Choice and Procrastination

Ted O'Donoghue
Department of Economics
Cornell University

and

Matthew Rabin
Department of Economics
University of California, Berkeley

March 17, 1999

Abstract

Recent models of procrastination due to self-control problems assume that a procrastinator considers just one option and is unaware of her self-control problems. We develop a model where a person chooses from a menu of options and is partially aware of her self-control problems. This menu model replicates earlier results and generates new ones. A person might forego completing an attractive option because she plans to complete a more attractive but never-to-be-completed option. Hence, providing a non-procrastinator additional options can induce procrastination, and a person may procrastinate worse pursuing important goals than unimportant ones.

Keywords: Choice, Naivete, Partial Naivete, Present-biased Preferences, Procrastination, Self Control, Sophistication, Time Inconsistency.

JEL Category: A12, B49, C70, D11, D60, D74, D91, E21

Acknowledgments: We thank Kitt Carpenter and Erik Eyster for valuable research assistance, and Erik Eyster and seminar participants at Harvard and the LSE for helpful comments. For help in interpreting Voltaire, we thank Roland Benabou, Alan Durrell, Erik Eyster, Geraldine Legendre, Julian LeGrand, Drazen Prelec, and Dan Rabin. For financial support, we thank the National Science Foundation (Award 9709485), and Rabin thanks the Russell Sage, Alfred P. Sloan, and MacArthur Foundations. This research was started while Rabin was a Fellow at the Center for Advanced Study in the Behavioral Sciences, supported by NSF Grant SBR-960123. He is grateful for the Center's hospitality and the NSF's support.

Mail: Ted O'Donoghue / Department of Economics / Cornell University / 414 Uris Hall / Ithaca, NY 14853-7601. Matthew Rabin / Department of Economics / 549 Evans Hall #3880 / University of California, Berkeley / Berkeley, CA 94720-3880. Email: edo1@cornell.edu and rabin@econ.berkeley.edu. CB handles: "golf boy" and "game boy". Web pages (with related papers): http://www.people.cornell.edu/pages/edo1/ and http://elsa.berkeley.edu/~rabin/index.html

"The better is the enemy of the good."

— Voltaire

1. Introduction

People procrastinate — we delay doing unpleasant tasks that we wish we would do sooner. Such procrastination can be very costly. We skip enjoyable events in mid-April because we procrastinate in completing our taxes; we die young because we procrastinate in quitting smoking, starting a diet, or scheduling a medical check-up; and we are denied tenure because of our own, co-authors', or journal referees' procrastination.

Recent formal models have shown how procrastination can arise from self-control problems, conceived of as a time-inconsistent taste for immediate gratification. These models assume that a potential procrastinator has only one task under consideration. In most situations, however, a person must decide not only *when* to complete a task, but also *which* task to complete, or how much effort to apply to a chosen task. If a person must buy a housewarming present for a friend, she can either buy a gift certificate or spend time shopping for an appropriate present. If she is planning to resubmit a paper for publication, she can either respond minimally to the editor's suggestions or expend more effort to respond thoroughly. If she is putting together a montage of Johnny Depp photos, she can either haphazardly throw together a few press clippings or work devoutly to construct the shrine that he deserves. If she is choosing how to invest some money, she can either thoughtlessly follow the advice of a friend, or thoroughly investigate investment strategies.

In this paper, we develop a model of procrastination where a person must choose not only when to do a task, but which task to do. The model makes a number of realistic predictions incompatible with the conventional assumption of time-consistent preferences. These include the possibilities that providing a person with an attractive new option can cause her to switch from doing something beneficial to doing nothing at all, and that a person may procrastinate more severely when pursuing important goals than unimportant ones.

This paper also develops a formal model of *partial naivete*, where a person is aware that she will have future self-control problems, but underestimates their magnitude. The literature on self-control problems has entirely focused on two assumptions regarding a person's beliefs about her future self-control problems: That she is *sophisticated* — fully aware of her future self-control problems — or that she is *naive* — fully unaware of her future self-control problems. We believe that introducing a model of partial naivete to the growing literature on time-inconsistent preferences is an important ancillary contribution of the paper. Economists have been predisposed to focus on complete sophistication; but since many of our predictions hold for *any* degree of naivete, our analysis suggests that this could be a methodological and empirical mistake even if people are mostly sophisticated.

In Section 2, we describe a formalization of time-inconsistent preferences originally developed by Phelps and Pollak (1968) in the context of intergenerational altruism, and later employed by Laibson (1994) to capture self-control problems within an individual: In addition to time-consistent discounting, a person always gives extra weight to current well-being over future well-being. These "present-biased preferences" imply that each period a person tends to pursue immediate gratification more than she would have preferred if asked in any prior period.

In Section 3 we present our model of task choice. We suppose that a person faces a menu of possible tasks, or different effort levels available for a given task. Each period she must either complete one of these tasks or do nothing, without being able to commit to future behavior. Completing a task requires that the person incur an immediate cost, but generates an infinite stream of delayed benefits; tasks may differ in both their costs and their benefits. We assume that at most one task can be done, and can be done at most once. We further assume that the person behaves optimally given her beliefs as to how she will behave in the future, where her beliefs reflect her (sophisticated, naive, or partially naive) perceptions of her future self-control problems.

Naivete about future self-control problems leads a person to be over-optimistic about how soon she would complete the task if she were to delay now, and hence is an important determinant of procrastination. Akerlof (1991) emphasizes the role of naivete in putting off unpleasant tasks, and O'Donoghue and Rabin (forthcoming *a*) show that even mild self-control problems can cause severe procrastination for a completely naive person, but not for a completely sophisticated person.[1] Section 3 fleshes out the logic behind these earlier results, and generalizes them by allowing

[1] Prelec (1989) discusses how time-inconsistent preferences can lead a person to avoid doing an unpleasant task. Because he does not look at a dynamic model, sophistication is not relevant. Fischer (1997) considers procrastination of a task that may take a while to complete. She assumes sophistication, though because she explores long-term projects she finds substantial procrastination is still possible. Akerlof (1991) does not frame his analysis of procrastination in terms of time-inconsistent preferences, but his model implicitly corresponds to a model of present-biased preferences, and he highlights the role of naive beliefs in generating severe procrastination. O'Donoghue and Rabin (forthcoming *a*) explicitly compare the naive to the sophisticated model; O'Donoghue and Rabin (1998, forthcoming *b*) explore procrastination with naive beliefs.

for both a menu of tasks and partial naivete. We show that for any specific environment, severe procrastination can occur only with a non-negligible degree of naivete; even so, for any degree of naivete, no matter how little, there exists an environment where a person with that degree of naivete procrastinates severely.

In Section 4, we turn to the core new results of this paper — those regarding the role of choice for procrastination. Choice affects procrastination because two different criteria determine which task a person plans to do and when she does it. A person plans to do the task with the highest long-run net benefit, taking into account her taste for immediate gratification. But whether a person ever completes that task depends on a comparison of its immediate cost to the benefits forgone by brief delay, and has very little to do with either its long-run benefit or the features of other tasks available.

The fact that a person virtually ignores the other tasks available when deciding whether to immediately complete her chosen task implies that providing the person with additional options may in fact induce procrastination. If a new option has a sufficiently high long-run net benefit, the person will plan to do this new option rather than what she would have otherwise done; and if this new option has a sufficiently large cost relative to its immediate benefit, the person now procrastinates. For example, a person might immediately invest her savings in her company's 401K plan when offered a single portfolio in which to invest, but might procrastinate when offered a list of portfolios. As Voltaire *should have* meant by the opening quote (but didn't), a person may never complete a good task because of persistent, but unfulfilled, aspirations to do a better job.[2]

The fact that a person plans to do a task based on its long-run benefit, but follows through based on its immediate cost, implies that factors which make it beneficial to opt for a high-cost task may increase the propensity to procrastinate. This has an important implication: A person may be more likely to procrastinate in pursuit of important goals than in pursuit of unimportant ones. The more important are a person's goals, the more ambitious are her plans. But the more ambitious are her plans — i.e., the higher is the effort she intends to incur — the more likely she is to procrastinate in executing those plans. We formalize this intuition by supposing the long-run net benefit of all tasks are increased either by making the person more patient or by increasing per-period benefits, and identify classes of situations where a sufficiently large increase in the long-run benefits of all tasks induces a person to procrastinate.

[2] Although we, like many people, interpreted Voltaire to be referring to procrastination, a proper reading of Voltaire's (1878) statement in the original Italian makes clear that he meant something more akin to "If it ain't broke, don't fix it."

Our model does not imply that people *always* procrastinate the most when pursuing their most important goals. Indeed, this possibility requires the combination of self-control problems, naivete, and multiple options; when a person has no self-control problems, or is completely aware of them, or has only one reasonable option, increasing the long-run net benefits of all tasks makes the person more likely to do a task. And even with all three factors present, increased importance can sometimes reduce procrastination. But our model shows that any presumption that people don't procrastinate on important tasks should be dismissed.[3]

We view it as neither a flaw nor a virtue that some of our results are paradoxical from the perspective of traditional economic analysis. Rather, we are interested in their economic relevance. In O'Donoghue and Rabin (1998), for instance, we argue with some calibration exercises that such issues can be an important determinant of whether and how a person invests her savings for retirement. Investing for retirement is perhaps the single most important economic decision that people (should) make. Our theoretical model matches what seems to be empirically true: In spite of — or perhaps *because* of — its immense importance, many people never get around to carefully planning their investment for retirement. We conclude the paper in Section 5 with a brief discussion of the results in that paper, as well as a discussion of how the intuitions in this paper might play out in extensions of our model, such as supposing a person must allocate time among more than one task, or can improve upon what she has done in the past.

2. Present-Biased Preferences and Beliefs

The standard economics model assumes that intertemporal preferences are *time-consistent*: A person's relative preference for well-being at an earlier date over a later date is the same no matter when she is asked. But there is a mass of evidence that intertemporal preferences take on a specific form of *time inconsistency*: A person's relative preference for well-being at an earlier date over a later date gets stronger as the earlier date gets closer.[4] In other words, the evidence suggests that people have self-control problems caused by a tendency to pursue immediate gratification in a way

[3] And as such, our model is another example where careful analysis does not bear out the commonplace conjecture that harmfully irrational behavior is eliminated by ("sufficiently") large stakes.

[4] See, for instance, Ainslie (1975, 1991, 1992), Ainslie and Haslam (1992*a*, 1992*b*), Loewenstein and Prelec (1992), Thaler (1991), and Thaler and Loewenstein (1992). While the rubric of "hyperbolic discounting" is often used to describe such preferences, the qualitative feature of the time inconsistency is more general, and more generally supported by empirical evidence, than the specific hyperbolic functional form.

that their "long-run selves" do not appreciate.

In this paper, we apply a very simple form of such *present-biased preferences*, using a model originally developed by Phelps and Pollak (1968) in the context of intergenerational altruism, and later used by Laibson (1994) to model time inconsistency within an individual.[5] Let u_t be the instantaneous utility a person gets in period t. Then her intertemporal preferences at time t, U^t, can be represented by the following utility function:

$$\text{For all } t,\ U^t(u_t, u_{t+1}, ..., u_T) \equiv \delta^t u_t \ + \ \beta \sum_{\tau=t+1}^{T} \delta^\tau u_\tau.$$

This two-parameter model is a simple modification of the standard one-parameter, exponential-discounting model. The parameter δ represents standard "time-consistent" impatience, whereas the parameter β represents a time-inconsistent preference for immediate gratification. For $\beta = 1$, these preferences are time-consistent. But for $\beta < 1$, at any given moment the person has an extra bias for now over the future. A convenient (though perhaps unrealistic) feature of these preferences is that β is irrelevant when a person evaluates future trade-offs — that is, if in period t a person evaluates trade-offs that involve only periods $t+1$, $t+2$, ..., T, then β does not affect her preferences.

To examine intertemporal choice given time-inconsistent preferences, one must ask what a person believes about her own future behavior. Two extreme assumptions have appeared in the literature: *Sophisticated* people are fully aware of their future self-control problems and therefore correctly predict how their future selves will behave, and *naive* people are fully *un*aware of their future self-control problems and therefore believe their future selves will behave exactly as they currently would like them to behave.[6]

While our main goal in this paper is to extend the analysis of procrastination to situations where a person has a choice among tasks to perform, an ancillary goal is to extend the analysis of time-inconsistent preferences beyond the extreme assumptions of sophistication and naivete. Hence, we also examine behavior for a person who is *partially naive* — she is aware that she has future self-control problems, but she underestimates their magnitude. To formalize this notion, let $\hat{\beta}$ be a person's beliefs about her future self-control problems — her beliefs about what her taste for

[5] This model has since been used by Laibson (1995, 1997), Laibson, Repetto, and Tobacman (1998), O'Donoghue and Rabin (forthcoming *a*, forthcoming *b*, 1998), Fischer (1997), and others.

[6] Strotz (1956) and Pollak (1968) carefully lay out these two assumptions (and develop the labels), but do not much consider the implications of assuming one versus the other. Fischer (1997) and Laibson (1994, 1995, 1997) assume sophisticated beliefs. O'Donoghue and Rabin (forthcoming *a*, 1998) consider both, and explicitly contrast the two.

immediate gratification, β, will be in all future periods. A sophisticated person knows exactly her future self-control problems, and therefore has perceptions $\hat{\beta} = \beta$. A naive person believes she will not have future self-control problems, and therefore has perceptions $\hat{\beta} = 1$. A partially naive person then has perceptions $\hat{\beta} \in (\beta, 1)$. In the next section, we shall define within our specific model a formal solution concept that applies to sophisticates, naifs, partial naifs, and (by setting $\beta = 1$) time-consistent agents. We then show in the context of our model how and when partial naivete leads to procrastination.[7]

3. The Model and Some Results

Suppose there is an infinite number of periods in which a person can complete a task, and each period the person chooses from the same menu of tasks, $X \subset \mathbb{R}^2_+$. We often illustrate our framework by assuming X is a finite set, but our analysis also holds for infinite X (in which case we assume it is closed). Task $x \in X$ can be represented by the pair (c, v), where if a person completes task x in period τ she incurs cost $c \geq 0$ in period τ and initiates a stream of benefits $v \geq 0$ in each period from period $\tau + 1$ onward.[8] While we discuss more realistic alternatives in the conclusion, throughout our analysis we assume that the tasks are mutually exclusive and final: The person can complete at most one task, and can complete that task at most once.

[7] For simplicity, we shall abstract away from some complications that might arise with partial naivete. First, we shall assume a person is always absolutely positive — though wrong when $\hat{\beta} > \beta$ — about her future self-control problems. We doubt that our qualitative results would change much if the person had probabilistic beliefs whose mean underestimated the actual self-control problem. But it is central to our analysis that a person not fully learn over time her true self-control problem, or, if she does come to recognize her general self-control problem, she still continues to underestimate it on a case-by-case basis. Second, we shall assume that all higher-order beliefs — e.g., beliefs about future beliefs — are also equal to $\hat{\beta}$. Hence, a person has what might be called "complete naivete about her naivete": A partially-naive person thinks she will be entirely aware in the future of what she now believes is the extent of her future self-control problems (since otherwise she predicts she will forget what she currently knows). While we think our modeling choice here is the most realistic and most tractable, alternatives are not without merit. We suspect that sometimes people do realize that they are too often over-optimistic, and predict their own future behavior based on such a prediction of future misperceptions. In O'Donoghue and Rabin (forthcoming *b*), in fact, we informally invoke such "metasophistication" as a potential justification for assuming that people might be *aware* that they are naive procrastinators while still being naive procrastinators. Perhaps such metasophistication could be formalized by saying that a person predicts her future self-control problem to be $\hat{\beta}$, and predicts her future *beliefs* about her future self-control problems to be $\hat{\beta}' > \hat{\beta}$. This would predict that people may wish to impose incentives on their future selves that will protect themselves from their own naivete.

[8] Our model could be interpreted as a reduced form of a more general model where there are also immediate rewards and delayed costs. With this interpretation, our underlying assumptions are that the immediate costs are larger than the immediate benefits and that the delayed benefits are larger than the delayed costs. Also note that throughout this paper we assume the rewards for completing a task are exogenously determined; O'Donoghue and Rabin (forthcoming *b*) consider how one might design incentives to combat procrastination.

Since in each period the person can either complete a task or do nothing, we define $A \equiv X \cup \{\emptyset\}$ to be the set of actions available each period. Action $x \in X$ means "complete task x", and action $\emptyset$ means "do nothing". We describe behavior by a "strategy" $\mathbf{s} \equiv (a_1, a_2, ...)$ which specifies an action $a_t \in A$ for each period t. Although the choice a_τ matters only if $a_t = \emptyset$ for all $t < \tau$, a strategy specifies an action for all periods.[9] In this environment, there are two relevant questions about a person's behavior: (1) When, if at all, does she complete a task? and (2) Which task does she complete? Given a strategy $\mathbf{s}$, let $\tau(\mathbf{s})$ denote the period in which the person completes a task, and let $x(\mathbf{s})$ denote the specific task that the person completes. Formally, $\tau(\mathbf{s}) = \min\{t \mid a_t \neq \emptyset\}$ and $x(\mathbf{s}) = a_{\tau(\mathbf{s})}$, with $\tau(\mathbf{s}) = \infty$ and $x(\mathbf{s}) = \emptyset$ if $a_t = \emptyset$ for all t. While the question of which task the person completes and when she completes that task are of obvious interest, we shall often focus only on whether the person *ever* completes *any* task. Hence, the strategy $\mathbf{s}^\emptyset \equiv (\emptyset, \emptyset, ..., \emptyset, ...)$ plays a prominent role in our analysis.

Our solution concept, "perception-perfect strategies", requires that at all times a person have reasonable beliefs about how she would behave in the future following any possible current action, and that she choose her current action to maximize her current preferences given these beliefs. Perception-perfect strategies depend on the three attributes of a person introduced in Section 2 — her self-control problem, β, her perceptions of future self-control problems, $\hat{\beta}$, and her time-consistent discounting, δ — which jointly determine her current preferences and beliefs about future behavior.

A person's current preferences are given by the present-biased preferences with parameters β and δ defined in Section 2. Denoting the period-t intertemporal utility function given strategy $\mathbf{s}$ by $U^t(\mathbf{s}, \beta, \delta)$, this means:

$$U^t(\mathbf{s}, \beta, \delta) = \begin{cases} -c + \frac{\beta\delta}{1-\delta} v & \text{if } x(\mathbf{s}) = (c, v) \text{ and } \tau(\mathbf{s}) = t \\ \beta\delta^{\tau(\mathbf{s})-t}\left(-c + \frac{\delta}{1-\delta} v\right) & \text{if } x(\mathbf{s}) = (c, v) \text{ and } \tau(\mathbf{s}) > t \\ 0 & \text{if } x(\mathbf{s}) = \emptyset \\ v + \frac{\beta\delta}{1-\delta} v & \text{if } x(\mathbf{s}) = (c, v) \text{ and } \tau(\mathbf{s}) < t. \end{cases}$$

9 Defining strategies to be independent of history is not restrictive in our model because a person's choice in period τ matters only for the history where $a_t = \emptyset$ for all $t < \tau$. Our definitions below also rule out "mixed strategies"; it is perhaps best to interpret our analysis as applying to "equilibrium" strategies for an infinite horizon that correspond to some "equilibrium" strategies for a long, finite horizon, which (generically) do not involve mixed strategies.

