Procrastination, Self-Imposed Deadlines and Other Commitment Devices

Alberto Bisin
New York University and NBER*

Kyle Hyndman
UT Dallas†

May 8, 2014

Abstract

In this paper we model a decision maker who must exert costly effort to complete a single task by a fixed deadline. Effort costs evolve stochastically in continuous time. The decision maker optimally waits to exert effort until costs are less than a given threshold, the solution to an optimal stopping time problem. We derive the solution to this model for three cases: (1) time consistent decision makers, (2) naïve hyperbolic discounters and (3) sophisticated hyperbolic discounters. Sophisticated hyperbolic discounters behave *as if* they were time consistent but instead have a smaller reward for completing the task. Moreover, they will often self-impose a deadline to ensure early completion of the task. We also discuss other forms of commitment that (sophisticated) present-biased decision makers may use, as well as the demand for commitment in other models of decision making.

1 Introduction

In this paper, we are concerned with characterizing the behavior of decision makers who must exert immediate effort to complete a task that provides a delayed reward in an environment in which the cost of completing the task (i.e., effort) follows a stochastic process. We consider three types of decision makers: exponential, sophisticated hyperbolic discounters and naive hyperbolic discounters, where sophisticates are aware of their present-bias, but naifs are not.[1] We then want to study how certain commitment devices alter behavior and which, if

*Department of Economics, New York University, 19 W 4th Street, 6th Floor, New York, NY 10012. E-mail: alberto.bisin@nyu.edu, url: http://www.econ.nyu.edu/user/bisina.

†Jindal School of Management, University of Texas at Dallas, 800 W. Campbell Rd. (SM31), Richardson, TX 75080. E-mail: KyleB.Hyndman@utdallas.edu, url: http://www.hyndman-honhon.com/Main/Kyle_Hyndman.html.

[1] O'Donoghue and Rabin (1999b) contains a detailed discussion of sophisticates and naifs.

any, of the three types of decision makers would self-impose an external commitment device. Because of their prominence in the literature, our main focus is on self-imposed deadlines but we also investigate other types of commitment, such as making a fixed payment conditional on successfully completing the task or imposing a penalty for not completing the task, which is at the heart of many commitment mechanisms (e.g., the smoking cessation study of Giné, Karlan, and Zinman (2010)).

Our model focuses on the problem faced by a single decision maker who has to decide when to complete a task before some fixed, final deadline. Completing the task provides a delayed benefit but requires the immediate exertion of effort. We assume that effort is costly and evolves according to a continuous time stochastic process. Thus, the decision maker faces an optimal stopping time problem. To study the behavior of decision makers with a present-bias, we adapt the model of Harris and Laibson (2013) — an elegant generalization of the standard $\beta - \delta$ quasi-hyperbolic model of Phelps and Pollak (1968) and Laibson (1994) to continuous time. In their model, a decision maker is divided into a present-self and a future-self — the latter being present-biased. The present-self exercises control for a random amount of time with transitions occurring with a constant hazard; particular attention is paid to the limit case in which the present-self maintains control for a vanishingly small amount of time.

Our characterization of behavior when decision makers face a fixed deadline, T, yields three results. First, a decision maker with standard, exponential time preferences adopts a threshold rule whereby she completes the task if the cost of doing so is below the threshold, which increases as the deadline approaches. Second, a naïve hyperbolic discounter will always delay until the last possible moment and will then only do the task if the cost is not too high. Finally, and most interestingly, we show that a sophisticated hyperbolic discounter behaves *as if* she were an exponential discounter but faced a smaller value to completing the task. That is, the sophisticated agent adopts a threshold which increases as the final deadline approaches.

Given our characterization of optimal behavior when faced with a fixed deadline, T, we next ask whether any of the decision makers would self-impose a deadline $D < T$. Using results from the option pricing literature, we are able to show that for exponential decision makers, the value is increasing in the time from the deadline. From this, we conclude that an exponential discounter would never self-impose a deadline $D < T$. Similarly, since the naïve hyperbolic discounter mistakenly believes that her future self is exponential, she will also never self-impose a deadline.

