Procrastination, partial naivete, and behavioral welfare analysis.

Geir B. Asheim*

August 20, 2008

Abstract

This paper has a dual purpose. First, I present a new modeling of partial naivete and apply this to procrastination. The decision maker is partially naive by perceiving that his current preferences may persist in the future. The behavioral implications of such partial naivete differ from those of related literature. Second, I suggest a general principle for welfare analysis in multi-self settings through a new application of Pareto-dominance, which reduces to the usual criterion for intertemporal choice if preferences are time consistent. In the case of procrastination, it leads to a clear welfare conclusion: Being partially naive reduces welfare.

Keywords and Phrases: Procrastination, partial naivete, time inconsistency, game theory, behavioral welfare economics.

JEL Classification Numbers: C70, D11, D60, D74, D91, E21

---

*Department of Economics, University of Oslo, P.O. Box 1095 Blindern, 0317 Oslo, Norway (e-mail: `g.b.asheim@econ.uio.no`).

This paper is part of the research activities at the ESOP centre at the Department of Economics, University of Oslo. I thank Ted O'Donoghue for helpful discussions, Zafer Akin, Kaushik Basu, Eddie Dekel, Martin Dufwenberg, Tore Ellingsen, Aanund Hylland, Botond Kőszegi, Tore Nilssen, Karine Nyborg, Andrés Perea and seminar participants in Atlanta (ESA 2006), Evanston (GAMES 2008), Maastricht, Oslo, Paris and Wien (ESEM 2007) for useful comments, Cornell University for hospitality, and the Research Council of Norway for financial support, in particular, through ESOP.

People procrastinate a task with immediate cost and future benefits, if they end up doing the task later than they *thought* they would and later than they *wish* they would.[1] Such procrastination is wide-spread, and has been made subject to economic analysis during the last couple of decades (see, e.g., Akerlof, 1991, and contributions that cite him).

The existence of this phenomenon leads to two sets of questions.

- How to formulate a *positive* theory of procrastination: Why do the future selves of a decision maker deviate from the behavior predicted by earlier selves of the decision maker?
- How to formulate a *normative* theory of procrastination: In what sense is the delayed execution of the task detrimental to the interests of the decision maker?

Procrastination is a problem of self-control. Two different lines of economic literature address self-control problems.

One line of literature studies *multi-self models*, which consider the *inter*temporal conflict between the decision maker's temporal selves. It is assumed that the decision maker is endowed with "present-biased preferences",[2] which gives rise to time-inconsistent preferences, meaning that plans that are currently optimal will not necessarily remain so in the future. The most prominent contribution on procrastination within this line of literature is O'Donoghue and Rabin (2001) (see also, e.g., Brocas and Carrillo, 2001, and Fischer, 1999).

Another line of literature studies *dual-self models*, which consider the *intra*temporal conflict between a long-run patient self and a sequence of short-run impulsive

[1]This is not meant to be an all-encompassing definition of the phenomenon of procrastination, but a description of the kind of procrastination that is the subject of the present paper.

[2]This phenomenon seems to be supported by experimental evidence, as reported by Loewenstein and Prelec (1992) and Frederick, Loewenstein and O'Donoghue (2003), although contested by Rubinstein (2003).

selves. The most prominent such contribution is Fudenberg and Levine (2006), with Thaler and Schefrin (1981) as an influential precursor. In this classification of the self-control problem literature, Gul and Pesendorfer (2001, 2004), Dekel, Lipman and Rustichini (2008) and other similar papers studying preferences over choice sets, with a desire to limit the available alternatives, can be considered as reduced forms as dual-self models, leading to similar positive and normative conclusions (cf. Fudenberg and Levine, 2006, Section VI).

To highlight that neither the analysis of O'Donoghue and Rabin (2001) nor the literature on dual-self models leads to a fully satisfactory theory of procrastination, both from a positive and normative perspective, I consider in this paper a simple stationary and deterministic environment where a decision maker has one task to perform, and his problem is to decide when to do the task, if at all.

In dual-self models, the utility of the long-run self is an obvious welfare criterion (Fudenberg and Levine, 2006, p. 1471). Moreover, it is easy to choose values for the task's immediate cost and future benefits so that delayed execution of the task is detrimental to the interests of the long-run self. In a deterministic setting like the one I will consider, dual-self models predict that the decision maker either acts at once or not all. As Fudenberg and Levine (2006, p. 1464) note, their dual-self model predicts delayed and stochastic execution only if there is structural uncertainty about the environment. Furthermore, in expectations, it can explain why the task will in fact be performed later than the decision maker originally thought, only if the decision maker has biased prior beliefs.

Procrastination, in the sense of the opening paragraph, is characterized by strategic uncertainty: The decision maker is uncertain about when his future selves will perform the task, even in a deterministic setting. Since dual-self models cannot easily handle such strategic uncertainty, it seems appropriate to consider multi-self models with present-biased and, thus, time-inconsistent preferences, as O'Donoghue and Rabin (2001) do.

By limiting their analysis to pure strategies, O'Donoghue and Rabin (2001) do not address procrastination with strategic uncertainty. Still, their analysis can easily be amended to allow for the modeling of strategic uncertainty, by considering multi-self equilibria in stationary mixed strategies (O'Donoghue and Rabin, 2001, footnote 14). Moreover, one can capture that the task will be performed later than originally expected, by introducing that the decision maker is *partial naive* in the sense of not being fully aware of the self-control problems that his present biased preferences lead to. However, due to the manner in which O'Donoghue and Rabin (2001) model partial naivete, any vestige of such partial naivete leads to a discontinuous change in behavior by resulting in perpetual delay (O'Donoghue and Rabin, 2001, footnote 20). This motivates *the first main contribution of this paper*, which goes beyond the specific analysis of procrastination: How to define partial naivete so that behavioral change is continuous in the present paper's model of procrastination.

My alternative way of modeling partial naivete is introduced in Section 1, defined formally in a general setting in Section 2, and applied to the specific model of procrastination in Section 3. This definition, unlike the one proposed by O'Donoghue and Rabin (2001), does not depend on a particular specification of present-biased preferences (like the (β, δ)–preferences described in Section 1). Furthermore, formally the notion can be analyzed by the concept of subgame-perfectness; hence, no underlying game-theoretic issues must be addressed before putting it to use.

This paper considers only stationary mixed strategies in its analysis of procrastination. This does not entail that I rule out that people overcome commitment problems by means of promises of self-reward and threats of self-punishment; in fact, it is an asset of the multi-self approach that it allows for the modeling of such self-enforcing schemes through non-stationary strategies. Rather, my point is that even an analysis of procrastination by means stationary mixed strategies requires, as I have argued above, a new positive theory of partial naivete and, as I will turn to next, a new normative theory by which to evaluate the welfare effects of such naivete.