The four cases in this equation correspond to the four different possibilities of when, relative to period t, the person completes the task. In the first case, the person completes task (c, v) now, and therefore she does not discount the immediate cost c by β, whereas she does discount the delayed reward $\frac{\delta}{1-\delta}v$ by β. In the second case, the person completes task (c, v) in the future, and therefore she discounts both the cost and reward by β. In the third case, the person never completes the task, and therefore her payoff is zero. In the fourth case, the person has completed task (c, v) prior to period t and, therefore, only experiences its stream of benefits. The fourth case is not relevant for behavior, but is relevant for welfare analysis.

The following terminology will prove useful in distinguishing between not doing any task merely because no task is worthwhile and "procrastinating":

Definition 1 Given β and δ, a task (c, v) is **β-worthwhile** if $\frac{\beta\delta}{1-\delta}v - c \geq 0$; and given X, the **β-best task** in X is $x^*(\beta, \delta, X) \equiv \arg\max_{(c,v)\in X} \left[\frac{\beta\delta}{1-\delta}v - c\right]$.[10]

A task is β-worthwhile if a person prefers doing it now to never doing anything *given* her taste for immediate gratification. Similarly, the β-best task is the task the person would most like to complete now *given* her taste for immediate gratification. Our somewhat conservative definition of "worthwhile" helps us to emphasize the severity of procrastination that our model predicts. Indeed, we define "procrastination" to mean never completing a task when there exists some task that the person would like to complete despite her taste for immediate gratification.[11]

Definition 2 A person **procrastinates** if she follows strategy $\mathbf{s}^{\emptyset}$ when there exists $x \in X$ that is β-worthwhile.

[10] Because the set $\mathbf{B} \equiv \arg\max_{(c,v)\in X} \left[\frac{\beta\delta}{1-\delta}v - c\right]$ need not be a singleton, $x^*(\beta, \delta, X)$ is not necessarily well-defined. If $\mathbf{B}$ is not a singleton, we define $x^*(\beta, \delta, X)$ to be the task $(c^*, v^*) \in \mathbf{B}$ such that $v^* = \max\{v|(c, v) \in \mathbf{B}\}$ — that is, the task in $\mathbf{B}$ with the largest reward (and therefore the largest cost). If either $\mathbf{B}$ is empty or $\max\{v|(c, v) \in \mathbf{B}\}$ does not exist, then the β-best task does not exist. For a given $(\beta, \hat{\beta}, \delta)$ combination, there exists a prediction consistent with the solution concept we define below if and only if the β-best task and the $\hat{\beta}$-best task both exist. If the menu of tasks X is finite, existence is guaranteed. If X is infinite, then letting $\bar{v}(c) = \max_{c' \leq c}\{v|(c', v) \in X\}$ be the maximal benefit that can be achieved for cost c or lower, our model makes a prediction if $\bar{v}(c)$ is defined for all c (i.e., the person cannot achieve an infinite reward for a finite cost) and $\lim_{c\to\infty} \frac{\bar{v}(c)}{c} < \frac{1-\delta}{\beta\delta}$.

[11] Our conservative definition of procrastination here differs from the definition in O'Donoghue and Rabin (forthcoming *a*, forthcoming *b*). As we show below, this conservative notion of procrastination means that some degree of naivete is necessary for procrastination. Since a person always does the task immediately if she thinks the alternative is never doing it, full sophistication eliminates the possibility that a person never completes a β-worthwhile task.

Let $\hat{\mathbf{s}}^t \equiv (\hat{a}^t_{t+1}, \hat{a}^t_{t+2}, ...)$ be a vector of period-t beliefs about future behavior, where $\hat{a}^t_\tau$ represents the action the person *believes* in period t that she would choose in period τ if she were to enter period τ not yet having completed a task. Given a person's beliefs $\hat{\mathbf{s}}^t$, we define $V^t(a_t, \hat{\mathbf{s}}^t, \beta, \delta)$ as the person's *perceived* period-t continuation utility conditional on having chosen $a_\tau = \emptyset$ for all $\tau < t$, choosing action a_t in period t, and following strategy $\hat{\mathbf{s}}^t$ beginning in period $t+1$. Then:

$$V^t(a_t, \hat{\mathbf{s}}^t, \beta, \delta) \equiv \begin{cases} -c + \frac{\beta\delta}{1-\delta}v & \text{if } a_t = (c,v) \\ \delta^\tau[-\beta c + \frac{\beta\delta}{1-\delta}v] & \text{if } a_t = \emptyset,\ \tau \equiv \min\{d > 0 \mid \hat{a}^t_{t+d} \neq \emptyset\} \text{ exists, and } \hat{a}^t_{t+\tau} = (c,v) \\ 0 & \text{if } a_t = \emptyset \text{ and } \hat{a}^t_{t+d} = \emptyset \text{ for all } d > 0. \end{cases}$$

With this notation, a person in period t chooses her current action a_t to maximize her current preferences V^t given her beliefs $\hat{\mathbf{s}}^t$.

To predict behavior in our model, however, we do not allow arbitrary beliefs. Rather, a person's beliefs should be a function of her perception of her future self-control problems, $\hat{\beta}$, in conjunction with some coherent theory of how she will behave given such self-control problems. We require beliefs to be dynamically consistent:

Definition 3 Given $\hat{\beta} \leq 1$ and δ, a set of beliefs $\{\hat{\mathbf{s}}^1, \hat{\mathbf{s}}^2, ...\}$ is **dynamically consistent** if
(i) For all $\hat{\mathbf{s}}^t$, $\hat{a}^t_\tau = \arg\max_{a \in A} V^\tau(a, \hat{\mathbf{s}}^t, \hat{\beta}, \delta)$ for all τ, and
(ii) For all $\hat{\mathbf{s}}^t$ and $\hat{\mathbf{s}}^{t'}$ with $t < t'$, $\hat{a}^t_\tau = \hat{a}^{t'}_\tau$ for all $\tau > t'$.

Definition 3 incorporates two aspects of dynamic consistency. First, each period's beliefs must be *internally consistent*: The beliefs must consist of a behavior path such that each period's action is optimal given that the person will stick to that behavior path in the future. Internal consistency implies that at all times the person perceives that in all future periods she will have "rational expectations" about her own behavior even further in the future. Second, the set of beliefs must be *externally consistent*: A person's beliefs must be consistent across periods, which means that a person's belief of what she will do in period τ must be the same in all $t < \tau$. This restriction rules out procrastination arising from a form of irrational expectations that goes beyond merely mispredicting self-control. For example, if in period 1 a person decides to delay based on a belief that she will complete a task in period 2 in order to avoid procrastination in period 3, then we do not allow this person to delay in period 2 based on a new belief that she will complete a task in period 3.[12]

Once we assume external consistency on beliefs, we can simplify our notation: Given $\hat{\beta}$ and δ, any set of dynamically consistent beliefs can be represented by a single vector of period-1 beliefs $\hat{\mathbf{s}}(\hat{\beta}, \delta) = (\hat{a}_2(\hat{\beta}, \delta), \hat{a}_3(\hat{\beta}, \delta), ...)$, because for all $t > 1$ external consistency requires that period-t beliefs be $\hat{\mathbf{s}}^t(\hat{\beta}, \delta) = (\hat{a}_{t+1}(\hat{\beta}, \delta), \hat{a}_{t+2}(\hat{\beta}, \delta), ...)$.

A perception-perfect strategy is a set of plans where in each period the person chooses an action to maximize her current preferences given dynamically consistent beliefs about future behavior:

Definition 4 A **perception-perfect strategy** for a $(\beta, \hat{\beta}, \delta)$ agent is $\mathbf{s}^{pp}(\beta, \hat{\beta}, \delta) \equiv (a_1(\beta, \hat{\beta}, \delta), a_2(\beta, \hat{\beta}, \delta), ...)$ such that there exists dynamically consistent beliefs $\hat{\mathbf{s}}(\hat{\beta}, \delta)$ where $a_t(\beta, \hat{\beta}, \delta) = \arg\max_a V^t(a, \hat{\mathbf{s}}(\hat{\beta}, \delta), \beta, \delta)$ for all t.

This definition encompasses as special cases the three cases previously studied in the literature: time consistency, sophisticated time inconsistency, and naive time inconsistency. A person with time-consistent preferences is characterized by $\hat{\beta} = \beta = 1$; for such a person a perception-perfect strategy maximizes long-run utility. That is, the unique perception-perfect strategy is $\mathbf{s}^* \equiv \arg\max_{\mathbf{s}} U^0(\mathbf{s})$. A completely sophisticated person is characterized by $\hat{\beta} = \beta < 1$; for such a person a perception-perfect strategy is identical to its corresponding dynamically consistent beliefs. A completely naive person is characterized by $\hat{\beta} = 1 > \beta$; for such a person the unique set of dynamically-consistent beliefs is $\mathbf{s}^*$, so at all times a completely naive person believes she will behave like a time-consistent person in the future. Definition 4 generalizes the previous literature by allowing for partial naivete.

To provide some intuition for our general results, we now carry out an extended analysis of a simple example where there is a singleton task menu; we shall return to variants of this example repeatedly. Suppose the menu of tasks is $X \equiv \{(c = 20, v = 10)\}$ and the discount parameters are $\beta = .5$ and $\delta = .9$. The values for the period-t intertemporal utility from completing task x in period $\tau \geq t$, which we denote by $\tilde{V}^t(\tau, x, \beta, \delta)$, are calculated in this example as follows:

[12] The restrictions imposed by external consistency essentially correspond to the additional restrictions which subgame-perfect equilibrium imposes beyond non-equilibrium backwards induction. By the same token, these restrictions would be unnecessary in generic, finite-period situations where "perceptual backwards induction" would yield a unique prediction. Previous analyses of time-inconsistent preferences, whether examining sophistication or complete naivete, have implicitly assumed that people have beliefs that are both internally and externally consistent. Without the assumption of external consistency, even a person who knows exactly her future self-control problems could fail to exhibit "rational expectations".

$$\begin{array}{lcccc}
\tilde{V}^t(\tau = t, x, \beta, \delta) & = & -20 + (.5)\frac{.9}{1-.9}10 & = & 25.00 \\
\tilde{V}^t(\tau = t+1, x, \beta, \delta) & = & (.5)(.9)\left[-20 + \frac{.9}{1-.9}10\right] & = & 31.50 \\
\tilde{V}^t(\tau = t+2, x, \beta, \delta) & = & (.5)(.9)^2\left[-20 + \frac{.9}{1-.9}10\right] & = & 28.35 \\
\tilde{V}^t(\tau = t+3, x, \beta, \delta) & = & (.5)(.9)^3\left[-20 + \frac{.9}{1-.9}10\right] & = & 25.52 \\
\tilde{V}^t(\tau = t+4, x, \beta, \delta) & = & (.5)(.9)^4\left[-20 + \frac{.9}{1-.9}10\right] & = & 22.96.
\end{array}$$

These preferences are prototypical of those that drive delay in our model. Since in our simple formulation present bias is a one-period phenomenon, if a person wants to complete a task, then the optimal time to do so, taking into account her taste for immediate gratification, is either today or tomorrow — putting off the task beyond tomorrow does nothing to satisfy her current taste for immediate gratification. If the optimal time is today, then the person clearly won't delay. But if, as in this example, the optimal time is tomorrow, then the person might delay. Whether she does so depends on both her payoffs in subsequent periods and her perception of future behavior. If the optimal time to complete the task is tomorrow, the person may also prefer other future dates to today. Indeed, in the example, the person prefers completing the task in two days or in three days to today. But if the person wants to complete the task at all, there is some maximum tolerable delay d^* such that for any $d \geq d^* + 1$ completing the task today is preferred to completing the task in d periods. In the example, $d^* = 3$. When there is a singleton task menu, the higher is β, the lower is the tolerance for delay. If, for instance, β were increased from .5 to .52 in the example above, then the maximum tolerable delay would decrease to $d^* = 2$.[13]

Now consider the role of the person's beliefs about future behavior. Suppose first that the person is completely sophisticated — she has beliefs $\hat{\beta} = \beta$ — and therefore accurately predicts her future behavior. Since in period t she completes the task if and only if she predicts that she would delay more than d^* periods, her perception-perfect strategy must in all periods involve planning to complete the task if and only if not doing so would cause at least $d^* + 1$ periods of further delay. In the $\beta = .5$, $\delta = .9$, $X = \{x = (c = 20, v = 10)\}$ example above, if $\hat{\beta} = \beta$ there are four perception-perfect strategies — $(x, \emptyset, \emptyset, \emptyset, x, \emptyset, \emptyset, \emptyset, x, ...)$, $(\emptyset, x, \emptyset, \emptyset, \emptyset, x, \emptyset, \emptyset, \emptyset, ...)$, $(\emptyset, \emptyset, x, \emptyset, \emptyset, \emptyset, x, \emptyset, \emptyset, ...)$, and $(\emptyset, \emptyset, \emptyset, x, \emptyset, \emptyset, \emptyset, x, \emptyset, ...)$ — and there are four perception-perfect outcomes — completing the task on the 1st, 2nd, 3rd, or 4th day. Similarly, with $\beta = .52$, a com-

[13] The relevant calculations are: $\tilde{V}^t(\tau = t, x, \beta, \delta) = -20 + (.52)\frac{.9}{1-.9}10 = 26.80$; $\tilde{V}^t(\tau = t+1, x, \beta, \delta) = (.52)(.9)\left[-20 + \frac{.9}{1-.9}10\right] = 32.76$; $\tilde{V}^t(\tau = t+2, x, \beta, \delta) = (.52)(.9)^2\left[-20 + \frac{.9}{1-.9}10\right] = 29.48$; and $\tilde{V}^t(\tau = t+3, x, \beta, \delta) = (.52)(.9)^3\left[-20 + \frac{.9}{1-.9}10\right] = 26.54$.

pletely sophisticated person will have three perception-perfect strategies, with corresponding outcomes of completing the task on the 1^{st}, 2^{nd}, or 3^{rd} day.[14]

Now consider a partial naif with $\beta = .5$ who has beliefs $\hat{\beta} = .52$. Because she perceives that she will behave in the future like a completely sophisticated person with a self-control problem of $\beta = .52$, and because such a sophisticate would tolerate a delay at most 2 days, a person who perceives a future self control of $\hat{\beta} = .52$ believes the most she will delay if she doesn't do the task now is 3 days. But since a person with $\beta = .5$ always prefers doing the task 3 days from now to doing it now, a partial naif with $\beta = .5$ and $\hat{\beta} = .52$ will procrastinate forever!

This example illustrates the basic logic that determines whether a person procrastinates: Even when a task is worthwhile, a person will delay indefinitely whenever she perceives that her future tolerance for delay will be at least one period less than her current (and actual future) tolerance for delay. Since for a completely sophisticated person the perceived future tolerance for delay is identical to the current tolerance for delay, a completely sophisticated person will never procrastinate. But, as this example illustrates, it needn't take much naivete to generate procrastination.

While there is only one task in this example, similar logic determines procrastination when there is a menu of tasks from which to choose. The main complication is that the person must choose which task to consider completing now, and she must predict which task she would complete in the future if she waits now. But clearly the person only considers completing the β-best task $x^*(\beta, \delta, X)$ now, and she perceives that in the future she will only consider completing the $\hat{\beta}$-best task $x^*(\hat{\beta}, \delta, X)$, and therefore in each period the person debates completing the β-best task now versus the $\hat{\beta}$-best task in the not-too-distant future. Hence, whether a person procrastinates boils down to comparing her current tolerance for delaying the β-best task now in favor of completing the $\hat{\beta}$-best task in the future to her perceived future tolerance for delay of the $\hat{\beta}$-best task.

Let $d(\beta|\hat{\beta})$ denote a person's current tolerance for delay as a function of her current self-control problem β and her perceived future self-control problem $\hat{\beta}$. Letting $x^*(\beta, \delta, X) = (c^*, v^*)$ and

14 Although we have ruled out "mixed strategies", it is worth noting what they look like. Whenever a completely sophisticated person prefers completing the β-best task tomorrow rather than today, there exists a mixed perception-perfect strategy when there is an infinite horizon. Let $x^*(\beta, \delta, X) \equiv (c^*, v^*)$, and let p satisfy

$$\frac{\beta\delta}{1-\delta}v^* - c^* = \beta\delta\left[\sum_{\tau=1}^{\infty}(1-p)^{\tau-1}p\delta^{\tau-1}\left(\frac{\delta}{1-\delta}v^* - c^*\right)\right] = \frac{\beta\delta p}{1-(1-p)\delta}\left(\frac{\delta}{1-\delta}v^* - c^*\right).$$

It is straightforward to show that preferring to complete the β-best task tomorrow rather than today implies that there exists a unique $p \in (0,1)$ that satisfies this condition, in which case it is a perception-perfect strategy to complete the β-best strategy with probability p in all periods. We conjecture but have not proven that this is the only mixed-strategy equilibrium for a completely sophisticated person.

$x^*(\hat{\beta}, \delta, X) = (\overline{c}, \overline{v})$, the current tolerance for delay is given by:

$$d(\beta|\hat{\beta}) \equiv \max \left\{ d \in \{0, 1, ...\} \;\middle|\; -c^* + \frac{\beta\delta}{1-\delta} v^* < \beta\delta^d \left(-\bar{c} + \frac{\delta}{1-\delta} \bar{v} \right) \right\}.$$

In words, $d(\beta|\hat{\beta})$ is the maximum delay such that a person with self-control problem β prefers doing the $\hat{\beta}$-best task in $d(\beta|\hat{\beta})$ periods over doing the β-best task now. It is worth noting that $\hat{\beta}$ affects $d(\beta|\hat{\beta})$ only insofar as $\hat{\beta}$ determines perceptions of which task will be chosen in the future, not when it will be chosen. Hence, if there is only one task, $d(\beta|\hat{\beta})$ is independent of $\hat{\beta}$.

Using the notation above, a person with beliefs $\hat{\beta}$ perceives that her future tolerance for delay is exactly $d(\hat{\beta}|\hat{\beta})$. Since $\hat{\beta} \geq \beta$, and since the more patient a person is the less tolerant she is of delay of a given task, it must be that $d(\hat{\beta}|\hat{\beta}) \leq d(\beta|\hat{\beta})$. Hence, whether a person procrastinates boils down to the following:

If $d(\hat{\beta}|\hat{\beta}) \;=\; d(\beta|\hat{\beta})$, the person will eventually complete a task.