Our main result from this analysis is to show that sophisticated hyperbolic decision makers may choose to self-impose a binding deadline $D < T$. The tension in this decision is, generally, between an immediate deadline or no deadline at all. On the one hand, sophisticates behave *as if* they were exponential but with a different value of completing the task. Thus, to preserve the option value of waiting for a lower cost, the latest possible deadline is attractive. However, from the perspective of the *ex ante* self, we show that there is a discontinuous benefit from completing the task immediately. If this benefit is high enough, they impose a deadline which ensures immediate completion; if not, they prefer the latest possible deadline in order to preserve the option value of waiting. We are able to prove this formally when the initial cost is known with certainty. When the initial costs are drawn from a continuous distribution, we are only able to show that both an immediate deadline and no deadline at all are local maxima. However, when the set of possible initial costs are discrete, intermediate deadlines may emerge. The trade-off at the heart of this decision is as follows: an early deadline encourages relatively low initial cost types to complete the task immediately; however, it destroys the option value to complete the task later for relatively high initial cost types. An intermediate deadline seeks to balance these two opposite effects.

The rest of the paper proceeds as follows. In the next section we briefly review the related literature. Section 3 outlines the model and provides the main theoretical results. In Section 4 we consider commitment by methods other than self-imposed deadlines. The two that we consider, motivated by Trope and Fishbach (2000), are making a fixed payment conditional on task completion and imposing a cost for not completing the task. In Section 5 we return to self-imposed deadlines but consider other models (one based on misperceptions and the other based on temptations à la Gul and Pesendorfer (2001, 2004)) which do not involve present-bias but in which a decision maker may, nonetheless, prefer to commit. Surprisingly, our results suggest that a decision maker with Gul and Pesendorfer (2001, 2004) preferences will never self-impose a binding deadline. Finally, Section 6 provides some concluding remarks.

2 Related Literature

There is a large literature studying time-inconsistent behavior in general and present-bias in particular dating to Strotz (1956). The quasi-hyperbolic model (i.e., the β/δ model) on which much of the subsequent work, including the present one, is based was first studied by Phelps and Pollak (1968). Subsequent work by Laibson (1994, 1997) and O'Donoghue and Rabin (1999a,b) developed the theory further and began introduced formally introduced the notion of commitment or other incentives to overcome problems due to procrastination.

In recent years a number of lab and field experiments have specifically looked at the role of commitment devices in overcoming self-control problems. There have been several papers which study the role of deadlines (and the willingness to commit to them). The seminal paper by Ariely and Wertenbroch (2002) showed that (i) a substantial fraction of people are willing to self-impose binding deadlines on themselves and that (ii) such deadlines lead to increased performance and less delay. However, later work by Burger, Charness, and Lynham (2011) failed to find that deadlines increased performance. Bisin and Hyndman (2014) showed that there was a strong demand for commitment via self-imposed deadlines but that, in contrast to Ariely and Wertenbroch (2002), they did not lead to improved performance. Comparing the estimates of present-bias obtained in their three-task and one-task treatments, Bisin and Hyndman (2014) also show that, despite a large fraction of apparently present-biased decision makers, repeated similar tasks may activate internal self-control, thereby rendering deadlines superfluous.

Demand for commitment has also been shown in other contexts. For example, Houser, Schunk, Winter, and Xiao (2010) showed that subjects are willing to commit to reduce their choice set — in particular, to remove internet access so that they cannot succumb to temptation to surf the internet. Trope and Fishbach (2000) examine students' demand for two commitment devices: making a fixed payment conditional on success, and imposing a cost for not completing the task. In both cases, there is strong demand for commitment.