As I point out in Section 1 and discuss in more detail in Section 4, the multi-self literature has not proposed any welfare criterion by which stationary mixed strategies can be ranked. This leads to *the second main contribution of this paper*, which also goes beyond the specific analysis of procrastination: How to formulate a plausible and general principle for welfare analysis which allows the ranking of stationary mixed strategies in the present paper's model of procrastination.

This alternative welfare criterion is introduced in Section 1, defined formally in in general multi-self settings in Section 4, and used to evaluate the welfare effects of procrastination in Section 5. It has the attractive feature of not being paternalistic; rather, it is based on the actual preferences of the decision maker when comparing costs and benefits accruing at different times. Moreover, it does not depend on a particular specification of present-biased preferences. Finally, it reduces to the usual criterion for intertemporal choice if preferences are time consistent.

As noted in Section 6, the simple model of procrastination is but a vehicle for demonstrating the need for and the consequences of (i) a new notion of partial naivete and (ii) a new principle for welfare analysis in multi-self models, which are the two main contributions of this paper. All proofs are relegated to an appendix.

1 Why procrastinate and what are the welfare effects?

Why do people procrastinate? Procrastination can be explained by endowing the decision maker with the following two properties:

(1) The decision maker has present-biased preferences. A common formulation (originating from Phelps and Pollak, 1968, and Elster, 1979, and employed by Laibson, 1994, 1997, and others) is to assume that the decision maker has preferences that, in addition to time-consistent discounting, give extra weight to current well-being over future well-being (so-called (β, δ)–preferences).

(2) The decision maker is not fully aware of the self-control problems that such

present-biased preferences lead to; instead he is partially naive about future preferences.

The present-biased preferences provide an incentive for the decision maker to postpone the task to the next period, while the deficient awareness of the self-control problems leads him to believe falsely that the task will actually be performed then. As shown by O'Donoghue and Rabin (2001), these two elements are sufficient to model behavior whereby the decision maker postpones the task period by period.

While obviously being a very important and influential contribution, the analysis presented by O'Donoghue and Rabin (2001) has features that might make it worthwhile to consider alternative ways to model procrastination, while keeping the key ingredients (1) and (2): (β, δ)–preferences and partial naivete.

The alternative modeling choices that I propose in the present paper can be illustrated through the following example. Consider a decision maker who performs a task at one of the stages 0, 1, 2, ..., or not at all. Performing the task leads to an immediate cost equal to 1, and enables the decision maker to reap benefits equal to 5 at the next stage. Assume a constant per-period discount factor, $\delta = 4/5$. Moreover, assume that future costs and benefits are discounted by an additional factor, $\beta = 1/2$, leading to present-biased preferences. The payoffs of the decision maker as a function of the number of stages he postpones the task are given in Table 1. Hence, the decision maker prefers to postpone the task one and only one period.

O'Donoghue and Rabin (2001) study pure strategies, leading to perhaps unconvincing patterns of behavior obtained through backward induction (cf. Fudenberg and Levine, 2006, p. 1464). E.g., in the example above, a sophisticated decision maker (i.e., who is fully aware of his future self-control problems) might decide to do the task now, because he believes that otherwise he will postpone the task for exactly 2 periods. Instead, I assume that the decision maker uses a stationary strategy, where the same mixed action is chosen if the task has not yet been performed. E.g., a sophisticated decision maker will do the task with probability 1/2, since then

Table 1: Payoffs of the decision maker as a function of postponement

Postponement	**Payoff**
0	$-1 + \frac{1}{2}\frac{4}{5}5 = 1$
1	$\frac{1}{2}\frac{4}{5}\big(-1 + \frac{4}{5}5\big) = \frac{6}{5}$
2	$\frac{1}{2}\big(\frac{4}{5}\big)^2\big(-1 + \frac{4}{5}5\big) = \frac{24}{25}$
...	...
n	$\frac{1}{2}\big(\frac{4}{5}\big)^n\big(-1 + \frac{4}{5}5\big) = \frac{3}{2}\big(\frac{4}{5}\big)^n$
...	...
∞	0

the payoff (= 1) he can ensure himself by doing the task now is equal to the expected payoff of doing the task at stages 1, 2, 3, ... with probability 1/2, 1/4, 1/8, ...

O'Donoghue and Rabin (2001) assume that the partially naive decision maker is consistently mistaken about future preferences, still quite sophisticated when constructing his decision rule. E.g., if the decision maker believes that his β at future stages will not be 1/2, but rather exceed 5/8 so that

$$-1 + \beta\tfrac{4}{5}5 \;>\; \beta\tfrac{4}{5}\big(-1 + \tfrac{4}{5}5\big)\,,$$

then he believes that he will do the task for sure at the next stage, entailing that he strictly prefers *not* to do the task now. However, at each new stage he discovers that his preferences are as present-biased as before, and that the task remains undone.

Here, I model partial naivete as follows: The decision maker *perceives that his present preferences will persist with positive probability* at the next stage, entailing that his future preferences will *not* be present-biased and effectively enabling him to commit to his present decision rule. E.g., if the decision maker believes with at least probability 1/2 that his present preferences will persist at the next stage, then it follows from the results of the present paper that in the unique stationary equilibrium the decision maker chooses to postpone the task for sure. Still, his understanding of the decision problem is not contradicted by the flow of events, since—as long as he is not fully naive—he perceives that there is at each stage a positive probability that the present preferences will *not* persist. Although related to the two-period preferences modeled by Eliaz and Spiegler (2006, Section 3),[3] this definition is original as a general suggestion for how to capture partial naivete.

What are the welfare consequences? O'Donoghue and Rabin (2001) along with most of the behavioral economics literature on time inconsistency employ two different welfare criteria: (i) "Long-run utility" meaning that one sets $\beta = 1$ for welfare analysis. (ii) The Pareto-criterion is used to evaluate two alternatives by comparing payoffs along the implemented paths at *all stages*. These welfare criteria are not uncontroversial—see, e.g., Bernheim and Rangel (2007, Section 2C), Gul and Pesendorfer (2005, Section 6.4) and O'Donoghue and Rabin (2001, footnote 21)—and may not be considered fully compelling. One can claim that "long-run utility" does not respect the preferences of the decision maker, and by applying the Pareto-criterion to contrast payoffs along the implemented paths one is lead to make comparisons across different histories. Also, as I will argue in Section 4, neither can rank the stationary strategies used by the decision maker in my model of procrastination.