If $d(\hat{\beta}|\hat{\beta}) \;\leq\; d(\beta|\hat{\beta}) - 1$, the person delays indefinitely.

We now formalize this logic. Lemma 1 formally characterizes the set of dynamically consistent beliefs:[15]

Lemma 1 For $\hat{\beta}$, δ, and X, any dynamically consistent beliefs $\hat{\mathbf{s}}(\hat{\beta}, \delta) \equiv (\hat{a}_2(\hat{\beta}, \delta), \hat{a}_3(\hat{\beta}, \delta), ...)$ must satisfy:

(1) For all t either $\hat{a}_t(\hat{\beta}, \delta) = \emptyset$ or $\hat{a}_t(\hat{\beta}, \delta) = x^*(\hat{\beta}, \delta, X)$, and

(2) If $x^*(\hat{\beta}, \delta, X)$ is not $\hat{\beta}$-worthwhile, then $\hat{a}_t(\hat{\beta}, \delta) = \emptyset$ for all t. Otherwise there exists $\tau \in \{2, 3, ..., d(\hat{\beta}|\hat{\beta}) + 2\}$ such that $\hat{a}_t(\hat{\beta}, \delta) = x^*(\hat{\beta}, \delta, X)$ if and only if $t \in \{\tau, \tau + (d(\hat{\beta}|\hat{\beta}) + 1), \tau + 2(d(\hat{\beta}|\hat{\beta}) + 1), ...\}$.

Lemma 1 states that the only task a person would ever expect to complete in the future is the $\hat{\beta}$-best task, and moreover that the person will expect to complete the $\hat{\beta}$-best task every $d(\hat{\beta}|\hat{\beta}) + 1$ periods. In other words, any dynamically consistent beliefs must be "cyclical", and whenever some task is $\hat{\beta}$-worthwhile the length of the cycle is finite. Notice, however, that whenever $d(\hat{\beta}|\hat{\beta}) > 0$,

[15] All proofs are in the Appendix.

the first date of completion is indeterminate, and therefore there are multiple dynamically consistent beliefs.[16]

Given that there can be multiple dynamically consistent beliefs, there can be multiple perception-perfect strategies. Many of our results will state properties of the entire set of perception-perfect strategies, which we denote by $S^{pp}(\beta, \hat{\beta}, \delta, X)$. Lemma 2 characterizes $S^{pp}(\beta, \hat{\beta}, \delta, X)$:

Lemma 2 For all β, $\hat{\beta}$, δ, and X, either $S^{pp}(\beta, \hat{\beta}, \delta, X) = \{\mathbf{s}^{\emptyset}\}$, or for every $\mathbf{s} \in S^{pp}(\beta, \hat{\beta}, \delta, X)$, $x(\mathbf{s}) = x^*(\beta, \delta, X)$, $\tau(\mathbf{s}) \leq d(\hat{\beta}|\hat{\beta}) + 1$, and if $\tau(\mathbf{s}) > 1$ then $\tau(\mathbf{s}) = \tau(\hat{\mathbf{s}})$ where $\hat{\mathbf{s}}$ is the corresponding set of dynamically consistent beliefs.

Lemma 2 establishes that either there is a unique perception-perfect strategy under which the person never completes a task, or in every perception-perfect strategy the person eventually completes task $x^*(\beta, \delta, X)$. That is, for given parameter values there can be indeterminacy solely in *when* the person completes a task, and not in either whether she completes a task, or which task she completes. The intuition for determinacy in whether a person will (eventually) complete a task should be clear from our earlier example. If $d(\hat{\beta}|\hat{\beta}) + 1 \leq d(\beta|\hat{\beta})$, then in all periods the maximum future delay that a person could possibly perceive is tolerable, and hence the person must delay in all periods. If $d(\hat{\beta}|\hat{\beta}) + 1 > d(\beta|\hat{\beta})$, in contrast, then in some period the person must perceive an intolerable delay, and she will therefore complete the task in that period. In the latter case, a multiplicity of perception-perfect strategies can arise because the period of completion depends on the specific dynamically consistent beliefs the person holds, which determine the first period in which she perceives an intolerable delay from waiting. The final part of Lemma 2 establishes that if the person delays but eventually completes a task, then she correctly predicts the period in which she will complete the task (although she incorrectly predicts which task she will complete in that period whenever the $\hat{\beta}$-best task differs from the β-best task).

Because the β-best task $x^*(\beta, \delta, X)$ does not depend on $\hat{\beta}$, an important implication of Lemma 2 is that while a person's awareness of her self-control is be a crucial determinant of whether or not

[16] The multiplicity does not arise from the type of reward-and-punishment supergame strategies that can be found in infinite-horizon models of time-inconsistent preferences; the strategies in the infinite-period model correspond to the set of strategies that are the limit of the strategies in the finite-period model as the number of periods becomes arbitrarily large, where (generically) each finite-period situation will have a unique perception-perfect strategy. The multiplicity in the limit comes from the fact that each perception-perfect strategy is "cyclical", so that (say) a person will plan to do the task the last period if not before, and the the fourth-to-last if not before, the seventh-to-last if not before, etc., and not do the task in other periods. Such strategies will therefore predict that in a 1008-period model the person does the task in period 2, but in a 1007-period model she does it in period 1.

she procrastinates, it plays no role in determining which task she does if she doesn't procrastinate. If and when a person decides to complete a task, her beliefs about what she would do in the future if she didn't complete the task are irrelevant — she chooses the task that currently seems best to her, $x^*(\beta, \delta, X)$.

Many of the intuitions we provide, and some of the sufficient conditions for procrastination we derive in the formal results below, simplify the person's thought process slightly: We imagine that the agent believes that she will choose the β-best task in the future rather than the $\hat{\beta}$-best task. This intuition is fully correct when there is only one task, in which case the $\hat{\beta}$-best and β-best task must of course be the same. When there are multiple tasks, the two may be different; however, because the $\hat{\beta}$-best task is better and therefore can only increase the person's maximum tolerable delay, if a person procrastinates believing she will do the β-best task in the future, she procrastinates given her actual beliefs.

Lemma 2 establishes that the parameters of the model fully determine whether or not the person does a task. Proposition 1 characterizes how whether a person ever does a task depends on the degree of sophistication $\hat{\beta}$:

Proposition 1 For all β, δ, and X:
(1) If no $x \in X$ is β-worthwhile, then $S^{pp}(\beta, \hat{\beta}, \delta, X) = \{\mathbf{s}^{\emptyset}\}$ for all $\hat{\beta}$, and
(2) If there exists $x \in X$ that is β-worthwhile, then either a) $S^{pp}(\beta, \hat{\beta}, \delta, X) \neq \{\mathbf{s}^{\emptyset}\}$ for all $\hat{\beta}$, or b) "generically" there exist β^* and β^{**} satisfying $\beta < \beta^* \leq \beta^{**} < 1$ such that $S^{pp}(\beta, \hat{\beta}, \delta, X) \neq \{\mathbf{s}^{\emptyset}\}$ for any $\hat{\beta} < \beta^*$ and $S^{pp}(\beta, \hat{\beta}, \delta, X) = \{\mathbf{s}^{\emptyset}\}$ for any $\hat{\beta} > \beta^{**}$.[17]

Part (1) of Proposition 1 merely states that if no task is β-worthwhile then the person doesn't complete a task regardless of her perceptions. Part (2) considers the role that naivete plays in procrastination in the more interesting case where some task is β-worthwhile. Because it determines her degree of over-optimism about how soon she will complete the task if she waits today, the degree to which a person is naive is an important determinant of whether she chooses to delay. If a person is sophisticated or nearly sophisticated, then she does not procrastinate. Intuitively, if a task is β-worthwhile, then in all periods the person prefers completing that task immediately to never completing any task. Since a sophisticated person correctly predicts future behavior, she cannot delay indefinitely because if she perceives that she will do so, then she completes the task

[17] The caveat "generically" is required for the result that $\beta^* > \beta$, which holds if we rule out knife-edge parameters where $x^*(\beta, \delta, X) \equiv (c, v)$ and $v - \frac{1}{\beta}c = \delta^{d(\beta|\beta)+1}\left[\frac{\delta}{1-\delta}v - c\right]$. In such cases, it could be that $\beta^* = \beta$.

immediately. If a person is nearly sophisticated, then her beliefs $\hat{\mathbf{s}}(\hat{\beta}, \delta)$ are nearly identical to a sophisticate's beliefs $\hat{\mathbf{s}}(\beta, \delta)$, and therefore she completes the task when the sophisticate does so. If a person is more naive, on the other hand, she may well be persistently optimistic enough that her taste for immediate gratification always induces her to delay.

In a slightly different framework, O'Donoghue and Rabin (forthcoming *a*) show that when there is a single task available, a pure sophisticate (i.e., $\hat{\beta} = \beta$) cannot procrastinate whereas a pure naif (i.e., $\hat{\beta} = 1$) might.[18] Proposition 1 generalizes this result in two ways: It replicates it for the case where there is a menu of tasks available, and establishes that partial naifs behave "in between" pure naifs and pure sophisticates. In terms of the second extension, Proposition 1 establishes a sort of continuity property: For a given environment, a person who is nearly completely sophisticated behaves like a person who is completely sophisticated, and a person who is nearly completely naive behaves like a person who is completely naive.[19]

Although Proposition 1 shows that in a *given* environment a person who is nearly sophisticated does not procrastinate, Proposition 2 establishes that for any departure from pure sophistication, there exists an environment where the person procrastinates:

Proposition 2 For all β, δ, and $\hat{\beta} > \beta$, there exists X such that the person procrastinates.

That is, any degree of naivete is sufficient to induce procrastination. A person procrastinates whenever she believes her future tolerance for delay will be one or more periods less than her current tolerance. If there is only one task, for instance, and the person barely prefers doing it tomorrow rather than today, then even for $\hat{\beta}$ very close to β the person perceives every day that she

---

[18] This statement is true using the definition of procrastination in this paper, but not for the slightly different definition in O'Donoghue and Rabin (forthcoming *a*).

[19] Because changing $\hat{\beta}$ can change the $\hat{\beta}$-best task, $d(\hat{\beta}|\hat{\beta})$ can be non-monotonic in $\hat{\beta}$, and therefore it is not necessarily the case that if the person procrastinates for some $\hat{\beta}$ then she must procrastinate for all $\hat{\beta}' > \hat{\beta}$. When there is a single task available, the $\hat{\beta}$-best task is necessarily independent of $\hat{\beta}$, so that if the person procrastinates for some $\hat{\beta}$ she procrastinates for all $\hat{\beta}' > \hat{\beta}$.

will do the task tomorrow, and thus she procrastinates forever.[20]

The terms we have used to describe our results — that people *procrastinate* on a task that is β-*worthwhile* — connote that this behavior is harmful to the person. To see why these terms might be appropriate, we now turn to formal welfare analysis. The meaning of the statement that somebody with time-inconsistent preferences is "hurting herself" has sometimes troubled researchers, since time-inconsistent preferences imply that a person evaluates her well-being differently at different times. Some researchers (e.g., Goldman (1979), Laibson (1994, 1995, 1997)) have avoided this problem by using a "Pareto criterion", under which one intertemporal stream of utilities is considered unambiguously better for a person than another only if it is preferred by the person from *all* time perspectives:

Definition 5 A strategy $\mathbf{s}$ is *Pareto-efficient* if there does not exist an alternative strategy $\mathbf{s}'$ such that $U^t(\mathbf{s}', \beta, \delta) \geq U^t(\mathbf{s}, \beta, \delta)$ for all t and $U^t(\mathbf{s}', \beta, \delta) > U^t(\mathbf{s}, \beta, \delta)$ for some t.

We show below that a person's behavior is Pareto-inefficient if it meets our definition of procrastination. But we shall also judge welfare by a second criterion that allows us to evaluate not just whether a person is hurting herself, but also how severely she is doing so.[21] A person's preferences from a long-run perspective are:

[20] Consider the implications for procrastination of allowing mixed strategies. "Generically", a person cannot be indifferent between doing the β-best task now and doing the $\hat{\beta}$-best task in some future period. Hence, a person can mix only if she has mixed beliefs. But our earlier discussion implies that if $d(\hat{\beta}|\hat{\beta}) > 0$, mixed beliefs indeed exist. However, according to these beliefs the (long-run) continuation payoff beginning next period must be just sufficient to make a person with self-control problem $\hat{\beta}$ indifferent between doing a task now versus waiting. But this means that a person with self-control problem $\beta < \hat{\beta}$ will strictly prefer to wait. Hence, we can conclude that whenever $d(\hat{\beta}|\hat{\beta}) > 0$, (under a more general definition) there exists a perception-perfect strategy based on mixed beliefs wherein the person procrastinates.

We do not focus on such strategies because they make the analysis somewhat trivial — e.g., procrastination for any $\hat{\beta} > \beta$ — and we don't feel that they are particularly realistic. We also remind the reader that such strategies are ruled out by a long, finite horizon.

[21] More generally, we feel that the Pareto criterion is too conservative an approach to intrapersonal welfare analysis. Just as for interpersonal comparisons where the Pareto criterion refuses to call a reallocation that barely hurts one person and enormously helps everyone else an improvement, the Pareto criterion refuses to rank strategies where one perspective barely prefers one strategy and all other perspectives vastly prefer a second strategy. Indeed, in this environment the Pareto criterion insists we never label as bad the act of immediately doing the task that has the highest *gross* benefit, regardless of its cost. For example, suppose there are two tasks, x_1, with $c_1 = 0$ and $v_1 = 1$, and x_2, with $c_2 = 1,000,000,000,000$ and $v_2 = 1.01$. Unless δ is very close to 1, doing task x_1 immediately is clearly better than doing task x_2 immediately; and yet for even very small values of δ, doing task x_2 immediately is not "Pareto-dominated" by x_1. Furthermore, the Pareto criterion's unwillingness to designate x_2 as inefficient does not depend on β, *and holds even if* $\beta = 1$. Hence, applying Pareto-agnosticism when evaluating time-*consistent* agents would mean that we would be agnostic about whether forcing a person to take action x_2 is a bad thing.

$$U^{LR}(\mathbf{s},\delta) \equiv \begin{cases} \delta^{\tau(\mathbf{s})-1}\left(-c+\frac{\delta}{1-\delta}v\right) & \text{if } x(\mathbf{s})=(c,v) \\ 0 & \text{if } x(\mathbf{s})=\emptyset. \end{cases}$$

Using this formulation, we define a person's "welfare loss" as the (normalized) difference between her actual long-run utility and her best possible long-run utility. We normalize the difference by dividing by c^*, the cost of the long-run-best task, so that the welfare loss does not depend arbitrarily on the unit used to measure costs and benefits.[22]

Definition 6 Let $U^* \equiv \max_{\bar{\mathbf{s}}} U^{LR}(\bar{\mathbf{s}},\delta)$ and let c^* be the cost of the task chosen (immediately) to maximize $U^{LR}(\bar{\mathbf{s}},\delta)$. If a person follows strategy $\mathbf{s}$, then her *welfare loss* is

$$WL(\mathbf{s},\delta) \equiv \frac{U^* - U^{LR}(\mathbf{s},\delta)}{c^*}.$$

In what follows, we shall say that a welfare loss of $WL(\mathbf{s},\delta) < \frac{1-\beta}{\beta}$ is "small". A welfare loss of $\frac{1-\beta}{\beta}$ corresponds to the maximum possible welfare loss from a single episode of pursuing immediate gratification. For instance, it is the maximum welfare loss a person can suffer when she does the β-best task rather than the long-run-best task in period 1, and it is the maximum welfare loss a person could suffer if she were hypothetically allowed to commit in period 1 to her most preferred lifetime behavior path. Our focus is on the more dramatic examples of harmful procrastination where a person repeatedly chooses to pursue immediate gratification rather than long-run welfare, in which case she can suffer welfare losses significantly larger than $\frac{1-\beta}{\beta}$.[23]

The following proposition characterizes how a person can hurt herself according to our two welfare criteria:

[22] Predicted behavior is homogenous of degree one — that is, if we were to double all rewards and costs, the set of perception-perfect strategies would not change.

[23] O'Donoghue and Rabin (forthcoming *a*) present limit results showing that a sophisticated person with a very small self-control problem (i.e., β close to 1) will not severely hurt herself, whereas a completely naive person can suffer arbitrarily severe welfare losses. Notice that a person can suffer a substantial welfare loss for small self-control problems (i.e., β close to 1) only if $WL(\mathbf{s}) > \frac{1-\beta}{\beta}$.

Proposition 3 For all β, $\hat{\beta}$, δ, and X:
(1) If $S^{pp}(\beta,\hat{\beta},\delta,X)=\{\mathbf{s}^{\emptyset}\}$, then
(a) $\mathbf{s}^{\emptyset}$ is Pareto-inefficient if and only if it reflects procrastination; and
(b) $WL(\mathbf{s}^{\emptyset})>\frac{1-\beta}{\beta}$ only if it reflects procrastination.
(2) If $S^{pp}(\beta,\hat{\beta},\delta,X)\neq\{\mathbf{s}^{\emptyset}\}$, then
(a) There exists $\mathbf{s}\in S^{pp}(\beta,\hat{\beta},\delta,X)$ that is Pareto-efficient and has $WL(\mathbf{s})<\frac{1-\beta}{\beta}$,
(b) Any $\mathbf{s}\in S^{pp}(\beta,\hat{\beta},\delta,X)$ is Pareto-inefficient if and only if $\tau(\mathbf{s})>d(\beta|\beta)+1$; and
(c) For any $\mathbf{s}\in S^{pp}(\beta,\hat{\beta},\delta,X)$, $WL(\mathbf{s})>\frac{1-\beta}{\beta}$ only if $\tau(\mathbf{s})>d(\beta|\beta)+1$.

Part (1) of Proposition 3 describes whether a person hurts herself when she never completes any task. Never completing a task can Pareto-harm a person if and only if it reflects procrastination, and yields significant welfare loss only if it reflects procrastination. If a person never completes a task merely because no task is β-worthwhile, then she is following her period-1 self's most preferred path of behavior, and therefore does not hurt herself by either criterion. In contrast, if she never completes a task when some task is β-worthwhile, then *every period-self* prefers to complete the β-best task in period 1 as opposed to doing nothing. In this case, never completing a task is clearly Pareto inefficient (although it does not necessarily cause a large welfare loss).