Even further into the field, Giné, Karlan, and Zinman (2010) allowed smokers who expressed an interest in quitting to make regular deposits into a bank account, with the money forfeited if they failed a test for nicotine after six months. They found that smokers in the commitment group were more likely to pass another surprise nicotine test even after 12 months. Other field studies have shown that certain commitment products may lead to higher savings (e.g., Thaler and Benartzi (2004), Ashraf, Karlan, and Yin (2006) and Duflo, Kremer, and Robinson (2011), etc).[2]

As noted in the introduction, our theoretical framework is based upon Harris and Laibson (2013). Several other papers have adopted this framework to address similar questions. In particular, Hsiaw (2013) applies this framework to study the interaction between goals and self-control, while Grenadier and Wang (2007) study the investment decisions of quasi-hyperbolic decision makers. These papers consider an infinite horizon model in which the decision maker can pursue an investment opportunity (by paying a fixed cost), but whose benefits evolve stochastically. The latter paper shows that naive decision makers will invest earlier than a time-consistent decision maker; moreover, a sophisticated decision maker will invest earlier still. In our model, it is the costs that evolve stochastically, while the benefit is fixed. Moreover, because of our emphasis on deadlines, we consider finite time horizons. In contrast to Grenadier and Wang (2007), we also show that sophisticated decision makers

[2]For a more detailed summary, the reader is referred to Bryan, Karlan, and Nelson (2010), who provide a comprehensive survey of the theoretical and experimental literature on self-control and commitment.

behave closer to the time-consistent benchmark than do naive decision makers.

3 The Model

A decision maker is faced with a task that needs to be completed by time T. If completed at or before time T, a task will pay a reward of $V > 0$. If the task is completed after its deadline no reward will be given. We assume that there is some delay in the payment upon completion of the task. That is, if the task is completed at time t, then the payment V is made at time $t + \Delta$, where $\Delta > 0$ is a constant.

Suppose also that the cost, x, of completing the task follows a geometric Brownian motion. That is:

$$dx = \sigma x \cdot dz \tag{1}$$

where z is a standard Weiner process and $\sigma > 0$ measures the standard deviation in costs per unit of time. We now proceed to the write down the general case for an exponential discounter. Later we will formulate and solve the problem of a hyperbolic discounter.

3.1 The Exponential Case

Let $W(x,t) = \sup_{0 \le \tau \le T-t} \mathbb{E}_{x,t}(e^{-\rho\tau} \max\{\bar{V} - x, 0\})$ denote the value of the task when the current cost is x, the current time is t (so that time to the deadline is $t^d := T - t$) and $\bar{V} := e^{-\rho\Delta} V$. The parameter $\rho > 0$ captures the time preferences of the decision maker.

This basic problem is formally equivalent to an American put option where the strike price is $\bar{V}$ and the current price of the underlying security is $x(t)$. Therefore, all of the tools and results derived from this literature will guide us here. In particular, the solution to this problem leads to a threshold function, $\bar{x}(t)$, which is decreasing the further from the deadline, such that for $x(t) \le \bar{x}(t)$ the decision maker will complete the task, while when the opposite inequality holds, the agent prefers to wait. It is well known that in the continuation region $W(x,t)$ is the solution to the following free boundary problem:

$$\begin{aligned}
\rho W(x,t) &= \frac{1}{2}\sigma^2 x^2 W_{xx}(x,t) + W_t(x,t) && (2)\\
W(\bar{x}(t),t) &= \bar{V} - \bar{x}(t) && (3)\\
W_x(\bar{x}(t),t) &= -1 && (4)\\
W(x,T) &= \max\{\bar{V} - x, 0\} && (5)
\end{aligned}$$

In the language of Dixit and Pindyck (1994), (3) represents the value matching condition and the threshold, while (4) is the smooth pasting condition at the boundary and (5) is the terminal condition describing optimal behaviour at the deadline.