Here I suggest to avoid these issues by using the Pareto-criterion to evaluate two alternatives by comparing payoffs at all *decision nodes*. By applying the Pareto-criterion in this fashion to my model of procrastination, stationary strategies—and thus, stategies for different levels of partial naivete—turn out to be Pareto-ranked. Hence, the analysis leads to a clear welfare conclusion: Being partially naive decreases welfare. Although formally related to dominance relations used elsewhere in game theory (e.g., in the literature on renegotiation-proofness in repeated games), this use of the Pareto-criterion is original for the purpose of welfare analysis.

[3]Indeed, Heidhues and Kőszegi (2008a, Section 6.2) refer to the present definition as "Eliaz-Spiegler-Asheim Partial Naivete". A similar notion of partial naivete is also proposed by Akin (2007, Section 5.2) using (β, δ)–preferences, although the behavioral implications are not explored.

2 A general model of partial naivete

Consider a T-stage decision problem where T can be finite or infinite. The set of *histories* is defined inductively as follows: The set of histories at the beginning of the first stage 0 is $H_0 = \{h_0\}$. Denote by H_t the set of histories at the beginning of stage t. At $h \in H_t$, the decision maker's action set is A^h. Define the set of histories at the beginning of stage $t+1$ as follows:

$$H_{t+1} = \{(h, a) \mid h \in H_t \text{ and } a \in A^h\}.$$

This concludes the induction. Let H denote the set of *decision nodes*. If $T < \infty$, then $H = \bigcup_{t=0}^{T-1} H_t$; otherwise, $H = \bigcup_{t=0}^{\infty} H_t$. Assume that, for all $h \in H$, A^h is non-empty and finite. A trivial decision is made at h if A^h is a singleton. Refer to $Z := H_T$ as the set of *outcomes*. Denote by $t(h)$ the stage which starts at $h \in H$. Denote by H^h the set of decision nodes equal to or following $h \in H$:

$$H^h := \{h' = (h_0, a_0, \dots, a_{t(h')-1}) \in H \mid t(h') \geq t(h) \text{ and } h = (h_0, a_0, \dots, a_{t(h)-1})\},$$

and denote by Z^h the set of outcomes following $h \in H$:

$$Z^h := \{z = (h_0, a_0, a_1, \dots) \in Z \mid h = (h_0, a_0, \dots, a_{t(h)-1})\},$$

The preferences of the decision maker at $h \in H$ are represented by a Bernoulli utility function $u^h : Z^h \to \mathbb{R}$ which assigns *payoff* to every outcome following h. The preferences are *time consistent* if, for all $h \in H$ and $a \in A^h$, $u^{(h,a)}$ is a positive affine transformation of u^h restricted to $Z^{(h,a)}$.

For every $h \in H$ and $h' \in H^h$, the decision maker perceives at h that his preferences at h' will equal $\tilde{u}^{h|h'}$. His system of perceived preferences over H^h are defined inductively as follows: (1) $\tilde{u}^{h|h} = u^h$, and (2) for every $h' \in H^h$ and $a \in A^{h'}$, $\tilde{u}^{h|(h',a)} = \tilde{u}^{h|h'}|_{Z^{(h',a)}}$ with probability p and $\tilde{u}^{h|(h',a)} = u^{(h',a)}$ with probability $1-p$. Refer to $p \in [0, 1]$ to as the *perceived preference persistence*, reflecting the probability with which the decision maker perceives at h that his preferences at stage t ($\geq t(h)$)

will persist at $t+1$. Note that since $\tilde{u}^{h|h} = u^h$ for every $h \in H$, there is no *actual* preference persistence, except if preferences are time consistent.

A *decision rule* at $h \in H$, r^h, assigns to every $h' \in H^h$ a mixed action $r^h(h') \in \Delta(A^{h'})$. Denote by R^h the set of decision rules at $h \in H$. A *decision rule* r^h is *optimal* at $h \in H$ if the probability measure over Z^h generated by r^h is weakly preferred to the probability measure over Z^h generated by any alternative decision rule $\tilde{r}^h$, when evaluating the probability measures by means of expected payoff. The decision rule at h specifies the behavior of the decision maker as long as the preferences at h, u^h, persist.

A *strategy* s assigns to every $h \in H$ a decision rule $s(h) \in R^h$. Denote by S the set of strategies. At h_0, the strategy s prescribes that the decision rule $s(h_0)$ be selected. At any later decision node $h \in H\backslash\{h_0\}$, the strategy s prescribes that $s(h)$ be selected if and only if prior preferences do not persist. Note that, with a positive perceived preference persistence p, it is essential that a strategy assigns to each $h \in H$ a decision rule rather than just a mixed action.

For every $h \in H$ and $h' \in H^h$, the strategy $s \in S$ and perceived preference persistence $p \in [0,1]$ generate a probability measure $\tilde{P}^{h|h'}(p,s)$ over $Z^{h'}$ as perceived at h, given that preferences persist from h to h'. Let $\tilde{v}^{h|h'} : [0,1] \times S \to \mathbb{R}$ assign to each $p \in [0,1]$ and $s \in S$ the perceived expected payoff $\tilde{v}^{h|h'}(p,s) = E_{\tilde{P}^{h|h'}(p,s)} u^h(z)$ that this combination of strategy and perceived preference persistence leads to at $h' \in H^h$ as perceived at $h \in H$, given that preferences persist from h to h'.

Denote by (r^h, s^{-h}) the strategy that is obtained from s by substituting r^h for $s(h)$ at h. We can now state our main definition.

Definition 1 A strategy $s \in S$ is a *multi-self subgame-perfect equilibrium (*MSSPE*) with perceived preference persistence* p if, for all $h \in H$, $h' \in H^h$ and $r^h \in R^h$,

$$\tilde{v}^{h|h'}(p,s) \geq \tilde{v}^{h|h'}(p,(r^h, s^{-h})) \,.$$

Hence, an MSSPE is a subgame-perfect equilibrium of the game induced by treating each self of the decision maker as a separate player having decision-making authority as long as his preferences persist, with the decision rules of the selves corresponding to the strategies of the players, and the strategy of the decision maker corresponding to the strategy profile of the induced game. In particular, existence results for the concept of subgame-perfect equilibrium can be applied.

In words, a strategy $s \in S$ is an MSSPE with perceived preference persistence p if there exists no decision nodes $h \in H$ and $h' \in H^h$ such that the decision maker at h gains at h' by selecting an alternative decision rule to the one assigned by s, given that preferences has persisted from h to h', s will be followed at all later decision nodes where prior preferences do not persist, and the decision maker perceives that the preference persistence probability is p.