Part (2) of Proposition 3 describes whether a person hurts herself in cases where she does eventually complete some task. Part (2a) establishes that in this case there exists at least one perception-perfect strategy under which the person does not harm herself. In particular, whenever $S^{pp}(\beta,\hat{\beta},\delta,X)\neq\{\mathbf{s}^{\emptyset}\}$, doing the β-best task in period 1 is always a perception-perfect strategy, and by definition doing the β-best task right away is not severe harm. But Parts (2b) and (2c) say that there may exist other perception-perfect strategies under which the person does harm herself because she delays longer than her tolerance for delaying the β-best task. Lemma 2 showed that the person correctly predicts when she will do the task; the possible source of harm, however, is that she incorrectly predicts she will complete the $\hat{\beta}$-best task, which can lead her to tolerate too long a delay. The condition in Parts (2b) and (2c) — $\tau(\mathbf{s})>d(\beta|\beta)+1$ — is precisely the condition that the period-1 self would have preferred to do the β-best task right away rather than delay doing the β-best task until period $\tau(\mathbf{s})$.[24]

[24] We are inclined to believe that the welfare harm from short-term delay due to a person's over-optimism about *which* task she'll do is small; it is the long-term (in the stationary model of this paper, infinite) delay due to over-optimism about when she'll do the task that we believe is the major source of harm. This intuition is admittedly rather impressionistic. To consider the relative importance of the two sources of harm, however, we pose the following question: Which "information" would be more valuable to a person, knowing when in the future she would actually do something (but not knowing what she would do), or knowing what exactly she would do (but not knowing when she would do it)? Our impression is that the former would be far more valuable.

Combining Proposition 3 with our earlier results yields conclusions about the role of naivete in causing welfare harm. Since a completely sophisticated person never procrastinates, and also correctly predicts which task she would in the future, Proposition 3 implies that a completely sophisticated person never severely hurts herself. But since Proposition 2 implies that anyone not completely sophisticated can procrastinate, Proposition 3 also implies that anyone not completely sophisticated can behave Pareto inefficiently. It does not, of course, say that a partially naive person always severely hurts herself by the long-run-utility criterion whenever she procrastinates. But the following Proposition shows that for any partially naive person there is no upper bound on how severely that person can hurt herself by procrastinating:

Proposition 4 For any β and any $\hat{\beta} > \beta$:
(1) For any δ, there exists X such that $S^{pp}(\beta, \hat{\beta}, \delta, X) = \{\mathbf{s}^{\emptyset}\}$ and $WL\left(\mathbf{s}^{\emptyset}\right) > \frac{1-\beta}{\beta}$, and
(2) For any $Z > 0$, there exist X and δ such that $S^{pp}(\beta, \hat{\beta}, \delta, X) = \{\mathbf{s}^{\emptyset}\}$ and $WL\left(\mathbf{s}^{\emptyset}\right) > Z$.

Hence, just as the behavioral results above extend earlier results about one-task, pure-naive procrastination, so too do these results extend earlier welfare results: A person can severely harm herself only if she is, to some degree, naive.

4. Choice and Procrastination

Section 3 shows that the principles developed in Akerlof (1991) and O'Donoghue and Rabin (forthcoming *a*) in the case of only one task and extreme naivete extend to multiple tasks and partial naivete. In this section, we turn to the core new results of our paper, which illustrate aspects of procrastination that pertain specifically to the presence of more than one option.

The implications of choice for procrastination derive from the relationship between the two aspects of a person's decision: Which task to do, and when to do it. These two aspects of a person's behavior are determined by different criteria. Indeed, the fact that there is so little relationship between the two aspects of the person's decision drives many results. The person decides which task to do according to long-term net benefits, choosing the task that maximizes $\frac{\beta\delta}{1-\delta}v - c$. But once she has chosen which task to do, whether she delays doing it has little to do with the long-term net benefit. Rather, it is (primarily) determined by comparing the cost of the chosen task to its *short-run*, per-period benefit.[25] To illustrate this point, Lemma 3 examines transformations of the person's choice set under which the long-term net benefits are held constant for all tasks:

Lemma 3 Consider any β, $\hat{\beta} > \beta$, δ, and X, and let X' and δ' satisfy $(c', v') \in X'$ if and only if there exists $(c, v) \in X$ such that $\frac{\beta\delta'}{1-\delta'}v' - c' = \frac{\beta\delta}{1-\delta}v - c$.

If $x^*(\beta, \delta, X) \equiv (c^*, v^*)$ and $x^*(\beta, \delta', X') \equiv (c^{*\prime}, v^{*\prime})$, then

(1) $\frac{\beta\delta'}{1-\delta'}v^{*\prime} - c^{*\prime} = \frac{\beta\delta}{1-\delta}v^* - c^*$; and

(2) If $\frac{v^{*\prime}}{c^{*\prime}} < \frac{1-\beta\delta'/\hat{\beta}}{\beta\delta'}$ then $S^{pp}(\beta, \hat{\beta}, \delta', X') = \{\mathbf{s}^{\emptyset}\}$.

For any transformation of the choice set which holds constant the long-term net benefits for all tasks, clearly the long-term net benefit of the β-best task are unchanged. Even so, if the transformation makes the β-best task sufficiently more costly relative to its *per-period benefit*, the person surely procrastinates with the transformed choice set (regardless of how the transformation affects $\frac{v}{c}$ for any other task $(c, v) \in X$). Lemma 3 therefore implies that no matter how large the long-run net benefit of the β-best task, a condition for procrastination is merely that the β-best task involve a cost that is sufficiently large relative to its per-period benefit.

Two classes of transformations nicely illustrate Lemma 3. First, suppose that $\delta' \equiv \delta$ and $X' \equiv \left\{(c', v') \mid (c, v) \in X,\ c' = c + n,\ \frac{\beta\delta}{1-\delta}v' = \frac{\beta\delta}{1-\delta}v + n\right\}$ for some $n \in \mathbb{R}$. This transformation holds the discount factor constant, while varying the cost and per-period benefit of each task in a way that holds the long-run net benefit of each task constant. Since for such transformations increasing the cost of each task also increases the ratio of cost to per-period benefit of each task, Lemma 3 implies that the person will not procrastinate when n is small enough, but will surely procrastinate if n is large enough. Second, suppose that $X' \equiv \left\{(c', v') \mid (c', v) \in X,\ \frac{\beta\delta'}{1-\delta'}v' = \frac{\beta\delta}{1-\delta}v\right\}$ and $\delta' \equiv \delta^{1/n}$ for some $n > 0$. This transformation holds the cost of each task constant, while varying the discount factor and the per-period benefit of each task in a way that holds the long-run net benefit of each task constant. Since for such transformations increased patience implies lower per-period benefits and, therefore, an increased ratio of cost to per-period benefit for each task, Lemma 3 implies the person will not procrastinate when n is small enough, but will surely

[25] The caveat "primarily" comes from the fact that whether a person completes the chosen task does depend on another task — namely, the $\hat{\beta}$-best task that she believes she would complete in the future. But since from today's perspective completing the $\hat{\beta}$-best task in the future can only look better than completing the β-best task in the future, this effect only makes procrastination more likely than connoted by the intuition we emphasize.

procrastinate if n is large enough.[26]

Our first main finding regarding the role of choice for procrastination is that providing additional options to a person who is not procrastinating can in fact induce procrastination. For example, consider a person with $\beta = .5$, $\hat{\beta} = 1$, and $\delta = .9$. If the only task available to this person is $x_1 \equiv (c_1 = 0, v_1 = 5.5)$, the person completes task x_1 immediately (since it is costless and the only reason a person ever delays is to put off incurring a cost). But if in addition to task x_1 this person also has available task $x_2 \equiv (c_1 = 20, v_1 = 10)$, she will do nothing. When tasks x_1 and x_2 are both available, task x_2 is both the β-best task (because $\frac{\beta\delta}{1-\delta}v_2 - c_2 > \frac{\beta\delta}{1-\delta}v_1 - c_1$) and the $\hat{\beta}$-best task (because $\frac{\hat{\beta}\delta}{1-\delta}v_2 - c_2 > \frac{\hat{\beta}\delta}{1-\delta}v_1 - c_1$). In her own mind, therefore, the person's choice boils down to when to do task x_2, and the availability of task x_1 is completely irrelevant. But, as is implicit from the example in Section 3, the person will never complete task x_2. Hence, the person procrastinates when tasks x_1 and x_2 are both available, even though she would complete x_1 if it were the only task available.

Of course, this example violates one of the core axioms of rational choice and revealed-preference theory — that additional options should not change choice among existing options. The source of this violation of revealed-preference theory is the person's naive belief that she will soon do one of the additional options, when in fact she won't. The person *intends* to adhere to the weak axiom of revealed preference, but fails to follow through. Proposition 5 formalizes the role of naivete in this phenomenon:

Proposition 5 For all β, $\hat{\beta}$, δ and X such that $S^{pp}(\beta, \hat{\beta}, \delta, X) \neq \{\mathbf{s}^{\emptyset}\}$:
(1) If $\hat{\beta} = \beta$, $S^{pp}(\beta, \hat{\beta}, \delta, X') \neq \{\mathbf{s}^{\emptyset}\}$ for all $X' \supset X$; and
(2) If $\hat{\beta} > \beta$, there exists task x' such that $S^{pp}(\beta, \hat{\beta}, \delta, X \cup \{x'\}) = \{\mathbf{s}^{\emptyset}\}$.

Part (1) of Proposition 5 establishes that a fully sophisticated person cannot be induced to procrastinate by providing more options. Since a sophisticate correctly predicts future behavior, she chooses to never complete a task only if there does not exist an option worth completing. If the initial menu contains an option worth completing, adding more options cannot change this condition, and therefore can never induce procrastination. Part (2) of Proposition 5 establishes, however,

[26] This second class suggests that making the time intervals between decisions negligible could induce procrastination, and that in a continuous-choice model with no moment-to-moment variation in the cost or benefit of doing the task, any partially naive person would for sure procrastinate. The caveat to this interpretation, however, is the assumption that it only takes one period to complete a task. A continuous-choice model where tasks could not be completed instantaneously would not necessarily guarantee procrastination.

that for *any* degree of naivete and any menu of tasks on which the person does not procrastinate, there exists a task that when added to this menu induces procrastination. The logic behind this result is exactly as in the example above: To induce procrastination, we merely add an option that is β-preferred to existing options but has a high cost relative to its per-period benefit.

There are of course welfare analogues to the behavioral results of Proposition 5. Additional choices can never hurt a sophisticate.[27] And there is no limit to how much additional choices can harm a partial naif, because no matter how well off some initial set of choices makes her, adding an option can induce procrastination.

Our second main finding regarding the role of choice for procrastination is that people may procrastinate more in pursuit of important goals than in unimportant ones, or equivalently that increasing importance can exacerbate procrastination. The basic intuition is simple: The more important a goal, the higher the cost the person wishes to incur in pursuit of that goal, and people tend to procrastinate more on higher-cost tasks.

To illustrate this point, we return to a variant of our earlier examples. Suppose $\beta = .6$, $\hat{\beta} = 1$, $\delta = .8$, and $X \equiv \{x_1 = (c_1 = 0, v_1 = 5.5), x_2 = (c_2 = 20, v_2 = 10)\}$. Algebra shows that x_1 is both the β-best task and the $\hat{\beta}$-best task, and since it is costless the person completes it immediately. Now suppose that δ increases from $.8$ to $.9$. This makes x_2 both the β-best and the $\hat{\beta}$-best task, but the person never completes it. Similarly, suppose once again that $\delta = .8$, but that the menu becomes $X' \equiv \{(c'_1 = 0, v'_1 = 11), (c'_2 = 20, v'_2 = 20)\}$ — the per-period benefit from each task is doubled. Again, x_2 is now both the β-best and the $\hat{\beta}$-best task, but the person never completes it. Both transformations make it more important for the person to do the task, in the sense that the present discounted value of rewards have been increased for each possible cost. With this increased importance, the person decides to do the more costly task, but this plan to incur high costs induces procrastination.

While this example illustrates that increasing importance can exacerbate procrastination, this phenomenon is not a universal. For instance, while doubling the per-period benefit of each task induces procrastination in this example by changing the person's preferred task from x_1 to x_2, doubling the per-period benefit of each task once more eliminates procrastination — this time by

[27] More precisely, two things are true: No perception-perfect strategy generated by the new choice set is strictly Pareto-dominated by any perception-perfect strategy generated by the old choice set, and for every perception-perfect strategy generated by the old choice set there exists a perception-perfect strategy in the new choice set that weakly Pareto-dominates it.

motivating the person to complete x_2 right away. The remainder of this section explores under what conditions increased importance induces procrastination.

The above example illustrates two ways in which a person's goals might become more important: the person might become more patient, or the per-period benefit from each task might become larger (while holding the cost of each task constant). The example also clearly shows how the importance-exacerbates-procrastination phenomenon relies on there being a menu of options. Indeed, in the one-task context, increasing importance reduces the likelihood of procrastination (with a minor caveat):

Proposition 6 Consider a person who faces singleton menu $X \equiv \{(c, v)\}$:

(1) When $\hat{\beta} = 1$, $S^{pp}(\beta, \hat{\beta}, \delta, (c, v)) \neq \{\mathbf{s}^{\emptyset}\}$ if and only if $\frac{v}{c} \geq \frac{1-\beta\delta}{\beta\delta}$ (or equivalently if and only if $\delta \geq \frac{c}{\beta v + \beta c}$); and

(2) When $\hat{\beta} < 1$, $S^{pp}(\beta, \hat{\beta}, \delta, (c, v)) \neq \{\mathbf{s}^{\emptyset}\}$ if $\frac{v}{c} \geq \frac{1-\beta\delta}{\beta\delta}$ (or equivalently if $\delta \geq \frac{c}{\beta v + \beta c}$), and $S^{pp}(\beta, \hat{\beta}, \delta, (c, v)) = \{\mathbf{s}^{\emptyset}\}$ if $\frac{v}{c} < \frac{1-\beta\delta/\hat{\beta}}{\beta\delta}$ (or equivalently if $\delta < \frac{c}{\beta v + \beta c/\hat{\beta}}$).

Part (1) of Proposition 6 says that when there is just one task, as the person becomes more patient (δ increases) or as the per-period benefit of the task increases relative to the cost ($\frac{v}{c}$ increases), a completely naive person always becomes less likely to procrastinate.[28] The intuition for these results is straightforward: A completely naive person always thinks she will do the task next period if she waits now, and as either a person becomes more patient or the magnitude of benefits relative to costs increases, the person becomes more and more likely to prefer doing the task now to doing it next period, and hence does not delay. Notice, however, that although increasing patience makes a completely naive person less likely to procrastinate, she may never do the task for any $\delta < 1$ (i.e., if $\frac{v}{c} < \frac{1-\beta}{\beta}$ then $\frac{v}{c} < \frac{1-\beta\delta}{\beta\delta}$ for all $\delta < 1$). In other words, there are tasks which a person would not do for any δ — those tasks with a particularly large ratio of cost to per-period benefit. In the example above, task x_2 is precisely such a task.[29]

Part (2) of Proposition 6 says that a similar result holds when $\hat{\beta} < 1$ in the sense that for δ or $\frac{v}{c}$ large enough the person completes the task and for δ or $\frac{v}{c}$ small enough the person does not

[28] Although Propostion 6 establishes that a completely naive person doesn't complete the task if and only if $\frac{v}{c} < \frac{1-\beta\delta}{\beta\delta}$, this does not always correspond to procrastination, since the task may not be β-worthwhile. Since not completing the task represents procrastination whenever $\frac{1-\delta}{\beta\delta} \leq \frac{v}{c} < \frac{1-\beta\delta}{\beta\delta}$, a more precise statement is that increasing δ or increasing $\frac{v}{c}$ always decreases the likelihood of procrastination over the range where the task is β-worthwhile.

[29] Note that any episode of procrastination as $\delta \to 1$ causes an infinite welfare loss, since $\lim_{\delta \to 1} WL(s^{\emptyset}, \delta) = \infty$ whenever there exists $(c, v) \in X$ such that $v > 0$.

complete the task. However, the caveat to the general result is that for a partially naive person there is a range where whether the person completes the task can be non-monotonic in either $\frac{v}{c}$ or δ.[30]

Proposition 6 says that, with the caveat of minor non-monotonicities for partial naifs, when there is a single task available, increasing importance makes a person less likely to procrastinate. This effect is also present when there are multiple options available in the sense that increasing importance makes a person less likely to procrastinate on any specific option. But with multiple options, increasing importance also makes costly tasks more attractive, which can make procrastination more likely. Hence, whether increased importance increases procrastination depends on the relative importance of these two forces. As illustrated by our earlier example, locally either force can dominate.

How does a person behave as her goals become really important? To answer this question, it is useful to introduce some notation. Define the maximal benefit available given a menu X by $v^{\max}(X) \equiv \sup\{v|(c,v) \in X\}$, with $v^{\max}(X) = \infty$ if $\{v|(c,v) \in X\}$ is unbounded above. Define the set of "productive tasks" — those tasks which yield benefits that cannot be obtained at lower

---

30 Because such non-monotonicities are a pervasive feature of our model and since we want to emphasize that our results that follow are *not* driven by these non-monotonicities, it is perhaps useful illustrate these non-monotonicities in detail with an example. Suppose $\beta = .5$, $\hat{\beta} = .58$, and $\delta = .9$, and consider comparative statics over $\frac{v}{c}$. For various ranges of $\frac{v}{c}$, the following table characterizes the person's current tolerance for delay (i.e., $d(\beta|\hat{\beta})$), her perceived future tolerance for delay (i.e., $d(\hat{\beta}|\hat{\beta})$), and whether she will (eventually) complete the task (i.e., whether $S^{pp}(\beta, \hat{\beta}, \delta, X) \neq \{s^{\emptyset}\}$):

$benefit/cost\ ratio$	$d(\beta\|\hat{\beta})$	$d(\hat{\beta}\|\hat{\beta})$	$S^{pp}(\beta, \hat{\beta}, \delta, X)$
$.283 < \frac{v}{c} < .292$	8	5	$= \{\mathbf{s}^{\emptyset}\}$
$.292 < \frac{v}{c} < .306$	8	4	$= \{\mathbf{s}^{\emptyset}\}$
$.306 < \frac{v}{c} < .324$	7	4	$= \{\mathbf{s}^{\emptyset}\}$
$.324 < \frac{v}{c} < .345$	6	4	$= \{\mathbf{s}^{\emptyset}\}$
$.345 < \frac{v}{c} < .348$	6	3	$= \{\mathbf{s}^{\emptyset}\}$
$.348 < \frac{v}{c} < .382$	5	3	$= \{\mathbf{s}^{\emptyset}\}$
$.382 < \frac{v}{c} < .408$	4	3	$= \{\mathbf{s}^{\emptyset}\}$
$.408 < \frac{v}{c} < .434$	4	2	$= \{\mathbf{s}^{\emptyset}\}$
$.434 < \frac{v}{c} < .521$	3	2	$= \{\mathbf{s}^{\emptyset}\}$
$.521 < \frac{v}{c} < .535$	2	2	$\neq \{\mathbf{s}^{\emptyset}\}$
$.535 < \frac{v}{c} < .696$	2	1	$= \{\mathbf{s}^{\emptyset}\}$
$.696 < \frac{v}{c} < .916$	1	1	$\neq \{\mathbf{s}^{\emptyset}\}$
$.916 < \frac{v}{c} < 1.222$	1	0	$= \{\mathbf{s}^{\emptyset}\}$
$1.222 < \frac{v}{c}$	0	0	$\neq \{\mathbf{s}^{\emptyset}\}$

This example illustrates that when $\hat{\beta} < 1$, the decrease in procrastination as $\frac{v}{c}$ increases need not be monotonic and how these non-monotonicities are driven by the discreteness of $d(\beta|\hat{\beta})$ and $d(\hat{\beta}|\hat{\beta})$. The example also illustrates the sufficient conditions in Proposition 6. We know that a person will complete a task when $\frac{v}{c}$ is large enough such that the person prefers completing the task now as opposed to next period — that is when her current tolerance for delay is $d(\beta|\hat{\beta}) = 0$. We know that a person will never complete the task when $\frac{v}{c}$ is small enough such that $d(\beta|\hat{\beta}) \geq d(\hat{\beta}|\hat{\beta})+1$. The proof of Proposition 1 provides more detail on these conditions.

cost — by

$$\tilde{X}(X) \equiv \{(c,v) \in X | \nexists (c',v') \in X \text{ with } c' \leq c,\, v' \geq v, \text{ and } (c',v') \neq (c,v)\}.$$

Define the maximal productive cost by $c^{\max}(X) \equiv \sup\{c | (c,v) \in \tilde{X}(X)\}$, with $c^{\max}(X) = \infty$ if $\{c | (c,v) \in \tilde{X}(X)\}$ is unbounded above.