3.1.1 Comparative Statics

The solution to this problem does not admit an explicit, closed form solution and one must instead rely on numerical techniques or analytic approximations. However, many properties of the value function and the threshold cost function are known. We now prove a few results and also provide some intuition for them. An intuitive discussion can be found in Hull (2005), while Peskir and Shiryaev (2006) contains an advanced treatment.

The main comparative statics are summarized as follows:

Proposition 1. *The following is known about the optimal threshold, $\bar{x}(t)$;*

(a) $\bar{x}(t)$ is increasing in t;

(b) $\bar{x}(t)$ is decreasing in σ; and

(c) A change in ρ has an ambiguous effect on $\bar{x}(t)$

Proof. We formally prove (a) and provide an intuitive discussion for (b) and (c).

(a) Notice that the gain function $\max\{\bar{V} - x, 0\}$ does not depend on time. Therefore, we may conclude that $W(x, t)$ is decreasing in t for each $x \in \mathbb{R}_{++}$. Suppose that $x > \bar{x}(t)$ for some t so that $W(x, t) - \max\{\bar{V} - x, 0\} > 0$. Now take any $t' \in [0, t)$. It follows that $W(x, t') - \max\{\bar{V} - x, 0\} \geq W(x, t) - \max\{\bar{V} - x, 0\} > 0$, which implies that $x > \bar{x}(t')$, which in turn implies that $\bar{x}(t)$ is increasing in t.

(b) The higher is volatility, σ, the more likely it is that there will be wide swings in the cost of task completion. However, the decision maker is shielded from cost increases (he can always *not* exercise the option) and benefits from cost decreases.

(c) An increase in ρ has a negative effect on the option value of waiting, which all else equal, would increase the threshold cost realization (see, e.g., Peskir and Shiryaev (2006)). However, since the reward is given by $\bar{V} = e^{-\rho\Delta}V$, an increase in ρ lowers the reward from completing the task, which, all else equal, would lower the cutoff, $\bar{x}(t)$, which is monotone increasing in the reward. Thus, the net effect is ambiguous, but as the time between task completion and payment becomes vanishingly small, the former effect will dominate.

□

3.1.2 The Optimal Deadline

Here we show that the decision maker prefers a later deadline to an earlier one. Let ω_τ denote the optimal policy given a deadline of τ and let $W^\tau(x,t;\omega)$ denote the value obtained by the decision maker at time t with cost x, facing a deadline of τ and following the policy ω. With this notation, we are now ready to demonstrate:

Proposition 2. *An exponential discounter prefers later to earlier deadlines.*

Proof. Consider two different decision problems, differentiated only by their deadlines. Assume that $\tau < \tau'$. We know that $W^{\tau'}(x,t;\omega_{\tau'}) \geq W^{\tau'}(x,t;\omega_\tau) \geq W^\tau(x,t;\omega_\tau)$. The first inequality comes from the optimality of the policy $\omega_{\tau'}$ when faced with deadline τ'. The latter inequality comes from the fact that exactly the same outcomes arise since the same policy is being implemented in the two situations; however, in the former case, the accumulated rewards from completing tasks are weakly higher because $\tau' > \tau$. Therefore, we have shown that the decision maker prefers later deadlines to earlier ones. □

3.2 The Hyperbolic Case

We now wish to move away from time consistent decision makers and move to a world of time inconsistent behaviour. We follow Rabin and assume that decision makers may either be naïve or sophisticated. Naïfs suffer from time inconsistency, but are themselves unaware of it. That is, at time t a naïf discounts a payment received in time $t+s$ by $\beta e^{-\rho s}$; however, she assumes that her time s self discounts according to $e^{-\rho(t'-s)}$. On the other hand, sophisticated decision makers recognize that they have self-control problems. Therefore, the decision problem is often modeled as a dynamic game with different versions of one's self at each period. In this way, a sophisticate anticipates what her future self will actually chose, but still suffers from a present bias.