If the decision maker has time-consistent preferences and an optimal decision rule exists at every decision node, then—independently of the perceived preference persistence—there is an MSSPE where an optimal decision rule is always chosen. On the other hand, if the decision maker has time-inconsistent preferences, then a strategy assigning an optimal decision rule at every decision node is an MSSPE when the perceived preference persistence is 1, not necessarily otherwise.

The above model nests both sophisticated behavior and fully naive behavior.

- By setting $p = 0$ we get ordinary sophistication: The decision maker acknowledges how his preferences will change, and selects at each decision node an optimal action, given the actions that the strategy prescribes at future decision nodes.

- By setting $p = 1$ we get full naivete: The decision maker perceives incorrectly at every decision node that his current preferences will persist, implying that the selected decision rule will be followed throughout the T-stage decision problem. However, when reaching the next stage, the decision maker discovers that his preferences have changed and he selects a new decision rule.

With $0 < p < 1$ we obtain a new theory of partial naivete: The decision maker perceives incorrectly that there are positive probabilities p, p^2, p^3, ... that his current preferences will persist and a selected decision rule will be followed at subsequent stages. The fact that his preferences do not persist does not contradict p, since it can be interpreted by the decision maker as a bad draw. In a Bayesian framework, the partially naive decision maker is rational if his prior distribution over values of p has all measure concentrated on one point so that updating does not change his perceived preference persistence. Although extreme, it is not irrational.

My alternative definition of partial naivete is general, in the sense that it can be applied to any system of conditional preferences. In comparison, O'Donoghue and Rabin's (2001) definition of partial naivete is tailored to (β, δ)–preferences, since it relies on the decision maker's point estimation of his future β.[4] In the next section I show how my alternative definition has different behavioral implications than the definition of partial naivete proposed by O'Donoghue and Rabin.

3 Procrastination as equilibrium behavior

In the present section I consider a special case of the general model presented in Section 2. I model a decision problem with an infinite number of stages, which concerns the timing of a task with immediate cost and future benefits. The set of decision nodes H is partitioned into two states ω^0 and ω^1: $H = \omega^0 \cup \omega^1$, with $\omega^0 \cap \omega^1 = \emptyset$. Here ω^0 denotes the state in which the decision maker has not performed the task, while ω^1 denotes the state in which the decision maker has performed the task. If the task has not been performed, i.e., $h \in \omega^0$, then the action set A^h equals

[4]In the context of (β, δ)–preferences it is possible to encompass both O'Donoghue and Rabin's (2001) modeling and the present paper's modeling of partial naivete by endowing the decision maker with some belief over his future value of β. This possibility is explored by Heidhues and Kőszegi (2008b). In contrast to the present definition, such a generalization does not extend to general time-inconsistent preferences.

$\{a^0, a^1\}$, where a^0 means to do nothing, while a^1 means to do the task. On the other hand, if the task has already been performed, i.e., $h \in \omega^1$, then the action set A^h equals $\{a^0\}$. Naturally, $(h, a^0) \in \omega^0$ and $(h, a^1) \in \omega^1$ if $h \in \omega^0$, while $(h, a^0) \in \omega^1$ if $h \in \omega^1$. With the assumption that the history at stage 0, h_0, is contained in ω^0—entailing that the task has not been performed at the root of the decision problem—the determination of the set of histories is complete. An outcome specifies in which stage the task is done; at stage 0, 1, 2, ..., or not at all.

To specify the decision maker's preferences, let $c > 0$ be the cost accruing at the stage in which the task is performed, and let $v > 0$ be the benefits accruing at the next stage. Assume that the decision maker has (β, δ)–preferences,

$$u^h(z) = \upsilon_{t(h)} + \beta \sum\nolimits_{t=t(h)+1}^{\infty} \delta^{t-t(h)} \upsilon_t$$

where $0 < \beta \leq 1$ and $0 < \delta < 1$, and where υ_t denotes the periodic payoff at stage t, with $\upsilon_t = -c$ if $a_t = a^1$, $\upsilon_t = v$ if $a_{t-1} = a^1$, and $\upsilon_t = 0$ otherwise. Hence, the payoff at decision node h of doing the task now is

$$-c \ + \ \beta\delta v \,,$$

the payoff at decision node h of doing the task $t - t(h) \geq 1$ stages from now is

$$\beta\delta^{t-t(h)}\big(-c + \delta v\big) \,,$$

while the payoff at decision node h of not doing the task at all is 0.

The following assumption on the decision maker's preferences ensures that this simple model serves as an illustration.

Assumption 1 It is better to perform the task now than not doing it at all, but even better to postpone the task to the next stage:

$$0 < -c + \beta\delta v < \beta\delta\big(-c + \delta v\big) \,.$$

The example of the introduction has parameter values ($\beta = 1/2$, $\delta = 4/5$, $c = 1$, and $v = 5$) that satisfy this assumption. Assumption 1 is satisfied if and only if

$$\beta\delta v \ > \ c \ > \ \beta\delta v \frac{1-\delta}{1-\beta\delta} \,.$$

Since $\beta < 1$ is necessary for these inequalities to hold, Assumption 1 implies that the decison maker has "present-biased preferences". With (β, δ)–preferences, it is w.l.o.g. to assume that the benefits accrue at the stage following the stage in which the task was performed, since v may represent the present value of future benefits discounted back to this stage by discount factor δ. For later use, denote by A the payoff ratio of (i) doing the task now and (ii) postponing it to the next period.

$$A := \frac{-c + \beta\delta v}{\beta\delta\big(- c + \delta v\big)} \, .$$

Assumption 1 implies that $A \in (0, 1)$.

Note that a decision rule specifies a non-trivial action choice only at decision nodes corresponding to the state in which the decision maker has not performed the task. Under the above assumptions the optimal decision rule is to do nothing *now*, and instead perform the task *at the next stage*. Of course, this is not time consistent. For a large perceived preference persistence p, a strategy that always assigns the optimal decision rule will turn out to constitute an MSSPE. For a small p this will not be the case, as the decision maker will perceive that the selected decision rule will not be followed at the next stage with a sufficiently high probability $1 - p$.

I will consider an MSSPE where the decision maker uses a stationary strategy, entailing that the same decision rule is assigned whenever prior preferences do not persist and the task has still not been performed. Hence, the decision maker is assumed *not* to let his selection depend on time *nor* to use self-enforcing schemes of self-reward and self-punishement to overcome his commitment problems.

When a stationary strategy is used, the equilibrium is characterized in terms of single decision rule. Note that a decision rule assigned by a stationary strategy constituting an MSSPE has the following two properties:

- Any such decision rule specifies that the task be performed at all future stages with probability 1, provided that prior preferences persist.