If the menu of tasks X is finite, then $c^{\max}(X) < \infty$, $v^{\max}(X) < \infty$, and $(c^{\max}(X), v^{\max}(X)) \in X$. In this case, as the person's goals become sufficiently important, $(c^{\max}(X), v^{\max}(X))$ becomes both the β-best task and the $\hat{\beta}$-best task. Hence, for sufficiently high importance, the person procrastinates if and only if she procrastinates on task $(c^{\max}(X), v^{\max}(X))$. Even when the menu of tasks X is infinite, an analogous logic holds whenever there is an upper bound on the maximal productive cost — that is, if $c^{\max}(X) < \infty$. Restricting attention to $v^{\max}(X) < \infty$ (because otherwise perception-perfect strategies do not exist), as the person's goals become very important, $(c^{\max}(X), v^{\max}(X))$ becomes (almost) both the β-best and $\hat{\beta}$-best task, in which case she makes her decision about whether to delay as if she were facing the single task $(c^{\max}(X), v^{\max}(X))$. Proposition 7 summarizes this logic.

Proposition 7 Consider a menu X with $c^{\max}(X) < \infty$ and $v^{\max}(X) < \infty$.

(1) If $\frac{v^{\max}(X)}{c^{\max}(X)} > \frac{1-\beta}{\beta}$, then there exists $\delta^* < 1$ such that $S^{pp}(\beta, \hat{\beta}, \delta, X) \neq \{\mathbf{s}^{\emptyset}\}$ for all $\delta > \delta^*$; and if $\frac{v^{\max}(X)}{c^{\max}(X)} < \frac{1-\beta/\hat{\beta}}{\beta}$, then there exists $\delta^{**} < 1$ such that $S^{pp}(\beta, \hat{\beta}, \delta, X) = \{\mathbf{s}^{\emptyset}\}$ for all $\delta > \delta^{**}$; and

(2) For any strictly increasing function $f : \mathbb{R}_+ \to \mathbb{R}_+$ that satisfies $\frac{f(v^{\max}(X))}{c^{\max}(X)} > \frac{1-\beta\delta}{\beta\delta}$, $S^{pp}(\beta, \hat{\beta}, \delta, X'(X)) \neq \{\mathbf{s}^{\emptyset}\}$ where $X'(X) \equiv \{(c, f(v)) | (c,v) \in X\}$.

Part (1) of Proposition 7 says that if the maximal productive task $(c^{\max}(X), v^{\max}(X))$ is one that a person would do in isolation if she were patient enough, then she does some task when facing menu X if she is patient enough; and if the maximal productive task $(c^{\max}(X), v^{\max}(X))$ is one that a person would not do in isolation for any δ, then she procrastinates when facing menu X if she is patient enough. Part (2) of Proposition 7 establishes that if the per-period benefits of tasks are made sufficiently large — large enough that the person would do the maximal productive task $(c^{\max}(X), f(v^{\max}(X)))$ in isolation — then the person for sure does some task.[31]

The logic behind part (1) of Proposition 7 is as follows: Increasing patience leads the person to plan to do a more costly task; but if she eventually plans to do a task which has a sufficiently

[31] We remind the reader that these results imply nothing about how the person would behave when her goals are only mildly important.

high cost relative to its per-period benefit, then she never does any task. When $c^{\max}(X) < \infty$, the latter condition holds if the maximal productive task $(c^{\max}(X), v^{\max}(X))$ is such a task. But in some instances, there may not be a maximal productive task — that is, there may be no limit to the costs that can be incurred to yield higher benefits. In such instances, it is in fact more likely that the person eventually plans to do a task which she would not do in isolation for any δ, and therefore it is more likely that the person procrastinates if she is patient enough. Indeed, under the reasonable assumption that the task menu X exhibits positive but decreasing marginal returns — that is, a given incremental increase in task cost has a smaller and smaller impact on per-period benefits — increasing patience is quite likely to induce procrastination. To formalize this claim, define $L(X) \equiv \lim_{c\to\infty} \sup\{\frac{v'}{c'} \mid (c', v') \in X \text{ and } c' \geq c\}$. $L(X)$ is the "limit ratio" of per-period benefits to costs as the person expends ever more effort, and can be loosely interpreted as the limit of the marginal return to additional effort.

Proposition 8 Suppose $c^{\max}(X) = \infty$.

(1) If $L(X) = 0$, then for all β and $\hat{\beta}$ such that $\hat{\beta} > \beta$, there exists $\delta^* < 1$ such that $S^{pp}(\beta, \hat{\beta}, \delta, X) = \{\mathbf{s}^{\emptyset}\}$ for all $\delta > \delta^*$; and

(2) For all β and $\hat{\beta} > \beta$ such that $L(X) < \frac{1-\beta/\hat{\beta}}{\beta}$, there exists $X' \equiv \{(c, v) \in X | c \leq \bar{c}\}$ for some $\bar{c} < \infty$ and $\delta^* < 1$ such that $S^{pp}(\beta, \hat{\beta}, \delta, X') = \{\mathbf{s}^{\emptyset}\}$ for all $\delta > \delta^*$.

Part (1) of Proposition 8 is the clearest and most striking example of how sufficiently high importance can exacerbate procrastination: When there is no upper bound on how much effort a person can productively put into a task, but the marginal return to additional effort eventually becomes arbitrarily small, then a sufficiently patient person with any degree of naivete surely procrastinates. This result reflects the intuitions developed in Lemma 3: As δ approaches 1, the value of even a small increase in per-period benefits becomes enormous, and hence the optimal task involves a very large cost. But the per-period benefit of the optimal task becomes very small relative to its cost, and therefore the person procrastinates.

Part (2) of Proposition 8 shows that a similar logic applies even if the marginal return to effort does not become arbitrarily small. As long as $L(X) < \frac{1-\beta/\hat{\beta}}{\beta}$, there exist productive tasks which have sufficiently high costs relative to their per-period benefits that the person would not do them for any δ. A problem arises, however, because $L(X) > 0$ implies that no perception-perfect strategy exists for δ close enough to 1; if the marginal returns to increased effort are everywhere non-negligible, a sufficiently patient person will eventually prefer an arbitrarily large cost.[32] Of course, existence is guaranteed if we truncate X by throwing out all tasks with very large costs. Part (2) establishes that we can always perform such a truncation in a way that yields procrastination if the person is patient enough — by leaving in the choice set some of those productive tasks with sufficiently high costs.

To illustrate examples where part (1) of Proposition 8 applies, consider cases where the task menu can be represented by a function — that is, when there exists a function $v : \mathbb{R}_+ \to \mathbb{R}_+$ such that $X = \{(c, v(c)) | c \in \mathbb{R}_+\}$. The following corollary follows directly from Proposition 8:

Corollary 1 Suppose a function $v : \mathbb{R}_+ \to \mathbb{R}_+$ is continuously differentiable with $v'(c) > 0$ for all c and $\lim_{c\to\infty} v'(c) = 0$. If $X = \{(c, v(c)) | c \in \mathbb{R}_+\}$, then for all β and $\hat{\beta}$ such that $\hat{\beta} > \beta$, there exists $\delta^* < 1$ such that $S^{pp}(\beta, \hat{\beta}, \delta, X) = \{s^{\emptyset}\}$ for all $\delta > \delta^*$.

This corollary indicates that there are some natural classes of task menus such that any degree of naivete induces procrastination if a person is patient enough. Examples include $v(c) = (a + bc)^d$, where $a \geq 0$, $b > 0$ and $d \in (0, 1)$; $v(c) = \ln(c + 1)$; and $v(c) = \frac{c}{ac+1}$, where $a > 0$.

Proposition 6 establishes that if the per-period benefits of a single task are made sufficiently large, then the person for sure does the task. Part (2) of Proposition 7 extends this result to the case where there is a menu of tasks X with a maximal productive task $(c^{\max}(X), v^{\max}(X))$. Since the person must eventually choose (almost) the maximal productive task, when its benefits become large enough, the person for sure does some task. We now consider what happens if the per-period benefits of a task are made sufficiently large when there is no maximal productive task. In such situations, it becomes possible that as the per-period benefits of all tasks are made sufficiently large, the person procrastinates for sure. But this can occur only if the effect of increasing importance making the person plan to do a more costly task dominates the effect of increasing importance eventually leading a person to do any specific task.

To illustrate, we consider multiplicative transformations of the benefits of all tasks, defining $X(k) \equiv \{(c, kv) | (c, v) \in X\}$. In fact, we believe such multiplicative transformations are of particular interest. For instance, if we interpret the task cost as the effort expended to find a good investment opportunity and the task benefits as the per-dollar return, then the factor k corresponds

[32] Formally, when $c^{\max}(X) = \infty$, existence given β, $\hat{\beta}$, δ, and X requires $L(X) \leq \frac{1-\delta}{\beta\delta}$.

to the quantity of funds that a person plans to invest.[33] The examples above which satisfy the conditions in Corollary 1 illustrate the indeterminate effect that increasing k has on procrastination. If $v(c) = (a + bc)^d$, where $a, b > 0$ and $d \in (0, 1)$, then there exists $k^* > 0$ such that $S^{pp}(\beta, \hat{\beta}, \delta, X(k)) = \{\mathbf{s}^\emptyset\}$ for all $k > k^*$.[34] If $v(c) = \frac{c}{ac+1}$, where $a > 0$, then there exists $k^* > 0$ such that $S^{pp}(\beta, \hat{\beta}, \delta, X(k)) \neq \{\mathbf{s}^\emptyset\}$ for all $k > k^*$. We have unfortunately found no useful general characterization of when multiplicative transformations of the benefits induce procrastination.

Consider why increasing per-period benefits has ambiguous effects on procrastination whereas increasing δ has unambiguous effects. Both increasing δ and increasing k cause the person to plan on a more costly task. But whereas increasing δ eventually has a negligible impact on the short-term benefits of completing the task immediately, increasing k can obviously have a significant impact on the per-period benefit of the chosen task even when k becomes very large. Formally, denoting the β-best task by $(c^*(\delta, k), v^*(\delta, k))$, Proposition 8 follows because the $\lim_{\delta \to 1} \frac{v^*(\delta,k)}{c^*(\delta,k)}$ is small. The corresponding condition for increasing k is $\lim_{k \to \infty} \frac{kv^*(\delta,k)}{c^*(\delta,k)}$ being small, but since k is in the numerator this term can remain large even as $\frac{v^*(\delta,k)}{c^*(\delta,k)}$ becomes small.

The indeterminate effects of increasing per-period benefits of course hold for more general transformations. Indeed, for any X such that a person completes a task, there exists a monotonic transformation of the benefits that induces procrastination — by making the β-best task significantly more costly without drastically increasing the per-period benefits of tasks. And for any X such that a person procrastinates, there exists a monotonic transformation of the benefits that induces completing a task — by significantly increasing the per-period benefits of all tasks without significantly increasing the cost of the β-best task. But one class of transformations always eliminates procrastination no matter X:

Lemma 4 Let $X(\eta) = \{(c, v + \eta) | (c, v) \in X\}$. Then for all β, $\hat{\beta}$, δ, and X, there exists η^* such that $S^{pp}(\beta, \hat{\beta}, \delta, X(\eta)) \neq \{\mathbf{s}^\emptyset\}$ for all $\eta > \eta^*$.

A sufficiently large additive transformation in which the benefits of all tasks are increased by the same absolute amount eliminates procrastination. This result drives home one final time the underlying logic behind the importance-exacerbates-procrastination results. These results are driven

[33] See O'Donoghue and Rabin (1998) for some calibrations of retirement behavior considering precisely this interpretation (although in a slightly different model).

[34] When $a = 0$, increasing per-period benefits has no effect on procrastination — that is, $S^{pp}(\beta, \hat{\beta}, \delta, X(k))$ is independent of k.

by the fact that increased importance may induce a person to plan to exert more effort. But an additive transformation of the benefits does not affect the optimal amount of effort to exert: The cost associated with the β-best and $\hat{\beta}$-best tasks are unchanged. Hence, when the benefits become large enough, the person for sure completes a task.

5. Discussion and Conclusion

In this paper, we have identified a number of features of naive procrastination. We believe the main lessons from our analysis apply beyond our specific model, and may be quite relevant in important economic contexts. In O'Donoghue and Rabin (1998), for instance, we calibrate a model of whether and how a person invests her savings for retirement. We argue that people may significantly delay transferring savings from their checking accounts into higher-interest accounts, even when the long-term benefits of doing so are enormous. For example, suppose a person is saving \$10,000 for retirement 30 years from now. If the person currently earns 1% interest in her checking account, and knows of an easy opportunity to earn 6% instead, it is well worth making the transfer. While the person may or may not procrastinate when the 6% account is her sole alternative, choice can greatly exacerbate her procrastination for the same reasons developed in this paper. Because investing for retirement is so important, she may decide that she should put in the effort to do it right — to find (say) a 6.2% account. No matter how cheap it is to transfer her money to the 6% account, if the total cost of transferring the money first from the 1% to the 6% account and then into the 6.2% account is greater than the cost of switching directly to the 6.2% account, the person may procrastinate transferring to the 6% account, because she persistently plans to transfer directly into the 6.2% account. Indeed, she may procrastinate searching for the 6.2% account for decades, making herself much poorer in retirement than she would have been had she settled for the very good option of investing in the (mere) 6% account.

Moreover, such procrastination can be exacerbated when the person has more money to invest, reflecting our importance-exacerbates-procrastination arguments. For example, the person may severely procrastinate when her principal is \$10,000, but not when her principal is only \$1,000. The logic is as in this paper: The person plans and executes a quick-and-easy investment strategy for the \$1,000, while she plans — but does *not* execute — a more ambitious investment strategy for the \$10,000. The calibration exercises in O'Donoghue and Rabin (1998) also support our claim in Section 3 that it needn't take much naivete to generate procrastination, and reinforce our worries that it would be a mistake for economists to focus solely on models of 100% sophistication.

We conclude with some conjectures about some realism-enhancing extensions of our model. While our model assumes that the person completes only one task, the person might, for instance, be working on a number of unrelated projects at the same time, and hence may have more than one task to complete. Having more than one task to complete can make procrastination less likely because it causes a person to perceive delay as more costly. If the person has other projects to do, then not doing a task today forces her either to delay this task for more than one day, or to delay the other projects she was planning to do tomorrow. The logic of procrastination says that a person procrastinates because she perceives the cost of delay to be small; if the person is busy, therefore, she is less likely to procrastinate.[35]

Notice that being busy is not at all the same as other types of high opportunity costs. Indeed, if the opportunity cost of doing a task is exogenously increased in all periods, a person is *more* likely to procrastinate. The reason that having other tasks to do mitigates procrastination is that it raises the perceived cost of delay *without* raising the cost of doing the task *immediately*. If every day a person chooses between doing her taxes versus playing tennis, the more she likes tennis (i.e., the higher her opportunity cost) the less likely she is to do her taxes. But if in addition to paying her taxes she must also paint the workbench, rotate the carburetors, or do other household chores, she might pay her taxes soon. Having only to pay her taxes makes her think that all she is foregoing by playing tennis today is playing tennis tomorrow. Having to do these other chores too means that delaying one's taxes delays completion of all chores by one day.[36]

A second realistic extension is to suppose that a person need not or cannot complete a task all at once. On many projects, a person can do a quick, cheap fix, initiating some benefits in the short run, and later come back and do a proper job to yield the rest of the benefits. If a person is writing a research paper, she needn't wait until she has the final version, with all the desired results, before distributing it. She can distribute a preliminary draft, labeled as such, telling readers her intention to produce a more complete paper in the near future. If a person is deciding how best to invest her

[35] As Voltaire might have said, "Se vuole aver' fatto una cosa immediatamente, la dia in mano a una persona molto ocupata."

[36] While assuming the person has many tasks to complete might suggest a decrease in procrastination, the importance-exacerbates-procrastination results are likely to generalize. For instance, if every project has $c^{\max}(X) = \infty$ and $L(X) = 0$, then no matter how many projects a person faces, she will surely procrastinate on all projects if she is patient enough.

money, she can put her money in an easy-to-initiate, relatively good investment in the short run, and then continue to search for the ideal investment.

Some preliminary analysis of such situations suggest two implications. First, if a person can improve on what she has done in the past, it becomes more likely that she does something — that is, that she at least does the quick fix. In the model of this paper, a person foregoes completing low-cost options in large part because (by assumption) doing so precludes doing the better job that the person plans to do in the future. While we noted above that adding a 6.2% investment opportunity may stop a person from making a 6% investment, there are also cases where adding a 6% investment opportunity to the 6.2% opportunity leads a person from never investing to making the 6% investment (which she perceives as a short-term fix). Since in many situations a person can do a short-run fix, our model may overstate the likelihood that a person does absolutely nothing.

While the presence of quick fixes makes it less likely that the person does nothing, however, it also makes it less likely that the person "completes" the task. If a person has done a quick fix in the past, then the short-term damage caused by delay in completing the task is relatively small, and therefore procrastination is more likely. If a person has already put her savings in the easy-to-initiate 6% account, the short-term damage caused by delay in finding the 6.2% account is smaller than if her savings were still in her 1% checking account. Similarly, once a person has taken half an hour to cover the roof with plastic to effectively stop the leaks, the cost of delay in fixing the roof is smaller than if the roof were uncovered. Hence, while our model overstates the likelihood that a person does nothing, it also overstates the likelihood that a person actually completes the task.