We adopt the analytical framework of Harris and Laibson (2013) in order to model hyperbolic discounting in continuous time. In particular, a decision maker born at time t can be divided into two *selves*: the present self, lasting from t to $t+\tau_t$, and the future self, lasting from $t+\tau_t$ to ∞. The length of time that the present self exercises control is then τ_t, which is stochastic and, as in Harris and Laibson (2013), is exponentially distributed with parameter $\lambda \geq 0$. The discount function of a decision maker born at time t can then be summarised as follows:

$$D_t(t') = \begin{cases} e^{-\rho(t'-t)} & \text{if } t' \in [t, t+\tau_t) \\ \beta e^{-\rho(t'-t)} & \text{if } t' \in [s+\tau_t, \infty) \end{cases}$$

where $\rho > 0$ and $\beta \in (0, 1]$.

3.2.1 Naïve Decision Makers

We first consider a naïve decision maker. As with the exponential case above, there will be a threshold cost realisation below which it is optimal to complete the task, and above which it is optimal to wait. Let $W^n(x,t)$ denote the value function of a naïf and $W^e(x,t)$ denote the value function of the exponential discounter. For any pair (x,t), in an interval of length dt, the value function, $W^n(x,t)$, must solve:

$$W^n(x,t) = \max\left\{\bar{V}_\lambda - x\ ,\quad e^{-\rho dt}\mathbb{E}\left[e^{-\lambda dt}W^n(x+dx,t+dt) \right.\right. \\ \left.\left. +(1-e^{-\lambda dt})\beta W^e(x+dx,t+dt)\right]\right\} \tag{6}$$

where $\bar{V}_\lambda = e^{-\lambda\Delta}\bar{V} + (1-e^{-\lambda\Delta})\beta\bar{V}$ is the expected discounted payment from completing the task immediately.[3] If the decision maker transforms into her future self, which occurs with probability $1 - e^{-\lambda dt}$, then she discounts the future by the extra factor β, but she anticipates that she will behave as the exponential discounter would — hence our inclusion of $W^e(x+dx,t+dt)$.

Now suppose that we are in the region for which it is optimal to wait. Multiply both sides of (6) by $e^{\rho dt}$ and subtract $W^n(x,t)$ to obtain:

$$\begin{aligned}(e^{\rho dt}-1)W^n(x,t) &= e^{-\lambda dt}\mathbb{E}[W^n(x+dx,t+dt) - W^n(x,t)] \\ &+ (1-e^{-\lambda dt})\mathbb{E}[\beta W^e(x+dx,t+dt) - W^n(x,t)]\end{aligned}$$

which, upon noting that $e^{\rho dt} \approx 1 + \rho dt$, $e^{-\lambda dt} \approx 1 - \lambda dt$ and applying Ito's Lemma, can be simplified to:

$$\begin{aligned}\rho dt W^n(x,t) &= (1-\lambda dt)\mathbb{E}[W^n_x dx + \frac{1}{2}W^n_{xx}dx^2 + W^n_t dt] \\ &+ \lambda dt[\beta W^e(x+dx,t+dt) - W^n(x+dx,t+dt)]\end{aligned}$$

Diving through by dt and taking the limit as $dt \to 0$, we arrive at:

$$\rho W^n(x,t) = \frac{1}{2}\sigma^2 x^2 W^n_{xx}(x,t) + W^n_t(x,t) + \lambda[\beta W^e(x,t) - W^n(x,t)] \tag{7}$$

[3]Since the subject must wait for $\Delta > 0$ units of time from completion to payment, there is a chance that the decision maker will transform into her future self, and so, from today's perspective, will discount with parameter β. Of course, as $\lambda \to \infty$, $\bar{V}_\lambda \to \beta\bar{V}$.

along with the following optimality conditions:

$$W^n(\bar{x}^n(t),t) = \bar{V}_\lambda - \bar{x}^n(t) \tag{8}$$
$$W^n_x(\bar{x}^n(t),t) = -1 \tag{9}$$
$$W^n(x,T) = \max\{\bar{V}_\lambda - x, 0\} \tag{10}$$