- Any such decision rule does not specify that the task be performed now with probability 1. The reason is that then the decision maker would believe that the task would be done with probability 1 at the next stage, independently of whether prior preferences persist. Hence, he would want to delay performing the task until the next stage, contradicting that the stationary strategy assigning this decision rule is an equilibrium.

For any decision rule r^h, let $r^h(h')$ denote the probability with which r^h specifies the choice of a^1 at $h' \in H^h$. Hence, if $h' \in H^h \cap \omega^0$, $r^h(h') \in [0,1]$, while if $h' \in H^h \cap \omega^1$, $r^h(h') = 0$. The following result characterizes the unique stationary MSSPE.

Proposition 1 *For given perceived preference persistence $p \in [0,1]$, there exists a unique stationary* MSSPE *s_p. The* MSSPE *s_p has the following properties:*

(1) If $h \in \omega^0$, then $s_p(h) = r_p^h$ where

$$r_p^h(h') = \begin{cases} q & \text{if} \quad h' = h, \\ 1 & \text{if} \quad h' \in H^h \backslash \{h\} \quad \text{and} \quad h' \in \omega^0 \\ 0 & \text{if} \quad h' \in H^h \backslash \{h\} \quad \text{and} \quad h' \in \omega^1 , \end{cases}$$

and $q \in (0,1)$ is determined by

$$A = \frac{p + q - pq}{1 - (1-p)(1-q)\delta} \tag{1}$$

if $A > p/\big(1-(1-p)\delta\big)$, while $q = 0$ if $A \leq p/\big(1-(1-p)\delta\big)$.

(2) If $h \in \omega^1$, then $s_p(h) = r_p^h$ where

$$r_p^h(h') = 0 \quad \text{for all} \quad h' \in H^h .$$

One may interpret q as the decision maker's belief about his future actions; it does not necessarily entail that the decision maker randomizes. This is consistent with the usual interpretation of Nash equilibrium as an equilibrium in beliefs.

Turn next to the comparative statics: How does the unique stationary MSSPE s_p vary with the perceived preference persistence? For the statement of this result, denote by $q(p)$ the probability determining the decision rule r_p^h assigned to $h \in \omega^0$ by s_p.

Proposition 2 *There exists $\bar{p} \in (0,1)$ such that*

$$q(p) \begin{cases} \in \ (0,1) & \text{if} \quad p \in [0, \bar{p}) \\ = \ 0 & \text{if} \quad p \in [\bar{p}, 1] \,. \end{cases}$$

The probability $q(p)$ is a continuous function for $p \in [0,1]$, continuously differentiable for $p \in (0, \bar{p})$, strictly decreasing function in p on $[0, \bar{p}]$, and constant in p on $[\bar{p}, 1]$. The critical perceived preference persistence $\bar{p}$ is determined by

$$A = \frac{\bar{p}}{1 - (1 - \bar{p})\delta} \,.$$

Proposition 2 entails that an increased perceived preference persistence p decreases the probability of doing task now, up to a critical level $\bar{p} \in (0,1)$, above which the task is postponed for sure.

Hence, the behavioral implications of the present modeling of procrastination (in the setting considered) can be described as follows:

(a) A sophisticated decision maker, having a perceived preference persistence equal to zero, will do the task in any stage with a probability between 0 and 1. This probability reflects the decision maker's own strategic uncertainty concerning whether the task will be performed at any future stage.

(b) A higher level of partial naivete, in the form of increased perceived preference persistence, lowers this probability *continuously* up to a critical level, above which the task is postponed for sure.

It is of interest to point out how these behavioral implications differ from those of O'Donoghue and Rabin (2001), on the one hand, and dual-self models, on the other hand, in the same setting.

With full sophistication, the modeling of O'Donoghue and Rabin (2001) is the same as mine, and behavioral implication (a) can be obtained within their framework by allowing for a stationary strategy implementing a mixed action as long as the task has not been performed. However, their modeling cannot replicate behavioral implication (b). To see this, assume that the decision maker is partially naive in the sense of O'Donoghue and Rabin (2001). Then, as O'Donoghue and Rabin (2001, footnote 20) explain, for parameter values satisfying Assumption 1, any degree of partial naivete leads to the task being postponed forever, provided that the decision maker believes a stationary strategy will be used by his future selves. Hence, going from full sophistication to a vestige of partial naivete in the sense of O'Donoghue and Rabin (2001) leads to a discontinuous change in behavior.[5] Heidhues and Kőszegi (2008a) show that this qualitative difference carry over to their model of self-control in the credit market.

When modeling the present setting by means of Fudenberg and Levine's (2006) dual-self model of impulse control or Gul and Pesendorfer's (2004) 'dynamic self-control' preferences, it holds generically that the decision maker, when faced with the binary choice of performing the task now or delaying the task to the next stage, will either do the task now, or not do the task at all. In a deterministic setting one will not observe that the task is postponed for a finite number of periods before, finally, being performed. Also, partial naivete, as such, cannot be explicitly modeled in the frameworks of Fudenberg and Levine (2006) and Gul and Pesendorfer (2004).

[5]E.g., with the parameter values used in the example of the introduction, suppose that the decision maker believes his future β will be 3/5, entailing partial naivete as his true β equals 1/2. If he believes that his future selves will apply a stationary strategy, then it follows (by setting $\beta = 3/5$ and $p = 0$ and applying Proposition 1) that the decision maker believes his future selves will perform the task in any stage with probability 7/8. This means that he now (and in the future) strictly prefers to postpone the task to the next stage, implying that the task will never be performed.