Conventional economic theory says that a person does something if she believes the benefits outweigh the costs. Our model and the proposed extensions highlight how, in addition to the overall costs and benefits, the timing of the costs and benefits can be an important determinant of whether and when a person undertakes some endeavor. Models with present-biased preferences assume that people engage in long-run cost/benefit analysis in formulating their *plans*. But they posit that people use a sort of immediate-cost/immediate-benefit analysis in deciding whether to do something *now*. This alternative conception of when people take actions challenges traditional axioms of choice. As an alternative to the conventional Weak Axiom of Revealed Preference, earlier papers generate what might be called the *Weak Axiom of Revealed Procrastination*: If we observe somebody never doing a single given task, we learn very little about whether she prefers to do that task. This paper generates what can be called the *Strong Axiom of Revealed Procrastination*: If we observe somebody never doing a task when she has a menu of tasks from which to choose, we learn even less.

Appendix: Proofs

Proof of Lemma 1: (1) If $a_t = (c, v) \in X$, then $V^t(a_t, \hat{\mathbf{s}}, \hat{\beta}, \delta) = \frac{\hat{\beta}\delta}{1-\delta}v - c$. By definition, $x^*(\hat{\beta}, \delta, X) = \arg\max_{(c,v)\in X} \frac{\hat{\beta}\delta}{1-\delta}v - c$, and therefore $x^*(\hat{\beta}, \delta, X) = \arg\max_{a\in A\setminus\{\emptyset\}} V^t(a, \hat{\mathbf{s}}, \hat{\beta}, \delta)$. Since any dynamically consistent beliefs $\hat{\mathbf{s}}(\hat{\beta}, \delta)$ must satisfy $\hat{a}_t(\hat{\beta}, \delta) = \arg\max_{a\in A} V^t(a, \hat{\mathbf{s}}, \hat{\beta}, \delta)$ for all t, it follows that for all t either $\hat{a}_t(\hat{\beta}, \delta) = \emptyset$ or $\hat{a}_t(\hat{\beta}, \delta) = x^*(\hat{\beta}, \delta, X)$.

(2) If $x^*(\hat{\beta}, \delta, X) \equiv (c, v)$ is not $\hat{\beta}$-worthwhile, then $\frac{\hat{\beta}\delta}{1-\delta}v - c < 0$, which implies $\frac{\hat{\beta}\delta}{1-\delta}v - c < \hat{\beta}\delta^\tau \left[\frac{\delta}{1-\delta}v - c\right]$ for all $\tau \in \{1, 2, ...\}$. For any proposed $\hat{\mathbf{s}}(\hat{\beta}, \delta)$ that satisfies for all t either $\hat{a}_t(\hat{\beta}, \delta) = \emptyset$ or $\hat{a}_t(\hat{\beta}, \delta) = x^*(\hat{\beta}, \delta, X)$, the latter inequality implies $\arg\max_{a\in A} V^t(a, \hat{\mathbf{s}}, \hat{\beta}, \delta) = \emptyset$, which implies $\hat{\mathbf{s}}(\hat{\beta}, \delta)$ must satisfy $\hat{a}_t(\hat{\beta}, \delta) = \emptyset$ for all t.

Suppose $x^*(\hat{\beta}, \delta, X) \equiv (c, v)$ is $\hat{\beta}$-worthwhile. The definition of $d(\hat{\beta}|\hat{\beta})$ implies that for any $d' \in \{1, 2, ..., d(\hat{\beta}|\hat{\beta})\}$, if $\hat{a}_t(\hat{\beta}, \delta) = x^*(\hat{\beta}, \delta, X)$ and $\hat{a}_{t-d}(\hat{\beta}, \delta) = \emptyset$ for all $d \in \{1, 2, ..., d' - 1\}$, then $\arg\max_{a\in A} V^{t-d'}(a, \hat{\mathbf{s}}, \hat{\beta}, \delta) = \emptyset$. The definition of $d(\hat{\beta}|\hat{\beta})$ also implies that for $d' = d(\hat{\beta}|\hat{\beta}) + 1$, if $\hat{a}_t(\hat{\beta}, \delta) = x^*(\hat{\beta}, \delta, X)$ and $\hat{a}_{t-d}(\hat{\beta}, \delta) = \emptyset$ for all $d \in \{1, 2, ..., d' - 1\}$, then $\arg\max_{a\in A} V^{t-d'}(a, \hat{\mathbf{s}}, \hat{\beta}, \delta) = x^*(\hat{\beta}, \delta, X)$. It follows that $\hat{\mathbf{s}}(\hat{\beta}, \delta)$ must have $\hat{a}_t(\hat{\beta}, \delta) = x^*(\hat{\beta}, \delta, X)$ every $d(\hat{\beta}|\hat{\beta}) + 1$ periods, and $\hat{a}_t(\hat{\beta}, \delta) = \emptyset$ otherwise. This condition can be satisfied only if $\min\{t \in \{2, 3, ...\}|\hat{a}_t(\hat{\beta}, \delta) = x^*(\hat{\beta}, \delta, X)\} \in \{2, 3, ..., d(\hat{\beta}|\hat{\beta}) + 2\}$. The result follows.

Q.E.D.

Proof of Lemma 2: A logic analogous to that in the proof of Lemma 1(1) implies that for any $\hat{\mathbf{s}}$, $\arg\max_{a\in A} V^t(a, \hat{\mathbf{s}}, \beta, \delta) \in \{\emptyset, x^*(\beta, \delta, X)\}$ for all t, which implies that any perception-perfect strategy must satisfy for all t either $a_t(\beta, \hat{\beta}, \delta) = \emptyset$ or $a_t(\beta, \hat{\beta}, \delta) = x^*(\beta, \delta, X)$. Moreover, given any dynamically consistent beliefs $\hat{\mathbf{s}}(\hat{\beta}, \delta)$, $a_t(\beta, \hat{\beta}, \delta) = x^*(\beta, \delta, X)$ if and only if
$V^t(x^*(\beta, \delta, X), \hat{\mathbf{s}}(\hat{\beta}, \delta), \beta, \delta) \geq V^t(\emptyset, \hat{\mathbf{s}}(\hat{\beta}, \delta), \beta, \delta)$.

If $x^*(\hat{\beta}, \delta, X)$ is not $\hat{\beta}$-worthwhile, then Lemma 1 implies the unique set of dynamically consistent beliefs is $\hat{\mathbf{s}}(\hat{\beta}, \delta) = \mathbf{s}^\emptyset$, in which case $V^t(\emptyset, \hat{\mathbf{s}}(\hat{\beta}, \delta), \beta, \delta) = 0$. If $x^*(\hat{\beta}, \delta, X) \equiv (c', v')$ is $\hat{\beta}$-worthwhile, then Lemma 1 implies that any dynamically consistent beliefs $\hat{\mathbf{s}}(\hat{\beta}, \delta)$ must yield for all t, $V^t(\emptyset, \hat{\mathbf{s}}(\hat{\beta}, \delta), \beta, \delta) = \beta\delta^d \left[\frac{\delta}{1-\delta}v' - c'\right]$ for some $d \in \{1, 2, ..., d(\hat{\beta}|\hat{\beta}) + 1\}$. Since $\frac{\hat{\beta}\delta}{1-\delta}v' - c' \geq 0$ implies $\frac{\delta}{1-\delta}v' - c' \geq 0$, it follows that $V^t(\emptyset, \hat{\mathbf{s}}(\hat{\beta}, \delta), \beta, \delta) \geq 0$ for all t.

If $x^*(\beta, \delta, X)$ is not β-worthwhile, then for any dynamically consistent beliefs $\hat{\mathbf{s}}(\hat{\beta}, \delta)$, $V^t(x^*(\beta, \delta, X), \hat{\mathbf{s}}(\hat{\beta}, \delta), \beta, \delta) < 0 \leq V^t(\emptyset, \hat{\mathbf{s}}(\hat{\beta}, \delta), \beta, \delta)$ for all t, and therefore $S^{pp}(\beta, \hat{\beta}, \delta, X) = \{\mathbf{s}^\emptyset\}$.

Suppose $x^*(\beta,\delta,X) \equiv (c^*,v^*)$ is β-worthwhile but $d(\hat\beta|\hat\beta)+1 \le d(\beta|\hat\beta)$. If $x^*(\beta,\delta,X)$ is β-worthwhile, then $x^*(\hat\beta,\delta,X)$ must be $\hat\beta$-worthwhile, in which case any dynamically consistent beliefs $\hat{\mathbf{s}}(\hat\beta,\delta)$ must yield for all t, $V^t(\emptyset,\hat{\mathbf{s}}(\hat\beta,\delta),\beta,\delta) = \beta\delta^d\left[\frac{\delta}{1-\delta}v'-c'\right]$ for some $d \in \{1,2,...,d(\hat\beta|\hat\beta)+1\}$. This implies $V^t(\emptyset,\hat{\mathbf{s}}(\hat\beta,\delta),\beta,\delta) \ge \beta\delta^{d(\hat\beta|\hat\beta)+1}\left[\frac{\delta}{1-\delta}v'-c'\right]$ for all t. Since $V^t(x^*(\beta,\delta,X),\hat{\mathbf{s}}(\hat\beta,\delta),\beta,\delta) = \frac{\beta\delta}{1-\delta}v^*-c^*$, $a_t(\beta,\hat\beta,\delta) = x^*(\beta,\delta,X)$ only if $\frac{\beta\delta}{1-\delta}v^*-c^* \ge \beta\delta^{d(\hat\beta|\hat\beta)+1}\left[\frac{\delta}{1-\delta}v'-c'\right]$. But since the definition of $d(\beta|\hat\beta)$ implies $\frac{\beta\delta}{1-\delta}v^*-c^* < \beta\delta^d\left[\frac{\delta}{1-\delta}v'-c'\right]$ for all $d \le d(\beta|\hat\beta)$, $d(\hat\beta|\hat\beta)+1 \le d(\beta|\hat\beta)$ implies $a_t(\beta,\hat\beta,\delta)=\emptyset$ for all t. Hence, if $x^*(\beta,\delta,X)$ is β-worthwhile but $d(\hat\beta|\hat\beta)+1 \le d(\beta|\hat\beta)$, $\mathbf{s}^\emptyset$ is the perception-perfect strategy for any dynamically consistent beliefs $\hat{\mathbf{s}}(\hat\beta,\delta)$, and thus $S^{pp}(\beta,\hat\beta,\delta,X)=\{\mathbf{s}^\emptyset\}$.

Suppose $x^*(\beta,\delta,X)$ is β-worthwhile and $d(\hat\beta|\hat\beta)+1 > d(\beta|\hat\beta)$. It is straightforward to show $d(\hat\beta|\hat\beta) \le d(\beta|\hat\beta)$, and therefore $d(\hat\beta|\hat\beta) = d(\beta|\hat\beta)$. Since the definition of $d(\beta|\hat\beta)$ implies $\beta\delta^{d(\beta|\hat\beta)}\left[\frac{\delta}{1-\delta}v'-c'\right] > \frac{\beta\delta}{1-\delta}v^*-c^* \ge \beta\delta^{d(\beta|\hat\beta)+1}\left[\frac{\delta}{1-\delta}v'-c'\right]$, for any dynamically consistent beliefs $\hat{\mathbf{s}}(\hat\beta,\delta)$ the perception-perfect strategy satisfies $a_t(\beta,\hat\beta,\delta)=x^*(\beta,\delta,X)$ if and only if $\min\left\{d\in\{1,2,...\}|\hat a_{t+d}(\hat\beta,\delta)=x^*(\hat\beta,\delta,X)\right\} = d(\hat\beta|\hat\beta)+1 = d(\beta|\hat\beta)+1$, and otherwise $a_t(\beta,\hat\beta,\delta) = \emptyset$. Hence, if $x^*(\beta,\delta,X)$ is β-worthwhile and $d(\hat\beta|\hat\beta)+1 = d(\beta|\hat\beta)$, then $\mathbf{s}^\emptyset \notin S^{pp}(\beta,\hat\beta,\delta,X)$. Finally, any $\mathbf{s}\in S^{pp}(\beta,\hat\beta,\delta,X)$ must satisfy $x(\mathbf{s}) = x^*(\beta,\delta,X)$ and

$$\tau(\mathbf{s}) = \min\left\{t\in\{1,2,...\}\,\middle|\,\left(\min\left\{d\in\{1,2,...\}|\hat a_{t+d}(\hat\beta,\delta)=x^*(\hat\beta,\delta,X)\right\}\right) = d(\hat\beta|\hat\beta)+1\right\}.$$

Clearly, either $\tau(\mathbf{s})=1$ (when $\tau(\hat{\mathbf{s}}) = d(\hat\beta|\hat\beta)+2$) or $\tau(\mathbf{s})=\tau(\hat{\mathbf{s}})$.

Q.E.D.

Proof of Proposition 1: This proof develops a series of properties that will be used in this and other proofs. Throughout this proof, we use notation $(c^*(\beta),v^*(\beta)) \equiv x^*(\beta,\delta,X)$.

(AA): $c^*(\beta)$ and $v^*(\beta)$ are nondecreasing in β.

Proof: Consider any β and $\beta' > \beta$. The definition of $x^*(\beta,\delta,X)$ implies $\frac{\beta\delta}{1-\delta}v^*(\beta)-c^*(\beta) \ge \frac{\beta\delta}{1-\delta}v^*(\beta')-c^*(\beta')$; the definition of $x^*(\beta',\delta,X)$ implies $\frac{\beta'\delta}{1-\delta}v^*(\beta)-c^*(\beta) \le \frac{\beta'\delta}{1-\delta}v^*(\beta')-c^*(\beta')$; and combining these inequalities yields

$$\frac{\beta\delta}{1-\delta}\left(v^*(\beta')-v^*(\beta)\right) \le \left(c^*(\beta')-c^*(\beta)\right) \le \frac{\beta'\delta}{1-\delta}\left(v^*(\beta')-v^*(\beta)\right).$$

This condition can hold only if $\left(v^*(\beta')-v^*(\beta)\right)\left(c^*(\beta')-c^*(\beta)\right) \ge 0$; and if $\left(v^*(\beta')-v^*(\beta)\right) < 0$, then $\beta' > \beta$ implies $\frac{\beta\delta}{1-\delta}\left(v^*(\beta')-v^*(\beta)\right) > \frac{\beta'\delta}{1-\delta}\left(v^*(\beta')-v^*(\beta)\right)$ and the condition cannot be satisfied. Hence, we must have $\left(v^*(\beta')-v^*(\beta)\right) \ge 0$ and $\left(c^*(\beta')-c^*(\beta)\right) \ge 0$.

(BB): $\left(\frac{\delta}{1-\delta}v^*(\beta) - c^*(\beta)\right)$ is nondecreasing in β.

Proof: Consider any β and $\beta' > \beta$. The proof of Property (AA) establishes $(v^*(\beta') - v^*(\beta)) \geq 0$ and $(c^*(\beta') - c^*(\beta)) \leq \frac{\beta'\delta}{1-\delta}(v^*(\beta') - v^*(\beta))$, which together imply $(c^*(\beta') - c^*(\beta)) \leq \frac{\delta}{1-\delta}(v^*(\beta') - v^*(\beta))$, which in turn yields $\left(\frac{\delta}{1-\delta}v^*(\beta) - c^*(\beta)\right) \leq \left(\frac{\delta}{1-\delta}v^*(\beta') - c^*(\beta')\right)$.

(CC): If $\hat{\beta} = 1$, then $S^{pp}(\beta, \hat{\beta}, \delta, X) = \{\mathbf{s}^{\emptyset}\}$ if and only if $\frac{\beta\delta}{1-\delta}v^*(\beta) - c^*(\beta) < \beta\delta\left[\frac{\delta}{1-\delta}v^*(\hat{\beta}) - c^*(\hat{\beta})\right]$.

Proof: When $\hat{\beta} = 1$, the unique set of dynamically consistent beliefs is

$$\hat{\mathbf{s}}(\hat{\beta}, \delta) = (x^*(\hat{\beta}, \delta, X), x^*(\hat{\beta}, \delta, X), x^*(\hat{\beta}, \delta, X), ...).$$

Hence, the person always compares doing the β-best task now to doing the $\hat{\beta}$-best task next period, which yields the condition wait if and only if $\frac{\beta\delta}{1-\delta}v^*(\beta) - c^*(\beta) < \beta\delta\left[\frac{\delta}{1-\delta}v^*(\hat{\beta}) - c^*(\hat{\beta})\right]$.

(DD): For any $\hat{\beta}$, $S^{pp}(\beta, \hat{\beta}, \delta, X) = \{\mathbf{s}^{\emptyset}\}$ if $\frac{\delta}{1-\delta}v^*(\beta) - \frac{1}{\beta}c^*(\beta) < \delta\left(\frac{\delta}{1-\delta}v^*(\hat{\beta}) - \frac{1}{\hat{\beta}}c^*(\hat{\beta})\right)$.

Proof: The proof of Lemma 2 establishes $S^{pp}(\beta, \hat{\beta}, \delta, X) = \{\mathbf{s}^{\emptyset}\}$ if $d(\hat{\beta}|\hat{\beta}) + 1 \leq d(\beta|\hat{\beta})$, and since $d(\beta|\hat{\beta})$ must satisfy

$$\delta^{d(\beta|\hat{\beta})}\left[\frac{\delta}{1-\delta}v^*(\hat{\beta}) - c^*(\hat{\beta})\right] > \frac{\delta}{1-\delta}v^*(\beta) - \frac{1}{\beta}c^*(\beta) \geq \delta^{d(\beta|\hat{\beta})+1}\left[\frac{\delta}{1-\delta}v^*(\hat{\beta}) - c^*(\hat{\beta})\right]$$

and $d(\hat{\beta}|\hat{\beta})$ must satisfy

$$\delta^{d(\hat{\beta}|\hat{\beta})}\left[\frac{\delta}{1-\delta}v^*(\hat{\beta}) - c^*(\hat{\beta})\right] > \frac{\delta}{1-\delta}v^*(\hat{\beta}) - \frac{1}{\hat{\beta}}c^*(\hat{\beta}) \geq \delta^{d(\hat{\beta}|\hat{\beta})+1}\left[\frac{\delta}{1-\delta}v^*(\hat{\beta}) - c^*(\hat{\beta})\right],$$

$d(\hat{\beta}|\hat{\beta}) + 1 \leq d(\beta|\hat{\beta})$ if $\frac{\delta}{1-\delta}v^*(\beta) - \frac{1}{\beta}c^*(\beta) < \delta\left(\frac{\delta}{1-\delta}v^*(\hat{\beta}) - \frac{1}{\hat{\beta}}c^*(\hat{\beta})\right)$.