It is intuitively obvious that $\bar{x}^n(t) \leq \bar{x}^e(t)$ for all t and for all $\lambda > 0$. This is so for two reasons. First, the reward, $\bar{V}_\lambda \leq \bar{V}$. Therefore, there are simply fewer cost realisations for which the naïve decision maker can profitably complete the task. Second, for all $\lambda > 0$, there is always a positive probability that the current self will relinquish control over task completion to her future self. Importantly, however, the current self believes that her future self will behave as an exponential discounter would. This means that she believes that her future self will be more likely to complete the task, making waiting optimal.

An interesting thing happens in the limit as $\lambda \to \infty$; namely, $W^n(x,t) = \beta W^e(x,t)$ in the waiting region for the naïve decision maker. That is,

$$W^n(x,t) = \max\{\beta\bar{V} - x, \beta W^e(x,t)\}$$

and so the decision maker will complete the task at time t provided that $x(t) \leq \beta[\bar{V} - W^e(x,t)]$. Now assume that $0 < x(t) \leq \bar{x}^e(t)$, so that the exponential discounter would rationally choose to complete the task. In this case, the hyperbolic discounter will complete the task if and only if $x(t) \leq \beta[\bar{V} - (\bar{V} - x(t))] = \beta x(t)$. Since $\beta < 1$, this can only be satisfied provided $x(t) = 0$. Therefore, we have proven that:

Proposition 3. *Unless $x(t) = 0$, the naïve decision maker will complete the task only at time T and only if $x(T) \leq \beta\bar{V}$.*

3.2.2 Sophisticated Decision Makers

Turn now to sophisticated decision makers. The approach is similar to the previous subsection, though there are a few conceptual differences due to sophistication. In particular, the current self anticipates the actions that will be taken by her future self and incorporates these decisions into the value function. Therefore, we can define a current value function $W^s(x,t)$ ("s" for sophisticated) and a continuation value function $w^c(x,t)$ ("c" for continuation).

As above, there will be a threshold cost $\bar{x}^s(t)$ below which it is optimal to complete the task, and above which it is optimal to wait. In the waiting region, the continuation value function can be expressed as the following partial differential equation:

$$\rho w^c(x,t) = \frac{1}{2}\sigma^2 x^2 w^c_{xx}(x,t) + w^c_t(x,t) \tag{11}$$

which has the same form as in (2), but does not also come explicitly endowed with boundary conditions, nor value-matching or smooth-pasting conditions. This is because w^c is not being maximised, but is instead being evaluated according to the policy function for W^s, which we now seek to derive. Outside of the waiting region, we also have $w^c(x,t) = \bar{V} - x$.

Consider now $W^s(x,t)$. It may be expressed as:

$$\begin{aligned} W^s(x,t) = \max\big\{\bar{V}_\lambda - x\ ,\quad & e^{-\rho dt}\mathbb{E}\left[e^{-\lambda dt}W^s(x+dx,t+dt)\right. \\ & \left. +(1-e^{-\lambda dt})\beta w^c(x+dx,t+dt)\right]\big\} \end{aligned}. \tag{12}$$

Running through the same steps as above, we can re-write (12) in the waiting region as:

$$\rho W^s(x,t) = \frac{1}{2}\sigma^2 x^2 W^s_{xx}(x,t) + W^s_t(x,t) + \lambda[\beta w^c(x,t) - W^s(x,t)]. \tag{13}$$