4 Welfare analysis in multi-self settings

To motivate the the general discussion of welfare analysis with time-inconsistent preferences, consider the model of Section 3 with $c = 1$ and $v = 5$. If the task is performed at stage 1, then the stream of periodic payoffs is

$$0,\ -1,\ 5,\ 0,\ 0,\ \dots\,,$$

while if the task is performed at stage 2, then the stream of periodic payoffs is

$$0,\ 0,\ -1,\ 5,\ 0,\ \dots\,.$$

With time-consistent preferences, say $\beta = 1$ and $\delta = 4/5$, the decision-maker at both stage 0 (i.e., at decision node h_0) and stage 1 (i.e., at decision node (h_0, a^0)) agree that performing the task at stage 1 (outcome z_1) is better than performing the task at stage 2 (outcome z_2):

$$u^{h_0}(z_1) = \tfrac{12}{5} > \tfrac{48}{25} = u^{h_0}(z_2) \quad \text{and} \quad u^{(h_0,a^0)}(z_1) = 3 > \tfrac{12}{5} = u^{(h_0,a^0)}(z_2)\,.$$

The situation is different with time-inconsistent preferences, say $\beta = 1/2$ and $\delta = 4/5$. Then the decision-maker at stage 0 (i.e., at decision node h_0) prefers to perform the task at stage 1, while the decision maker at stage 1 (i.e., at decision node (h_0, a^0)) prefers to perform the task at stage 2:

$$u^{h_0}(z_1) = \tfrac{6}{5} > \tfrac{24}{25} = u^{h_0}(z_2) \quad \text{and} \quad u^{(h_0,a^0)}(z_1) = 1 < \tfrac{6}{5} = u^{(h_0,a^0)}(z_2)\,. \tag{2}$$

This conflict of interests in the case of time-inconsistent preferences is the motivation for using the Pareto-criterion in such cases. While with time-consistent preferences, all selves on the path leading up to a decision (in this case, whether to do the task at stage 1 or 2) agree on what to do, this is the not the case with time-inconsistent preferences. So with time-inconsistent preferences, we cannot in this case rank the alternatives due to the conflict of interests between the selves at h_0 and (h_0, a^0).

Note that all other decision nodes are irrelevant for deciding between z_1 and z_2. At decision nodes where the task has been performed (e.g., after (h_0, a^1)), there is no task to perform, while when stage 2 or later has been reached, it is no longer possible to perform the task at stage 1. The general principle is to compare decision rules at each and every decision node. A decision rule at h_0 implementing z_1 is

$$r_1^{h_0}(h) = \begin{cases} 1 & \text{if} \quad h \in H^{(h_0,a^0)} \cap \omega^0 \\ 0 & \text{otherwise,} \end{cases}$$

while a decision rule at h_0 implementing z_2 is

$$r_2^{h_0}(h) = \begin{cases} 1 & \text{if} \quad h \in H^{(h_0,a^0,a^0)} \cap \omega^0 \\ 0 & \text{otherwise.} \end{cases}$$

Hence, $r_1^{h_0}$ ($r_2^{h_0}$ respectively) specifies to perform the task at stage 1 (2 respectively) and, if it is not done by then, to do it as soon as possible. These two decision nodes implement the same outcome at all decision nodes except for $h = h_0$ and $h = (h_0, a^0)$. It follows from (2) that the decision maker at $h = h_0$ prefers $r_1^{h_0}$ to $r_2^{h_0}$, while the opposite is the case for the decision maker at $h = (h_0, a^0)$.

Turning now to the general model of Section 2 and taking explicitly into account that the decision maker is not able to commit to one decision rule throughout the decision problem, one must compare alternative strategies (not decision rules) *at each and every decision node.* I will take the position that the comparisons should be based on the actions that will actually be taken throughout the decision problem, not the actions that the decision maker naively thinks he will take.

Since there is no actual preference persistence, it follows that for every $h \in H$, a strategy $s \in S$ generates an actual probability measure $P^h(s)$ over the set of outcomes Z^h, independently of the perceived preference persistence. Let $v^h : S \to \mathbb{R}$ assign to each $s \in S$ the actual expected payoff $v^h(s) = E_{P^h(s)} u^h(z)$ that this strategy leads to. A formal definition of Pareto-dominance in the context of our general model of Section 2 can now be stated.

Definition 2 A strategy $s' \in S$ *Pareto-dominates* another strategy $s'' \in S$ if

$$v^h(s') \geq v^h(s'') ,$$

for all $h \in H$, with strict inequality for some $h' \in H$.

I use this criterion to evaluate the welfare effects of partial naivete in Section 5 and show how it can used to rank stationary strategies in the model of Section 3 under Assumption 1.

It is of interest to note that Pareto-dominance as defined in Definition 2 is formally related to the dominance relation used to define concepts of renegotiation-proofness in repeated games (see, e.g., Farrell and Maskin, 1989, and their concept of weak renegotiation-proofness). I have in Asheim (1997) used the same dominance relation to refine of the concept of MSSPE in the context of individual planning with time-inconsistent preferences. In that paper I explain and exploit the analogy between, on the one hand, a single decision maker with inconsistent preferences revising his decision rule and, on the other hand, the grand coalition in a repeated game renegotiating away from a continuation equilibrium that punishes all players. However, neither the literature on renegotiation-proofness nor my modeling of consistent planning in Asheim (1997) is concerned with welfare analysis.

In the multi-self literature on time inconsistency and procrastination, other welfare criteria have been applied. In the case of (β, δ)–preferences, the most common practice seems to be to set $\beta = 1$ and treat $\sum_{t=0}^{\infty} \delta^t v_t$ as the welfare criterion, where v_t denotes the periodic payoff at stage t. O'Donoghue and Rabin (1999, 2001) use this welfare criterion, which they refer to as "a person's long-run utility", while Gul and Pesendorfer (2005, Section 6.4) question its justification.

Measuring welfare by means of "long-run utility" can be motivated by a paternalistic concern for the well-being of the decision maker, where $\beta < 1$ is interpreted as a defect of his decision making capabilities. This entails the normative position that β *should* equal 1.

"Long-run utility" utility can also be provided with a non-paternalistic justification, if we are willing to assert the following:

- Assume that the decision maker has no decision to make at stage 0 (i.e., at the root of the decision tree h_0, the action set A^{h_0} is a singleton), entailing that "life" starts at stage 1.
- Adopt the normative position that *only* the self of the decision maker at the root of the decision tree h_0 has normative significance.

The combination of these two points means that β plays no rule and the ranking of strategies are made in accordance with "long-run utility". Hence, they imply in the setting of the general model of Section 2 that $v^{h_0}(s)$ measures the welfare generated by strategy s, provided that A^{h_0} is a singleton.

The present paper is based on the presumption that the problem of the decision maker is not his preferences, but his naivete, and his lack of access to a commitment mechanism. Hence, for the purpose of this paper, I do not adopt a paternalistic position. Furthermore, even if A^{h_0} is a singleton, so that there is no decision to make at stage 0, it might be difficult to argue that the self at the root of the decision tree should be treated as a dictator, trumping the interests of all future selves.

If one attempts to apply "long-run" utility for the ranking of stationary strategies in the model of Section 3 under Assumption 1, it leads to the result that q, the probability with which the task is performed, should equal 1. However, for any history at which the task has not yet been performed, the decision maker strictly prefers setting $q < 1$.[6] Hence, if one respects the preferences of the decision maker, "long-run" utility is not suitable for ranking stationary strategies in this setting.