(EE): For any $\hat{\beta}$, $S^{pp}(\beta, \hat{\beta}, \delta, X) \neq \{\mathbf{s}^{\emptyset}\}$ if $\frac{\beta\delta}{1-\delta}v^*(\beta) - c^*(\beta) \geq \beta\delta\left(\frac{\delta}{1-\delta}v^*(\hat{\beta}) - c^*(\hat{\beta})\right)$.

Proof: The proof of Lemma 2 establishes $S^{pp}(\beta, \hat{\beta}, \delta, X) \neq \{\mathbf{s}^{\emptyset}\}$ if $d(\hat{\beta}|\hat{\beta}) = d(\beta|\hat{\beta})$, which must hold if $d(\beta|\hat{\beta}) = 0$, and $d(\beta|\hat{\beta}) = 0$ if and only if $\frac{\beta\delta}{1-\delta}v^*(\beta) - c^*(\beta) \geq \beta\delta\left(\frac{\delta}{1-\delta}v^*(\hat{\beta}) - c^*(\hat{\beta})\right)$.

Continuity Property A: $\left(\frac{\beta\delta}{1-\delta}v^*(\beta) - c^*(\beta)\right)$ is continuous in β.

Proof: $\left(\frac{\beta\delta}{1-\delta}v^*(\beta) - c^*(\beta)\right) = \max_{(c,v)\in X}\left(\frac{\beta\delta}{1-\delta}v - c\right)$, which is continuous if X is closed.

Continuity Property B: For every $\varepsilon > 0$ there exists $\beta' > \beta$ such that for all $\hat{\beta} \in (\beta, \beta')$, $\left(\frac{\delta}{1-\delta}v^*(\hat{\beta}) - c^*(\hat{\beta})\right) - \left(\frac{\delta}{1-\delta}v^*(\beta) - c^*(\beta)\right) < \varepsilon$.

Proof: The definition of $x^*(\beta, \delta, X)$ implies $\frac{\beta\delta}{1-\delta}v^*(\beta) - c^*(\beta) \geq \frac{\beta\delta}{1-\delta}v^*(\hat{\beta}) - c^*(\hat{\beta})$, which implies $\left(\frac{\delta}{1-\delta}v^*(\hat{\beta}) - c^*(\hat{\beta})\right) - \left(\frac{\delta}{1-\delta}v^*(\beta) - c^*(\beta)\right) \leq \frac{1-\beta}{\beta}\left(c^*(\hat{\beta}) - c^*(\beta)\right)$. It is therefore sufficient to show that for any $\bar{\varepsilon} > 0$ there exists $\beta' > \beta$ such that for all $\hat{\beta} \in (\beta, \beta')$, $\left(c^*(\hat{\beta}) - c^*(\beta)\right) < \bar{\varepsilon}$. Define $\beta^* \equiv \inf\{\beta' > \beta | c^*(\beta') > c^*(\beta)\}$. If either β^* doesn't exist (because $c^*(\beta') = c^*(\beta)$ for all $\beta' > \beta$) or $\beta^* > \beta$, the result follows. Suppose $\beta^* = \beta$. If we define $\hat{c} \equiv \lim_{\beta' \to \beta^+} c^*(\beta')$ and $\hat{v} \equiv \lim_{\beta' \to \beta^+} v^*(\beta')$, both of which exist since c^* and v^* are nondecreasing, then we must have $\frac{\beta\delta}{1-\delta}\hat{v} - \hat{c} \geq \frac{\beta\delta}{1-\delta}v^*(\beta) - c^*(\beta)$, since otherwise there would exist a neighborhood $\hat{X}$ of $(\hat{c}, \hat{v})$ such that $(c^*(\beta), v^*(\beta))$ is $\hat{\beta}$-preferred to any $x \in \hat{X}$ for $\hat{\beta}$ close enough to β. Given the definition of $x^*(\beta, \delta, X)$, we can conclude that $(\hat{c}(\beta), \hat{v}(\beta)) = (c^*(\beta), v^*(\beta))$, and the result follows.[37]

Proof of Part (1): Follows directly from the proof of Lemma 2.

Proof of Part (2): Suppose $S^{pp}(\beta, \hat{\beta}, \delta, X) \neq \{\mathbf{s}^\emptyset\}$ for $\hat{\beta} = 1$, in which case Property (CC) implies $\frac{\beta\delta}{1-\delta}v^*(\beta) - c^*(\beta) \geq \beta\delta\left[\frac{\delta}{1-\delta}v^*(1) - c^*(1)\right]$. Since Property (BB) implies $\left[\frac{\delta}{1-\delta}v^*(\hat{\beta}) - c^*(\hat{\beta})\right] \leq \left[\frac{\delta}{1-\delta}v^*(1) - c^*(1)\right]$ for all $\hat{\beta} < 1$, it follows that $\frac{\beta\delta}{1-\delta}v^*(\beta) - c^*(\beta) \geq \beta\delta\left[\frac{\delta}{1-\delta}v^*(\hat{\beta}) - c^*(\hat{\beta})\right]$ for all $\hat{\beta}$, in which case Property (EE) implies $S^{pp}(\beta, \hat{\beta}, \delta, X) \neq \{\mathbf{s}^\emptyset\}$ for all $\hat{\beta}$.

Suppose $S^{pp}(\beta, \hat{\beta}, \delta, X) = \{\mathbf{s}^\emptyset\}$ for $\hat{\beta} = 1$, in which case Property (CC) implies $\frac{\delta}{1-\delta}v^*(\beta) - \frac{1}{\beta}c^*(\beta) < \delta\left[\frac{\delta}{1-\delta}v^*(\hat{\beta}) - \frac{1}{\hat{\beta}}c^*(\hat{\beta})\right]$ for $\hat{\beta} = 1$. Then Property (DD) and Continuity Property A imply that there exists $\beta^{**} < 1$ such that $S^{pp}(\beta, \hat{\beta}, \delta, X) = \{\mathbf{s}^\emptyset\}$ for all $\hat{\beta} > \beta^{**}$.

Lemma 1 implies $S^{pp}(\beta, \hat{\beta}, \delta, X) \neq \{\mathbf{s}^\emptyset\}$ for $\hat{\beta} = \beta$. "Generically"[38], $d(\beta|\beta)$ satisfies

$$\delta^{d(\beta|\beta)}\left[\frac{\delta}{1-\delta}v^*(\beta) - c^*(\beta)\right] > \frac{\delta}{1-\delta}v^*(\beta) - \frac{1}{\beta}c^*(\beta) > \delta^{d(\beta|\beta)+1}\left[\frac{\delta}{1-\delta}v^*(\beta) - c^*(\beta)\right].$$

Since $d(\beta|\hat{\beta})$ must satisfy

$$\delta^{d(\beta|\hat{\beta})}\left[\frac{\delta}{1-\delta}v^*(\hat{\beta}) - c^*(\hat{\beta})\right] > \frac{\delta}{1-\delta}v^*(\beta) - \frac{1}{\beta}c^*(\beta) \geq \delta^{d(\beta|\hat{\beta})+1}\left[\frac{\delta}{1-\delta}v^*(\hat{\beta}) - c^*(\hat{\beta})\right],$$

Continuity Property B implies there exists $\beta' > \beta$ such that $d(\beta|\hat{\beta}) = d(\beta|\beta)$ for all $\hat{\beta} \in (\beta, \beta')$. Similarly, since $d(\hat{\beta}|\hat{\beta})$ must satisfy

$$\delta^{d(\hat{\beta}|\hat{\beta})}\left[\frac{\delta}{1-\delta}v^*(\hat{\beta}) - c^*(\hat{\beta})\right] > \frac{\delta}{1-\delta}v^*(\hat{\beta}) - \frac{1}{\hat{\beta}}c^*(\hat{\beta}) \geq \delta^{d(\hat{\beta}|\hat{\beta})+1}\left[\frac{\delta}{1-\delta}v^*(\hat{\beta}) - c^*(\hat{\beta})\right],$$

Continuity Properties A and B imply there exists $\beta'' > \beta$ such that $d(\hat{\beta}|\hat{\beta}) = d(\beta|\beta)$ for all $\hat{\beta} \in (\beta, \beta'')$. If $\beta^* \equiv \min\{\beta', \beta''\}$, then $d(\beta|\hat{\beta}) = d(\hat{\beta}|\hat{\beta}) = d(\beta|\beta)$ for all $\hat{\beta} < \beta^*$, and therefore $S^{pp}(\beta, \hat{\beta}, \delta, X) \neq \{\mathbf{s}^\emptyset\}$ for all $\hat{\beta} < \beta^*$. Since it is clear that $\beta^* \leq \beta^{**}$, the result follows.

Q.E.D.

---

[37] This last step relies on our assumption that if the set $\mathbf{B} \equiv \arg\max_{(c,v)\in X}\left[\frac{\beta\delta}{1-\delta}v - c\right]$ is not a singleton, then $x^*(\beta, \delta, X)$ is the task $(c^*, v^*) \in \mathbf{B}$ such that $v^* = \max\{v|(c,v) \in \mathbf{B}\}$.

[38] By "generically", we mean ruling out knife-edge parameters where $\frac{\delta}{1-\delta}v^*(\beta) - \frac{1}{\beta}c^*(\beta) = \delta^{d(\beta|\beta)+1}\left[\frac{\delta}{1-\delta}v^*(\beta) - c^*(\beta)\right]$.

Proof of Proposition 2: It is sufficient to show there exists a singleton such X. Suppose $X \equiv \{(c,v)\}$. The task is β-worthwhile if $\frac{\beta\delta}{1-\delta}v - c > 0$ or $\frac{v}{c} > \frac{1-\delta}{\beta\delta}$. Property (DD) from the proof of Proposition 1 implies $S^{pp}(\beta,\hat{\beta},\delta,X) = \{\mathbf{s}^{\emptyset}\}$ if $\frac{\beta\delta}{1-\delta}v - c < \beta\delta\left[\frac{\delta}{1-\delta}v - \frac{1}{\hat{\beta}}c\right]$, or $\frac{v}{c} < \frac{1-\beta\delta/\hat{\beta}}{\beta\delta}$. For any $(\beta,\hat{\beta},\delta)$ such that $\hat{\beta} > \beta$, we have $\frac{1-\beta\delta/\hat{\beta}}{\beta\delta} > \frac{1-\delta}{\beta\delta}$, which implies that there exists (c,v) such that $\frac{1-\beta\delta/\hat{\beta}}{\beta\delta} > \frac{v}{c} > \frac{1-\delta}{\beta\delta}$. The result follows.

Q.E.D.

Proof of Proposition 3: (1a) If $\mathbf{s}^{\emptyset}$ does not reflect procrastination, then no task in X is β-worthwhile. For any alternative strategy $\mathbf{s} \neq \mathbf{s}^{\emptyset}$, $U^{\tau(\mathbf{s})}(\mathbf{s},\beta,\delta) < 0 = U^{\tau(\mathbf{s})}(\mathbf{s}^{\emptyset},\beta,\delta)$, which implies that $\mathbf{s}^{\emptyset}$ is Pareto-efficient. If $\mathbf{s}^{\emptyset}$ reflects procrastination, then task $x^*(\beta,\delta,X) \equiv (c,v)$ is β-worthwhile. For any alternative strategy $\mathbf{s}$ satisfying $\tau(\mathbf{s}) = 1$ and $x(\mathbf{s}) = x^*(\beta,\delta,X)$, $U^t(\mathbf{s},\beta,\delta) \geq 0 = U^t(\mathbf{s}^{\emptyset},\beta,\delta)$ for all t, which implies that $\mathbf{s}^{\emptyset}$ is Pareto-inefficient.

(1b) Suppose $\mathbf{s}^{\emptyset}$ does not reflect procrastination, so no task in X is β-worthwhile. If task $x^*(1,\delta,X) \equiv (c^*,v^*)$ is not $(\beta = 1)$-worthwhile, then $WL(\mathbf{s}^{\emptyset},\delta) = 0$, and otherwise $WL(\mathbf{s}^{\emptyset},\delta) = \frac{\delta}{1-\delta}\frac{v^*}{c^*} - 1$. But since task (c^*,v^*) is not β-worthwhile, we know $\frac{\beta\delta}{1-\delta}v^* - c^* < 0$ or $\frac{\delta}{1-\delta}\frac{v^*}{c^*} < \frac{1}{\beta}$. Hence, $WL(\mathbf{s}^{\emptyset},\delta) < \frac{1}{\beta} - 1 = \frac{1-\beta}{\beta}$.

(2a) Lemma 2 implies that if $S^{pp}(\beta,\hat{\beta},\delta,X) \neq \{\mathbf{s}^{\emptyset}\}$ then there exists $\mathbf{s} \in S^{pp}(\beta,\hat{\beta},\delta,X)$ such that $\tau(\mathbf{s}) \leq d(\beta|\beta) + 1$. The result is then a direct implication of parts (2b) and (2c).

(2b) First, consider any $\mathbf{s} \in S^{pp}(\beta,\hat{\beta},\delta,X)$ with $\tau(\mathbf{s}) \leq d(\beta|\beta) + 1$, where Lemma 2 implies $x(\mathbf{s}) = x^*(\beta,\delta,X)$. This strategy is Pareto-efficient because (i) any alternative $\mathbf{s}' \neq \mathbf{s}$ with $\tau(\mathbf{s}') = \tau(\mathbf{s})$ and $x(\mathbf{s}') = x^*(\beta,\delta,X)$ yields $U^t(\mathbf{s}',\beta,\delta) = U^t(\mathbf{s},\beta,\delta)$ for all t; (ii) any alternative $\mathbf{s}' \neq \mathbf{s}$ with $\tau(\mathbf{s}') = \tau(\mathbf{s})$ and $x(\mathbf{s}') \neq x^*(\beta,\delta,X)$ yields $U^{\tau(\mathbf{s})}(\mathbf{s}',\beta,\delta) < U^{\tau(\mathbf{s})}(\mathbf{s},\beta,\delta)$; (iii) any alternative $\mathbf{s}' \neq \mathbf{s}$ with $\tau(\mathbf{s}') > \tau(\mathbf{s})$ yields $U^{\tau(\mathbf{s}')}(\mathbf{s}',\beta,\delta) < U^{\tau(\mathbf{s}')}(\mathbf{s},\beta,\delta)$ because having completed $x^*(\beta,\delta,X)$ in the past is better than completing any task now; and (iv) any alternative $\mathbf{s}' \neq \mathbf{s}$ with $\tau(\mathbf{s}') < \tau(\mathbf{s})$ yields $U^{\tau(\mathbf{s}')}(\mathbf{s}',\beta,\delta) < U^{\tau(\mathbf{s}')}(\mathbf{s},\beta,\delta)$ because $\tau(\mathbf{s}) \leq d(\beta|\beta) + 1$ implies completing $x^*(\beta,\delta,X)$ in $\tau(\mathbf{s}) - \tau(\mathbf{s}')$ periods is better than completing any task now.

Second, consider any $\mathbf{s} \in S^{pp}(\beta,\hat{\beta},\delta,X)$ with $\tau(\mathbf{s}) > d(\beta|\beta)+1$, where again Lemma 2 implies $x(\mathbf{s}) = x^*(\beta,\delta,X)$. Consider alternative strategy $\mathbf{s}' \neq \mathbf{s}$ with $\tau(\mathbf{s}') = 1$ and $x(\mathbf{s}') = x^*(\beta,\delta,X)$. The definition of $d(\beta|\beta)$ implies $U^1(\mathbf{s}',\beta,\delta) > U^1(\mathbf{s},\beta,\delta)$. Moreover, since having completed $x^*(\beta,\delta,X)$ in the past is better than completing $x^*(\beta,\delta,X)$ now or in the future, $U^t(\mathbf{s}',\beta,\delta) \geq U^t(\mathbf{s},\beta,\delta)$ for all $t \geq 2$. It follows that $\mathbf{s}$ is Pareto-inefficient.

(2c) Consider any $\mathbf{s} \in S^{pp}(\beta, \hat{\beta}, \delta, X)$ with $\tau(\mathbf{s}) \le d(\beta|\beta) + 1$, where again Lemma 2 implies $x(\mathbf{s}) = x^*(\beta, \delta, X) \equiv (c, v)$. Letting $x^*(1, \delta, X) \equiv (c^*, v^*)$,
$WL(\mathbf{s}, \delta) = \frac{1}{c^*}\left[\left(\frac{\delta}{1-\delta}v^* - c^*\right) - \delta^{\tau(\mathbf{s})-1}\left(\frac{\delta}{1-\delta}v - c\right)\right]$. If $\tau(\mathbf{s}) = 1$, then
$\delta^{\tau(\mathbf{s})-1}\left(\frac{\delta}{1-\delta}v - c\right) = \frac{\delta}{1-\delta}v - c > \frac{\delta}{1-\delta}v - \frac{1}{\beta}c$. If $\tau(\mathbf{s}) \in \{2, 3, ..., d(\beta|\beta) + 1\}$, then the definition of $d(\beta|\beta)$ implies $\delta^{\tau(\mathbf{s})-1}\left(\frac{\delta}{1-\delta}v - c\right) > \frac{\delta}{1-\delta}v - \frac{1}{\beta}c$. Hence, $WL(\mathbf{s}, \delta) < \frac{1}{c^*}\left[\left(\frac{\delta}{1-\delta}v^* - c^*\right) - \left(\frac{\delta}{1-\delta}v - \frac{1}{\beta}c\right)\right]$. Finally, $x^*(\beta, \delta, X) \equiv (c, v)$ implies $\frac{\delta}{1-\delta}v - \frac{1}{\beta}c \ge \frac{\delta}{1-\delta}v^* - \frac{1}{\beta}c^*$, which yields $WL(\mathbf{s}, \delta) < \frac{1-\beta}{\beta}$.

Q.E.D.

Proof of Proposition 4: (1) It is sufficient to prove each result for a singleton X. Suppose $X \equiv \{(c, v)\}$. Property (DD) from the proof of Proposition 1 implies that $S^{pp}(\beta, \hat{\beta}, \delta, X) = \{\mathbf{s}^\emptyset\}$ if $\frac{\delta}{1-\delta}v - \frac{1}{\beta}c < \delta\left(\frac{\delta}{1-\delta}v - \frac{1}{\hat{\beta}}c\right)$, or $\frac{v}{c} < \frac{1-\beta\delta/\hat{\beta}}{\beta\delta}$. As long as $\frac{v}{c} > \frac{1-\delta}{\delta}$, $WL(\mathbf{s}^\emptyset, \delta) = \frac{1}{c}\left[\left(\frac{\delta}{1-\delta}v - c\right) - (0)\right] = \frac{\delta}{1-\delta}\frac{v}{c} - 1$ (and otherwise $WL(\mathbf{s}^\emptyset, \delta) = 0$). Hence, given β, $\hat{\beta} > \beta$, and δ, for any $\varepsilon > 0$ there exists $X \equiv \{(c, v)\}$ such that $S^{pp}(\beta, \hat{\beta}, \delta, X) = \{\mathbf{s}^\emptyset\}$ and $WL(\mathbf{s}^\emptyset, \delta) > \left(\frac{1-\beta\delta/\hat{\beta}}{\beta}\frac{1}{1-\delta} - 1\right) - \varepsilon$. Since $\hat{\beta} > \beta$ implies $\left(\frac{1-\beta\delta/\hat{\beta}}{\beta}\frac{1}{1-\delta} - 1\right) > \frac{1-\beta}{\beta}$, the result follows.