Of course, optimality requires the following extra conditions:

$$W^s(\bar{x}^s(t),t) = \bar{V}_\lambda - \bar{x}^s(t) \tag{14}$$

$$W^s_x(\bar{x}^s(t),t) = -1 \tag{15}$$

$$W^s(x,T) = \max\{\bar{V}_\lambda - x, 0\}. \tag{16}$$

Therefore, for each $\lambda > 0$, the equilibrium solution for the hyperbolic discounter is given by a solution to (11) and (13) – (16). Note also that as $\lambda \to \infty$, we have that $\beta w^c(x,t) = W^s(x,t)$ in the waiting region. In fact, this key insight allows us to fully characterise the solution. Since $w^c(x,t)$ satisfies the partial differential equation (11) and since $W^s(x,t) = \beta w^c(x,t)$, implies that $W^s(x,t)$ must also satisfy (11). That is, $W^s(x,t)$ satisfies: $\rho W^s(x,t) = \frac{1}{2}\sigma^2 x^2 W^s_{xx}(x,t) + W^s_t(x,t)$, which together with the boundary conditions (14) – (16) is identical to the decision problem faced by an exponential discounter in which the reward for completing the task is $\beta\bar{V}$ instead of $\bar{V}$. Formally stated,

Proposition 4. *In the limit as $\lambda \to \infty$, the sophisticated hyperbolic discounter behaves as if he were an exponential discounter in which the reward for completing the task is $\beta\bar{V}$.*

3.2.3 The Optimal Deadline

It is obvious that the naïve decision maker will not self-impose a deadline since she believes she will behave like an exponential decision maker in the future, and it was shown in Proposition 2 that exponential decision makers prefer later deadlines. However, this need not be the case for a sophisticated decision maker since the present self knows that the future self will not complete the task as soon as she would like. The only way for the sophisticated decision maker to influence when the task gets completed is to set an earlier deadline.

We begin with the case in which the decision maker *knows with certainty* what the cost realisation will be at time 0 (*i.e.,* $x(0)$ is known). We claim the following:

Proposition 5. *There exists some initial cost realisation, $\hat{x}$, such that if $x(0) \leq \hat{x}$, the decision maker sets any deadline τ such that $x(0) \leq \bar{x}^s(0, \tau)$. Otherwise, the decision maker sets the longest possible deadline.*

Proof. Observe that if the decision maker completes the task at time 0, then her utility is simply $\beta(\bar{V} - x(0))$. On the other hand, if she does not immediately complete the task, then her expected utility is $\beta w^c(x(0), 0; \tau)$. From Proposition 4, we know that $\beta w^c(x(0), 0; \tau) = W^s(x(0), 0; \tau)$ and that W^s is the value function of an exponential decision maker with a reward for completion of $\beta\bar{V}$. Therefore, we know that W^s is monotonically increasing in τ.

It is also relatively easy to see that as $\tau \to \infty$, $W(x, 0; \tau) \to Ax^\alpha$ for $x > \bar{x}(0, \infty) = \frac{\alpha}{\alpha-1}\beta\bar{V}$, where α is the negative root of the fundamental quadratic: $\frac{\sigma^2}{2}\alpha(\alpha - 1) - \rho = 0$ and $A = -\frac{\bar{x}(0,\infty)^{1-\alpha}}{\alpha}$. Therefore, $\hat{x}$ is given by the unique solution to $Ax^\alpha = \beta(\bar{V} - x)$. Then, since the function $\bar{x}^s(0, \tau)$ is monotonic in τ, we can invert it to determine $\hat{\tau}$, where $\hat{\tau}$ is the *weakest* deadline such that the decision maker immediately completes the task at time 0. Of course, any deadline $\tau' < \hat{\tau}$ will also ensure immediate completion of the task and would also represent an optimal decision for $x \leq \hat{x}$. □

The intuition for this is as follows: the *ex ante* self, who, when considering whether or not to impose a deadline, discounts exponentially faces a tradeoff between two extremes. On the one hand, she knows that if the task is not completed immediately, it is optimal to have the latest possible deadline to preserve option value. This follows since, by Proposition 4, the sophisticated decision maker behaves as if she were exponential. On the other hand, however, there is a “discontinuous” benefit to completing the task immediately since the *ex ante* self values immediate task completion at $\beta(\bar{V} - x(0))$ but completion at time $t > 0$ at only $e^{-\rho t}(\beta\bar{V} - x(t))$. Therefore, if the commitment benefit of an immediate deadline is larger than the gain in option value from no deadline at all, the decision maker will impose an immediate deadline.