Some researchers (e.g., Goldman, 1979; O'Donoghue and Rabin, 2001) have used the Pareto criterion differently from Definition 2. In their use of the concept, one

[6] With $c = 1$, $v = 5$, $\beta = \frac{1}{2}$ and $\delta = \frac{4}{5}$, any $q \in (\frac{1}{2}, 1)$ is better than $q = 1$, with the most preferred stationary strategy being determined by $q = (-1 + \sqrt{15})/4 \approx 0.72$.

stream of periodic payoffs is considered unambiguously as good as another if it is weakly preferred by the decision maker *at all stages*. Returning to the two streams of periodic payoffs given at the beginning of this section, they compare not only $u^{h_0}(z_1)$ to $u^{h_0}(z_2)$ and $u^{(h_0,a^0)}(z_1)$ to $u^{(h_0,a^0)}(z_2)$, but also $u^{(h_0,a^0,a^1)}(z_1) = 5$ to $u^{(h_0,a^0,a^0)}(z_2) = 1$ and $u^{(h_0,a^0,a^1,a^0)}(z_1) = 0$ to $u^{(h_0,a^0,a^0,a^1)}(z_2) = 5$. Hence, at stage 0 the decision maker prefers z_1 to z_2, at stage 1 he switches to preferring z_2 to z_1, at stage 2 he wishes the task had already been done and reswitches back to preferring z_1 to z_2, while finally, at stage 3 he thinks the later execution date is better and prefers once more z_2 to z_1.

However, the comparisons at stages 2 and 3 are made across different histories. That such comparisons lead to nonsensical results is illustrated by O'Donoghue and Rabin (2001) in their footnote 21. And as O'Donoghue and Rabin (2001, footnote 21) observe, this use of the Pareto-criterion leads to a conflict of interest between different selves even if the decision maker has time-consistent preferences. In contrast, the concept of Pareto-dominance proposed in Definition 2 reduces to the usual criterion for intertemporal choice with time-consistent preferences. Moreover, Definition 2 applies the actual preferences of the decision maker when comparing costs and benefits accruing at different future stages.[7]

In the comparison of stationary strategies in the model of Section 3 under Assumption 1, it is straightforward to see that, when going from one q to another, the decision maker will lose out at some stage. This entails that a version of the Pareto criterion that compares the decision maker's actual expected payoffs at different *stages* cannot be used to rank stationary strategies in this setting.

[7] A non-paternalistic position entails that only the preferences of the decision maker should matter. With forward-looking preferences, the decision maker's preferences can compare only future consequences; past consequences cannot be taken into account. The issues that arise if the decision maker is endowed with preferences also over past consequences are not discussed here, as they are unrelated to the usual modeling of time-inconsistent preferences in the multi-self model.

5 Welfare effects of procrastination

In this section I use Pareto-dominance as proposed in Definition 2 for the welfare analysis of procrastination. In particular, I do comparative statics of the equilibrium strategies established in Proposition 1 w.r.t. perceived preference persistence p.

Following Definition 2, one strategy s' Pareto-dominance another s'' if $v^h(s') \geq v^h(s'')$ for all $h \in H$, with strict inequality for some $h' \in H$. When applied to the equilibria of the special model of Section 3, Definition 2 is simple to apply. The reason is that, for each perceived preference persistence p, the MSSPE s_p of Proposition 1 is a stationary strategy, entailing that the same decision rule is assigned whenever prior preferences do not persist and the task has still not been performed. Hence, the welfare indicator $v^h(s_p)$ is the same at all decision nodes h at which the task has not yet been performed (i.e., $h \in \omega^0$). Since there is only one feasible decision rule at decision nodes h at which the task has already been performed (i.e., $h \in \omega^1$), and thus, the problem of selecting a decision rule is trivial at such nodes, we need only be concerned with the common value of $v^h(s_p)$ for all $h \in \omega^0$.

The main welfare result is stated by the following proposition.

Proposition 3 *The welfare indicator $v^h(s_p)$ is strictly decreasing in p on $[0, \bar{p}]$ and constant in p on $[\bar{p}, 1]$ for all $h \in \omega^0$, and constant in p on $[0, 1]$ for all $h \in \omega^1$.*

Proposition 3 yields an unambiguous welfare conclusion: *Having a positive perceived preference persistence reduces welfare, since such partial naivete leads to welfare reducing procrastination.*

The present analysis can also be used to evaluate the welfare effects of introducing a commitment device. E.g., suppose that a mechanism was offered, enabling the decision maker at stage t (where $t = 0, 1, \ldots$) to make a costless commitment to performing the task at stage $t + 1$. If the perceived preference persistence is smaller than 1, then a stationary MSSPE leads to the choice of the commitment mechanism at any decision node where the task has not already been performed. Furthermore, Definition 2 entails that this equilibrium welfare dominates any other stationary strategy, since comparisons need only be made at decision nodes where the task has still not been performed, and the decision maker prefers to make the commitment at all such nodes. This means that policy offering the decision maker an opportunity to choose such a commitment device is welfare enhancing. In particular, according to Definition 2 it is irrelevant that the decision maker will regret his commitment when being forced to perform the task, as long as his selves up to the point at which the commitment was made unanimously agreed.

6 Concluding remarks

In this paper I have applied a simple model of procrastination as a vehicle for demonstrating the need for and the consequences of (i) a new notion of partial naivete and (ii) a new principle for welfare analysis in multi-self models. These are the two main contributions of this paper, with significance beyond the problem of procrastination.

The new notion of partial naivete has been used to explain why people end up doing tasks later than they originally thought. This notion does not depend on a particular specification of present-biased preferences and can applied without addressing any underlying game-theoretic issues.

The new principle for welfare analysis in multi-self models has been used to conclude that such procrastination leads to people performing tasks later than they wish they would. This principle has the attractive feature of being based on the actual preferences of the decision maker when comparing costs and benefits accruing at different times. Furthermore, it reduces to the usual criterion for intertemporal choice if preferences are time consistent.