(2) This result is a straightforward extension of the proof of part (1) — if we can choose X and δ, then we can make $WL(\mathbf{s}^\emptyset, \delta)$ arbitrarily large by choosing δ sufficiently close to 1.

Q.E.D.

Proof of Lemma 3: (1) Since $(c^*, v^*) \equiv \arg\max_{(c,v)\in X}\left(\frac{\beta\delta}{1-\delta}v - c\right)$ and $(c^{*\prime}, v^{*\prime}) \equiv \arg\max_{(c,v)\in X'}\left(\frac{\beta\delta'}{1-\delta'}v - c\right)$, $\frac{\beta\delta}{1-\delta}v^* - c^* \neq \frac{\beta\delta'}{1-\delta'}v^{*\prime} - c^{*\prime}$ would contradict the premise that $(c', v') \in X'$ if and only if there exists $(c, v) \in X$ such that $\frac{\beta\delta'}{1-\delta'}v' - c' = \frac{\beta\delta}{1-\delta}v - c$.

(2) Follows directly from the following property:

Property (FF): Let $x^*(\beta, \delta, X) = (c, v)$. For any $\hat{\beta}$, $S^{pp}(\beta, \hat{\beta}, \delta, X) = \{\mathbf{s}^\emptyset\}$ if $\frac{\delta}{1-\delta}v - \frac{1}{\beta}c < \delta\left(\frac{\delta}{1-\delta}v - \frac{1}{\hat{\beta}}c\right)$.

Proof: Letting $x^*(\hat{\beta}, \delta, X) = (c', v')$, Property (DD) from the proof of Proposition 1 says $S^{pp}(\beta, \hat{\beta}, \delta, X) = \{\mathbf{s}^\emptyset\}$ if $\frac{\delta}{1-\delta}v - \frac{1}{\beta}c < \delta\left(\frac{\delta}{1-\delta}v' - \frac{1}{\hat{\beta}}c'\right)$. Since $x^*(\hat{\beta}, \delta, X) = (c', v')$ implies $\frac{\delta}{1-\delta}v' - \frac{1}{\hat{\beta}}c' \ge \frac{\delta}{1-\delta}v - \frac{1}{\hat{\beta}}c$, the result follows.

Q.E.D.

Proof of Proposition 5: (1) Proposition 1 establishes for $\hat{\beta} = \beta$, for any menu X, $S^{pp}(\beta, \hat{\beta}, \delta, X) \neq \{\mathbf{s}^{\emptyset}\}$ if and only if there exists $x \in X$ that is β-worthwhile. $X' \supset X$ implies that if there exists $x \in X$ that is β-worthwhile, then there exists $x \in X'$ that is β-worthwhile, and the result follows.

(2) Define $x^*(\beta, \delta, X) \equiv (c^*, v^*)$, and consider $x' \equiv (c', v')$ with $c' > c^*$. If $\frac{\beta\delta}{1-\delta}v' - c' > \frac{\beta\delta}{1-\delta}v^* - c^*$, then $x^*(\beta, \delta, X \cup x') = x'$. If in addition $\frac{\delta}{1-\delta}v' - \frac{1}{\beta}c' < \delta\left(\frac{\delta}{1-\delta}v' - \frac{1}{\hat{\beta}}c'\right)$, then Property (FF) from the proof of Lemma 3 implies $S^{pp}(\beta, \hat{\beta}, \delta, X) = \{\mathbf{s}^{\emptyset}\}$. We can rewrite the first inequality as $\frac{v'-v^*}{c'-c^*} > \frac{1-\delta}{\beta\delta}$, and the second inequality as $\frac{v'}{c'} < \frac{1-\beta\delta/\hat{\beta}}{\beta\delta}$. Since $\hat{\beta} > \beta$ implies $\frac{1-\beta\delta/\hat{\beta}}{\beta\delta} > \frac{1-\delta}{\beta\delta}$, no matter what is (c^*, v^*), there exists $(c', v') \in \mathbb{R}^2_+$ with $c' > c^*$ that satisfies both properties.

Q.E.D.

Proof of Proposition 6: (1) When $X \equiv \{(c, v)\}$, Property (CC) from the proof of Proposition 1 becomes if $\hat{\beta} = 1$, then $S^{pp}(\beta, \hat{\beta}, \delta, X) = \{\mathbf{s}^{\emptyset}\}$ if and only if $\frac{\beta\delta}{1-\delta}v - c < \beta\delta\left[\frac{\delta}{1-\delta}v - c\right]$, which can be rearranged as $\frac{v}{c} < \frac{1-\beta\delta}{\beta\delta}$ or $\delta < \frac{c}{\beta v + \beta c}$.

(2) When $X \equiv \{(c, v)\}$, Property (EE) from the proof of Proposition 1 becomes for any $\hat{\beta}$, $S^{pp}(\beta, \hat{\beta}, \delta, X) \neq \{\mathbf{s}^{\emptyset}\}$ if $\frac{\beta\delta}{1-\delta}v - c \geq \beta\delta\left[\frac{\delta}{1-\delta}v - c\right]$, which can be rearranged as $\frac{v}{c} \geq \frac{1-\beta\delta}{\beta\delta}$ or $\delta \geq \frac{c}{\beta v + \beta c}$. When $X \equiv \{(c, v)\}$, Property (DD) from the proof of Proposition 1 becomes for any $\hat{\beta}$, $S^{pp}(\beta, \hat{\beta}, \delta, X) = \{\mathbf{s}^{\emptyset}\}$ if $\frac{\delta}{1-\delta}v - \frac{1}{\beta}c < \delta\left[\frac{\delta}{1-\delta}v - \frac{1}{\hat{\beta}}c\right]$, which can be rearranged as $\frac{v}{c} < \frac{1-\beta\delta/\hat{\beta}}{\beta\delta}$ or $\delta < \frac{c}{\beta v + \beta c/\hat{\beta}}$.

Q.E.D.

Proof of Proposition 7: (1) Define $x^*(\beta, \delta, X) \equiv (c^*(\delta), v^*(\delta))$. We first prove $\lim_{\delta\to 1} c^*(\delta) = c^{\max}(X)$ and $\lim_{\delta\to 1} v^*(\delta) = v^{\max}(X)$. It is straightforward to show that c^* and v^* must be non-decreasing in δ (the logic is exactly analogous to that used to prove Property (AA) in the proof of Proposition 1), which implies $\lim_{\delta\to 1} c^*(\delta)$ and $\lim_{\delta\to 1} v^*(\delta)$ both exist. Note that for any (c, v) and (c', v') satisfying $c' > c$ and $v' > v$, there exists $\delta' < 1$ such that $\frac{\beta\delta}{1-\delta}v' - c' > \frac{\beta\delta}{1-\delta}v - c$ for all $\delta > \delta'$. Since for any $(c, v) \in X$ with $c < c^{\max}(X)$ there exists $(c', v') \in X$ with $c' > c$ and $v' > v$, for any $(c, v) \in X$ with $c < c^{\max}(X)$ we must have $\lim_{\delta\to 1} c^*(\delta) > c$. Finally, since $c^*(\delta) \leq c^{\max}(X)$ for all δ implies $\lim_{\delta\to 1} c^*(\delta) \leq c^{\max}(X)$ we can conclude $\lim_{\delta\to 1} c^*(\delta) = c^{\max}(X)$. It then follows directly that $\lim_{\delta\to 1} v^*(\delta) = v^{\max}(X)$.

Given $\lim_{\delta\to 1} c^*(\delta) = c^{\max}(X)$ and $\lim_{\delta\to 1} v^*(\delta) = v^{\max}(X)$, $\lim_{\delta\to 1} \frac{v^*(\delta)}{c^*(\delta)} = \frac{v^{\max}(X)}{c^{\max}(X)}$. Then $\frac{v^{\max}(X)}{c^{\max}(X)} < \frac{1-\beta/\hat{\beta}}{\beta}$ implies there exists $\delta^* < 1$ such that for all $\delta > \delta^*$, $\frac{v^*(\delta)}{c^*(\delta)} < \frac{1-\beta/\hat{\beta}}{\beta}$. Since $\frac{1-\beta/\hat{\beta}}{\beta} <$

$\frac{1-\beta\delta/\hat{\beta}}{\beta\delta}$ for any $\delta < 1$, Property (FF) from the proof of Lemma 3 implies $S^{pp}(\beta, \hat{\beta}, \delta, X) = \{\mathbf{s}^{\emptyset}\}$ for all $\delta > \delta^*$.

To prove that if $\frac{v^{\max}(X)}{c^{\max}(X)} > \frac{1-\beta}{\beta}$ then there exists $\delta^* < 1$ such that $S^{pp}(\beta, \hat{\beta}, \delta, X) \neq \{\mathbf{s}^{\emptyset}\}$ for all $\delta > \delta^*$, we prove that if $\frac{v^{\max}(X)}{c^{\max}(X)} > \frac{1-\beta\delta}{\beta\delta}$ then $S^{pp}(\beta, \hat{\beta}, \delta, X) \neq \{\mathbf{s}^{\emptyset}\}$, from which the result follows because $\frac{v^{\max}(X)}{c^{\max}(X)} > \frac{1-\beta}{\beta}$ implies there exists $\delta^* < 1$ such that $\frac{v^{\max}(X)}{c^{\max}(X)} > \frac{1-\beta\delta}{\beta\delta}$ for all $\delta > \delta^*$. Let $x^*(\beta, \delta, X) \equiv (c^*, v^*)$ and $x^*(\hat{\beta}, \delta, X) \equiv (c', v')$. $x^*(\hat{\beta}, \delta, X) = (c', v')$ implies $\frac{\hat{\beta}\delta}{1-\delta}v' - c' \geq \frac{\hat{\beta}\delta}{1-\delta}v^{\max}(X) - c^{\max}(X)$, or $\left(\frac{\hat{\beta}\delta}{1-\delta}\frac{v'}{c'} - 1\right) c' \geq \left(\frac{\hat{\beta}\delta}{1-\delta}\frac{v^{\max}(X)}{c^{\max}(X)} - 1\right) c^{\max}(X)$, and then $c' \leq c^{\max}(X)$ implies $\frac{v'}{c'} > \frac{v^{\max}(X)}{c^{\max}(X)}$. Hence, $\frac{v'}{c'} > \frac{1-\beta\delta}{\beta\delta}$, which implies $\frac{\beta\delta}{1-\delta}v' - c' \geq \beta\delta\left[\frac{\delta}{1-\delta}v' - c'\right]$. Then $x^*(\beta, \delta, X) = (c^*, v^*)$ implies $\frac{\beta\delta}{1-\delta}v^* - c^* \geq \frac{\beta\delta}{1-\delta}v' - c' \geq \beta\delta\left[\frac{\delta}{1-\delta}v' - c'\right]$, in which case Property (EE) from the proof of Proposition 1 implies $S^{pp}(\beta, \hat{\beta}, \delta, X) \neq \{\mathbf{s}^{\emptyset}\}$.

(2) f increasing implies that $c^{\max}(X'(X)) = c^{\max}(X)$ and $v^{\max}(X'(X)) = f(v^{\max}(X))$. Hence, $\frac{v^{\max}(X'(X))}{c^{\max}(X'(X))} = \frac{f(v^{\max}(X))}{c^{\max}(X)} > \frac{1-\beta\delta}{\beta\delta}$, in which case it follows from the proof of part (1) that $S^{pp}(\beta, \hat{\beta}, \delta, X) \neq \{\mathbf{s}^{\emptyset}\}$.

Q.E.D.

Proof of Proposition 8: First note that for any X with $c^{\max}(X) = \infty$ and $L(X) < \frac{1-\beta/\hat{\beta}}{\beta}$, there exists $\underline{c} < \infty$ such that any $(c, v) \in X$ with $c > \underline{c}$ has $\frac{v}{c} < \frac{1-\beta/\hat{\beta}}{\beta}$. Since $\frac{1-\beta\delta/\hat{\beta}}{\beta\delta} > \frac{1-\beta/\hat{\beta}}{\beta}$ for any $\delta < 1$, any $(c, v) \in X$ with $c > \underline{c}$ has $\frac{v}{c} < \frac{1-\beta\delta/\hat{\beta}}{\beta\delta}$ for all $\delta < 1$. Given Property (FF) from the proof of Lemma 3, if $x^*(\beta, \delta, X)$ and $x^*(\hat{\beta}, \delta, X)$ both exist and $x^*(\beta, \delta, X) = (c, v)$ for some $c > \underline{c}$, then $S^{pp}(\beta, \hat{\beta}, \delta, X) = \{\mathbf{s}^{\emptyset}\}$.

(1) $c^{\max}(X) = \infty$ and $L(X) = 0$ imply that $x^*(\beta, \delta, X)$ and $x^*(\hat{\beta}, \delta, X)$ both exist. Defining $x^*(\beta, \delta, X) \equiv (c^*(\delta), v^*(\delta))$ as in the proof of Proposition 7, it remains to show that there exists $\delta^* < 1$ such that $c^*(\delta) > \underline{c}$ for all $\delta > \delta^*$. Since in the proof of Proposition 7 the argument that $\lim_{\delta\to 1} c^*(\delta) = c^{\max}(X)$ does not rely on $c^{\max}(X) < \infty$, the result follows.

(2) Construct $X' \equiv \{(c, v) \in X | c \leq \bar{c}\}$ by choosing $\bar{c} > \underline{c}$ such that there exists $(c, v) \in \tilde{X}(X)$ with $\underline{c} < c < \bar{c}$. X closed implies X' is closed and bounded and therefore compact, and hence $x^*(\beta, \delta, X')$ and $x^*(\hat{\beta}, \delta, X')$ both exist. Moreover, since by construction there exists $(c, v) \in \tilde{X}(X)$ with $\underline{c} < c < \bar{c}$, then $c^{\max}(X') > \underline{c}$, from which it follows that there exists $\delta^* < 1$ such that $c^*(\delta) > \underline{c}$ for all $\delta > \delta^*$, and therefore $S^{pp}(\beta, \hat{\beta}, \delta, X) = \{\mathbf{s}^{\emptyset}\}$ for all $\delta > \delta^*$.

Q.E.D.

Proof of Lemma 4: First note that $x^*(\beta, \delta, X) = (c^*, v^*)$ implies $x^*(\beta, \delta, X(\eta)) = (c^*, v^* + \eta)$ and $x^*(\hat{\beta}, \delta, X) = (c', v')$ implies $x^*(\beta, \delta, X(\eta)) = (c', v' + \eta)$. Property (EE) from the proof of Proposition 1 then implies that for any η, $S^{pp}(\beta, \hat{\beta}, \delta, X(\eta)) \neq \{\mathbf{s}^{\emptyset}\}$ if $\frac{\beta\delta}{1-\delta}(v^* + \eta) - c^* \geq \beta\delta\left(\frac{\delta}{1-\delta}(v' + \eta) - c'\right)$, or $\eta \geq \frac{\delta}{1-\delta}v' - \frac{1}{1-\delta}v^* + \frac{1}{\beta\delta}c^* - c'$. The result follows.

Q.E.D.

References

Ainslie, G. (1975). "Specious Reward: A Behavioral Theory of Impulsiveness and Impulse Control." *Psychological Bulletin* 82(4): 463-496.

Ainslie, G. (1991). "Derivation of 'Rational' Economic Behavior from Hyperbolic Discount Curves." *American Economic Review* 81(2): 334-340.

Ainslie, G. (1992). *Picoeconomics: The strategic interaction of successive motivational states within the person*. Cambridge University Press; New York.

Ainslie, G. and N. Haslam (1992*a*). "Self-control." in *Choice over time*, G. Loewenstein and J. Elster, eds., Russell Sage Foundation, New York, NY: 177-209.

Ainslie, G. and N. Haslam (1992*b*). "Hyperbolic Discounting." in *Choice over time*, G. Loewenstein and J. Elster, eds., Russell Sage Foundation, New York: 57-92.

Akerlof, G. A. (1991). "Procrastination and Obedience." *American Economic Review* 81(2): 1-19.

Fischer, C. (1997). "Read This Paper Even Later: Procrastination with Time-Inconsistent Preferences." Dissertation, University of Michigan.

Goldman, S. M. (1979). "Intertemporally Inconsistent Preferences and the Rate of Consumption." *Econometrica* 47(3): 621-626.

Laibson, D. (1994). "Essays in Hyperbolic Discounting." Economics, MIT.

Laibson, D. (1995). "Hyperbolic Discount Functions, Undersaving, and Savings Policy." Mimeo, Harvard University.

Laibson, D. (1997). "Hyperbolic Discounting and Golden Eggs." *Quarterly Journal of Economics* 112, 443-477.

Laibson, D., A. Repetto, and J. Tobacman (1998). "Self-Control and Saving for Retirement." *Brookings Papers on Economic Activity*, forthcoming.

Loewenstein, G. and D. Prelec (1992). "Anomalies in Intertemporal Choice: Evidence and an Interpretation." *Quarterly Journal of Economics* 107(2): 573-597.

O'Donoghue, T. and Rabin, M. (forthcoming *a*). "Doing It Now or Later." *American Economic Review*, forthcoming, March 1999.

O'Donoghue, T. and Rabin, M. (forthcoming *b*). "Incentives for Procrastinators." *Quarterly Journal of Economics*, forthcoming.

O'Donoghue, T. and Rabin, M. (1998). "Procrastination in Preparing for Retirement." Mimeo, Cornell University and U.C. Berkeley.

Phelps, E. S. and R. A. Pollak (1968). "On Second-Best National Saving and Game-Equilibrium Growth." *Review of Economic Studies* 35: 185-199.

Pollak, R. A. (1968). "Consistent Planning." *Review of Economic Studies* 35: 201-208.

Prelec, D. (1989). "Decreasing Impatience: Definition and Consequences." Mimeo, Harvard Business School.

Strotz, R. H. (1956). "Myopia and Inconsistency in Dynamic Utility Maximization." *Review of Economic Studies* 23(3): 165-180.

Thaler, R. H. (1991). "Some Empirical Evidence on Dynamic Inconsistency." *Quasi Rational Economics*, Russell Sage Foundation: 127-133.

Thaler, R. and G. Loewenstein (1992). "Intertemporal Choice." in *The Winner's Curse: Paradoxes & Anomalies of Economic Life*, R. Thaler, ed., Free Press, New York: 92-106.

Voltaire, Jean Francois Marie Arouet de (1878). "L'art dramatique." in *Œuvres completes de Voltaire*, v. 17, Garnier Freres.