upwards in unfamiliar situations. Another important factor here is that, under the assumption that the agent is risk averse, bid values below optimal are likely to increase faster than above optimal bid values are likely to decrease (see proof below).

Consider the true payoff function[5], dropping the indexation to Ψ for the time being:

$$PO = V_E - B(V_E) + \left(\bar{T} - B(V_E)\right) MU_P \tag{5}$$

Let $B(V_E)^*$ maximise this function and PO_* be the maximised function. Thus any values of $B(V_E)$ above or below $B(V_E)^*$ will result in a lower payoff. If $B(V_E) > B(V_E)^*$, a higher payoff could be achieved be choosing higher L while if $B(V_E) < B(V_E)^*$ a higher payoff could be gained by allocating higher E assuming that $V_E > MU_P$, as, by definition, $B(V_E) < B(V_E)^*$ results in the payoff to E being lost.

Define $B(V_E)^+$ as being infinitesimally greater than $B(V_E)^*$ which is itself infinitesimally greater than $B(V_E)^-$ Substituting $B(V_E)^+$ into (5) and simplifying gives

$$PO_+ = -B(V_E)^+ (1 + MU_P) + \bar{T}MU_P + V_E \tag{6}$$

Repeating the procedure for $B(V_E)^-$ gives

$$PO_- = -B(V_E)^- (1 + MU_P) + \bar{T}MU_P \tag{7}$$

By definition, PO_* is greater than PO_+ and PO_-. Furthermore, $PO_+ > PO_-$ for all values of V_E such that $V_E > (1 + MU_P)\left(B(V_E)^+ - B(V_E)^-\right)$. What this demonstrates is that for appropriate values of V_E, excessive consumption of L is more costly than excessive E. Furthermore, the restriction on these values is very lenient as, by definition, $B(V_E)^+ - B(V_E)^-$ is very small. Returning to the learning process, this implies that for the risk averse agent, $|\gamma^+| > |\gamma^-|$ for values of $B\left(\hat{V}_E\right)$ equidistant from $B(V_E)^*$ and thus sub-optimal bid values will increase faster than super-optimal bid values will decrease.

[5] The true payoff function describes the payoff that the agent actually receives at the end of the period.

The fact that sub-optimal bid values are more costly is important as it predicts a role for bundling. In order to prove this, it is necessary to establish that the agent, when considering a sequence, $\psi_1, ..., \psi_n$ from Ψ will select a proportionally higher value of $B\left(\hat{V}_E\right)$ when forced to precommit in ψ_1 than when each ψ is considered separately. This implies bundling as the agent is more likely to choose a higher value of $B\left(\hat{V}_E\right)$ when committing to a course of action over multiple choices in the set Ψ.

Assume that the agent faces three choices in Ψ, ψ_1, ψ_2, ψ_3. There are two possible cases with respect to bundling. In the no-bundling case, the agent sets $B\left(\hat{V}_E\right)$ at the start of each ψ_n. In the bundling case, the agent commits in ψ_1 to a value of $B\left(\hat{V}_E\right)$ that is the same in $\psi_{1,...,3}$. Furthermore, assume in both cases that the agent is equally likely to choose $B\left(\hat{V}_E\right)_1 < B\left(\hat{V}_E\right)^*$ and $B\left(\hat{V}_E\right)_1 > B\left(\hat{V}_E\right)^*$, so $\gamma_2 = \gamma^-$ with probability 0.5 and $\gamma_2 = \gamma^+$ with probability 0.5.

In the no-bundling case, the probability that $B\left(\hat{V}_E\right)_2 > B\left(\hat{V}_E\right)^*$ is greater than 0.5 as a positive revision is expected in ψ_2 due to the fact that $(0.5\gamma^+) + (0.5\gamma^-) > 0$ as $|\gamma^+| > |\gamma^-|$.

For the no-bundling case:

$$B\left(\hat{V}_E\right)_1 = b\left(\hat{V}_E\right)$$
$$B\left(\hat{V}_E\right)_2 = b\left(\hat{V}_E\right) + (0.5\gamma^+) + (0.5\gamma^-)$$
$$B\left(\hat{V}_E\right)_3 = b\left(\hat{V}_E\right) + (P_L\gamma^+) + (P_H\gamma^-)$$

Where $P_L < P_H$ as the expected revision for ψ_2 is positive thus increasing the probability of over-estimation in ψ_3.

In the no-bundling case then, the evolution of $B\left(\hat{V}_E\right)$ can be expressed as

$$B\left(\hat{V}_E\right)_{1,2,3}^{T(NB)} = b\left(\hat{V}_E\right) + \left(b\left(\hat{V}_E\right) + (0.5\gamma^+) + (0.5\gamma^-)\right) + b\left(\hat{V}_E\right) + (P_L\gamma^+) + (P_H\gamma^-)$$

where $B\left(\hat{V}_E\right)_{1,2,3}^{T(NB)}$ is the total bid value across the three elements of Ψ.