Appendix: Proofs

Proof of Proposition 1. The discussion preceding the proposition entails that only q is to be determined. Denote by $V(p,q)$ the expected present value of future payoffs conditional on the task not having been performed yet and not being performed now, discounted back to the next stage using $(1,\delta)$–discounting, under the assumptions that (i) preferences persist at the next stage with probability p, and (ii) a stationary strategy is applied, prescribing a decision rule which specifies that the task be performed with probability q now and with probability 1 at the next stage. Since

$$V(p,q) = (p+q-pq)\big(-c+\delta v\big) + (1-p)(1-q)\delta V(p,q)\,,$$

we obtain that

$$V(p,q) = \frac{p+q-pq}{1-(1-p)(1-q)\delta}\big(-c+\delta v\big)\,.$$

It follows that s_p is an MSSPE if and only if one of the three cases holds:

(a) $q=0$ and $-c+\beta\delta v \le \beta\delta V(p,q)$

(b) $q\in(0,1)$ and $-c+\beta\delta v = \beta\delta V(p,q)$

(c) $q=1$ and $-c+\beta\delta v \ge \beta\delta V(p,q)$

Since by Assumption 1 case (c) can never hold, we are left with the two remaining cases. If $p=1$, then by Assumption 1 we must be in case (a), with $q=0$ and $A<1=p/\big(1-(1-p)\delta\big)$. If $p\in[0,1)$, then

$$\frac{p+q-pq}{1-(1-p)(1-q)\delta}$$

is an increasing function of q, and it follows that q is uniquely determined as specified in the proposition. This establishes the proposition. ■

Proof of Proposition 2. Write $f(p) := p/\big(1-(1-p)\delta\big)$. Since $A\in(0,1)$ (by Assumption 1), $f(0)=0$, $f(1)=1$, and $f(\cdot)$ is continuous and strictly increasing, we obtain that $\bar{p}\in(0,1)$ is uniquely determined. It now follows from Proposition 1 that $q(p)$ is a continuous function for $p\in[0,1]$, $q(p)=(0,1)$ if $p\in[0,\bar{p})$ and $q(p)=0$ if $p\in[\bar{p},1]$. Finally, for $p\in(0,\bar{p})$, $q(p)$ is continuously differentiable and

$$q'(p) = -\frac{1-q(p)}{1-p} < 0\,,$$

since $A=\big(p+q-pq\big)/\big(1-(1-p)(1-q)\delta\big)$ in this range of p values. ■

Proof of Proposition 3. The statements for $h \in \omega^1$ are trivially true. Hence, consider $h \in \omega^0$. Applying the notation of Section 3, we have that

$$v^h(s_p) = q(p)\big(-c + \beta\delta v\big) + (1 - q(p))\beta\delta V(0, q(p)) \tag{3}$$

if $h \in \omega^0$. Since, by Proposition 2, $q(p)$ is a continuous function for $p \in [0, 1]$ and continuously differentiable for $p \in (0, \bar{p})$, it follows that $v^h(s_p)$ is a continuous function of p for $p \in [0, 1]$ and continuously differentiable for $p \in (0, \bar{p})$.

Propositions 1 and 2 imply that, for $p \in [0, \bar{p})$,

$$-c + \beta\delta v \quad = \quad \beta\delta V(p, q(p)) \quad \geq \quad \beta\delta V(0, q(p))$$

since

$$\frac{p + q - pq}{1 - (1 - p)(1 - q)\delta}$$

is a nondecreasing function of p. Hence, it follows from (3) that, for $p \in (0, \bar{p})$,

$$\frac{dv^h(s_p)}{dp} \quad < \quad q'(p)\Big(\big(-c + \beta\delta v\big) - \beta\delta V(0, q(p))\Big) \quad \leq \quad 0$$

since $q'(p) < 0$ by Proposition 2 and

$$\frac{q}{1 - (1 - q)\delta}$$

is an increasing function of q. Hence, $v^h(s_p)$ is decreasing in p on $[0, \bar{p}]$. ■

References

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

Akin, Z. (2007), Time inconsistency and learning in bargaining games, *International Journal of Game Theory* **36**, 275–299.

Asheim, G.B. (1997), Individual and collective time consistency, *Review of Economic Studies* **64**, 427–443.

Bernheim, D. and Rangel, A. (2007), Behavioral public economics: Welfare and policy analysis with non-standard decision-makers, in Diamond, P. and Vartiainen, H. (eds.), *Economic Institutions and Behavioral Economics*. Princeton University Press, Princeton.

Brocas, I. and Carrillo, J.D. (2001), Rush and procrastination under hyperbolic discounting and interdependent activities, *Journal of Risk and Uncertainty* **22**, 141–164.

Dekel, E., Lipman, B.L., and Rustichini, A. (2008), Temptation-driven preferences, mimeo, Northwestern Univesity, Boston University and University of Minnesota.

Eliaz, K. and Spiegler, R. (2006), Contracting with diversly naive agents, *Review of Economic Studies* **73**, 689–714.

Elster, J. (1979), *Ulysses and the Sirens: Studies in Rationality and Irrationality*, Cambridge University Press, Cambridge, UK.

Farrell, J. and Maskin, E. (1989), Renegotiation in repeated games, *Games and Economic Behavior* **1**, 327–360.

Fischer, C. (1999), Read this paper even later: Procrastination with time-inconsistent preferences, Discussion Paper 99-20, Resources for the Future.

Frederick, S., Loewenstein, G., and O'Donoghue, T. (2003), Discounting and time preference: A critical review, *Journal of Economic Literature* **40**, 351–401.

Fudenberg, D. and Levine, D. (2006), A dual-self model of impulse control, *American Economic Review* **96**, 1449–1476.

Goldman, S.M. (1979), Intertemporally inconsistent preferences and the rate of consumption, *Econometrica* **47**, 621–626.

Gul, F. and Pesendorfer, W. (2001), Temptation and self-control, *Econometrica* **69**, 1403–1435.

Gul, F. and Pesendorfer, W. (2004), Self-control and the theory of consumption, *Econometrica* **72**, 119–158.

Gul, F. and Pesendorfer, W. (2005), The case for mindless economics, mimeo, Princeton University.

Heidhues, P. and Kőszegi, B. (2008a), Exploiting naivete about self-control in the credit market, mimeo, University of Bonn and University of California Berkeley

Heidhues, P. and Kőszegi, B. (2008b), Futile attempts at self-control, paper presented at 2008 EEA Congress, Milan.

Laibson, D. (1994), *Essays in Hyperbolic Discounting*, Ph.D. Thesis, Massachusetts Institute of Technology.

Laibson, D. (1997), Golden eggs and hyperbolic discounting, *Quarterly Journal of Economics* **112**, 443–477.

Loewenstein, G. and Prelec, D. (1992), Anomalies in Intertemporal Choice, *Quarterly Journal of Economics* **107**, 573–597.

O'Donoghue, T. and Rabin, M. (1999), Doing it now or later, *American Economic Review* **89**, 103–124.

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

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

Rubinstein, A. (2003), Economics and psychology? The case of hyperbolic discounting, *International Economic Review* **44**, 1207–1216.

Thaler, R.H. and Shefrin, H.M. (1981), An economic thoery of self-control, *Journal of Political Economy* **89**, 392–406.