In the bundling case,

$$B\left(\hat{V}_E\right)_1 = B\left(\hat{V}_E\right)_2 = B\left(\hat{V}_E\right)_3 = b\left(\hat{V}_E\right) + (0.5\gamma^+) + (0.5\gamma^-)$$
$$B\left(\hat{V}_E\right)_{1,2,3}^{T(B)} = 3\left(b\left(\hat{V}_E\right) + (0.5\gamma^+) + (0.5\gamma^-)\right)$$

To show that the agent improves her payoff by bundling, it must be established that

$$B\left(\hat{V}_E\right)_{1,2,3}^{T(B)} > B\left(\hat{V}_E\right)_{1,2,3}^{T(NB)}$$

For this to be the case:

$$2\left(\left(0.5\gamma^+\right) + \left(0.5\gamma^-\right)\right) > \left(P_L\gamma^+\right) + \left(P_H\gamma^-\right)$$

$$\therefore \gamma^+ - P_L\gamma^+ > -\left(\gamma^- - P_H\gamma^-\right)$$

$$\therefore \gamma^+\left(1 - P_L\right) > -\left(\gamma^-\left(1 - P_H\right)\right)$$

As $P_H + P_L = 1$

$$\gamma^+ P_H > -\left(\gamma^- P_L\right)$$

But $-\left(\gamma^- P_L\right) > 0$ as $P_L > 0$ and $\gamma^- < 0$

Additionally, $\gamma^+ P_H > -\left(\gamma^- P_L\right)$ as $|\gamma^+| > |\gamma^-| > 0$ and $P_H > P_L > 0$.

Therefore,

$$B\left(\hat{V}_E\right)_{1,2,3}^{T(B)} > B\left(\hat{V}_E\right)_{1,2,3}^{T(NB)}$$

as the difference between the two is positive. This establishes that bundled bid values will be proportionately higher than unbundled bid values.

As the agent learns from experience to better estimate $B\left(\hat{V}_E\right)$, if she discovers that her initial choices over Ψ failed to optimize, she should adjust $B\left(\hat{V}_E\right)$. This is equivalent to learned bundling.

5. Conclusion

The agent described by our model chooses so as to maximize her utility. Two factors may prevent her from succeeding. One of these is uncertainty about the time she will need to accomplish projects. The other is that she lacks ex ante knowledge about which projects will come along for her, and while waiting for this information she makes whimsical choices. Sometimes these whimsical choices will undermine her policy choices over projects. Uncertainty about these conflicts can lead to sub-optimal time allocations, in which some projects are started too soon and some are started too late.

In the context of this model, we suggest that the basic source of the ambivalence described by Ainslie is that people do not have, and would not 'rationally' aim to have, complete ex ante preferences that put whimsical and policy preferences in a single ordering. A typical person is not best-off deriving all of her whimsical choices from her policy preferences, for the consequences of this manifest as obsessiveness, rigidity and self-repression. But a person who lets whimsical choices undermine policy preferences without limit is at best a drifter through life and more typically an addict. Most people are aware that the tension between whimsical and policy preferences requires management. But because rationality underdetermines the *summum bonum*, most people procrastinate and are frequently anxious about doing so.

We do not say that rationality is *silent* on the management of preference scales with incompletely commensurable trade-off margins. People encounter sequences of projects with similar structures and learn both about these structures and about their own capacities with respect to them. Application of this learning results in reward bundling. Procrastinators learn various techniques of self-control, which Ainslie has characterized with fresh descriptive richness following a long literary and philosophical tradition. By means of these techniques, most addicts eventually bring their self-destructive consumption under control. As addiction counselers emphasize, risk never entirely goes away. But this is not a special feature of addict psychology; it characterizes all people who ever procrastinate, which is to say, almost everyone.

References

[1] Ainslie, G. (1992). Picoeconomics. Cambridge: Cambridge University Press.

[2] Ainslie, G. (2001). Breakdown of Will. Cambridge: Cambridge University Press.

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

[4] Ainslie, G. (2011a). Pure hyperbolic discount curves predict 'eyes-open' self control. Theory and Decision, forthcoming.

[5] Ainslie, G. (2011b). Free will as recursive self-prediction. In G. Graham & J. Poland, eds., Addiction and Responsibility. Cambridge, MA: MIT Press.

[6] Becker, G., & Murphy, K. (1988). A theory of rational addiction. Journal of Political Economy 96: 675-700.

[7] Bénabou, R., & Tirole, J. (2004). Willpower and personal rules. Journal of Political Economy 112: 848-886.

[8] 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.

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

[10] Fudenberg, D., & Levine, J. (2006). A dual-self model of impulse control. American Economic Review, 96, 1449-1476.

[11] 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.

[12] Gul, F., & Pesendorfer, W. (2001). Temptation and self-control. Econometrica 69: 1403-1435.

[13] Gul, F., & Pesendorfer, W. (2004). Self-control, revealed preference and consumption choice. Review of Economic Dynamics 7: 243-264.

[14] Herrnstein, R. (1997). The Matching Law. Cambridge, MA: Harvard University Press.

[15] Heyman, G. (2009). Addiction: A Disorder of Choice. Cambridge, MA: Harvard University Press.

[16] Laibson, D. (1997). Golden eggs and hyperbolic discounting. Quarterly Journal of Economics 112: 443-477.

[17] Mazur, J. (1987). An adjusting procedure for studying delayed reinforcement. In M. Commons, J. Mazur, J. Nevin, and H. 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.

[18] McClure, S., Laibson, D., Loewenstein, G., and Cohen, J. (2004). Separate neural systems value immediate and delayed monetary rewards. Science 306: 503-507.

[19] O'Donoghue, T., and Rabin, M. (1999). Incentives for procrastinators. Quarterly Journal of Economics 114: 769-816.

[20] O'Donoghue, T., and Rabin, M. (2001). Choice and procrastination. Quarterly Journal of Economics 116: 121-160.

[21] Phelps, E., & Pollack, R. (1968). On second-best national saving and game equilibrium growth. Review of Economic Studies 35: 201-208.

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

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

[24] Ross, D. (2011). Addictive, impulsive and other counter-normative consumption. In V. Wells & G. Foxall, eds., New Developments in Consumer Behaviour. London: Edward Elgar.

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

[26] Tice, D., & Baumeister, R. (1997). Longitudinal study of procrastination, performance, stress and health: The costs and benefits of dawdling. Psychological Science 8: 454-458.