Now suppose that $x(0)$ is a random variable, the realisation of which is drawn from a continuous distribution function F such that $\hat{x} \in \text{supp}(F)$. Intuitively, for realisations $x \leq \hat{x}$, the decision maker would like to set a very strict deadline to ensure immediate completion of the task, while for realisations $x > \hat{x}$, the decision maker prefers no deadline at all.

A bit more formally, observe that we may write the *ex ante* expected utility of the decision maker as follows:

$$\bar{W}^s(\tau) = \mathbb{E}_x[\beta w^c(x,0;\tau)] = \int_0^{\bar{x}^s(0,\tau)} \beta(\bar{V}-x)dF(x) + \int_{\bar{x}^s(0,\tau)}^{\infty} \beta w^c(x,0;\tau)dF(x) \tag{17}$$

where τ is the deadline, $\bar{x}^s(0,\tau)$ is the optimal threshold at time 0, $w^c(x,0;\tau)$ denotes the continuation value of the decision maker. Importantly, notice that $\beta w^c(x,0;\tau) = W^s(x,0;\tau)$ and for $x \in (\bar{x}^s(0,\tau), \bar{x}^s(0,\tau)+\epsilon)$, $W^s(x,0;\tau) \approx \beta V - x < V - x$. Notice also that W^s and $\bar{x}^s$ are both continuous in τ, meaning that a maximum is sure to exist. Formally, using Leibniz Rule, after some simplifications, we obtain:

$$\frac{\partial \bar{W}^s(\tau)}{\partial \tau} = (1-\beta)\bar{x}^s(0,\tau)f(\bar{x}^s(0,\tau))\frac{\partial \bar{x}^s}{\partial \tau} + \int_{\bar{x}^s(0,\tau)}^{\infty} W^s_\tau(x,0;\tau)dF(x). \tag{18}$$

There are two opposing effects in (18). If τ is increased slightly, then the threshold for immediate task completion is lowered, leading to a utility loss of $(1-\beta)\bar{x}^s(0,\tau)f(\bar{x}^s(0,\tau))$ since some initial cost types no longer complete the task at time 0. On the other hand, by increasing τ, the decision maker is giving herself more opportunities to complete the task in the future. Thus she experiences a gain, which is represented by the integral. Carr, Jarrow, and Myneni (1992) has shown that $\frac{\partial \bar{x}^s}{\partial \tau}$ is negative and tends to $-\infty$ as $\tau \to 0$. This implies that an immediate deadline is a local maximum of $\bar{W}^s(\tau)$. Furthermore, since $\frac{\partial \bar{x}^s}{\partial \tau} \to 0$ as $\tau \to \infty$, while the integral remains positive, an infinite deadline will be another "local" maximum. However, without further assumptions, we cannot say which of these two, if any, will be the global maximum, or if there is an intermediate global maximum.

Given the inability above to formally show whether a sophisticated hyperbolic discounter will set a deadline when there is uncertainty over the initial cost realisation, we seek to answer this question numerically. Our task is greatly facilitated by Proposition 4, since it tells us that we need to solve for the value function of an exponential decision maker with a value of completing the task of $\beta\bar{V}$, which amounts to pricing an American put option with strike price $\beta\bar{V}$. To do this, we use the binomial tree method of Cox, Ross, and Rubinstein (1979) to solve for $P(x,\beta\bar{V},t)$ which is the price of an American put option with "asset price" (*i.e.,* cost of completing the task), x, "strike price" (*i.e.,* value of completing the task), $\beta\bar{V}$, and time until expiration t. Then, the *ex ante* value of completing the task at time 0 is:

$$W^s(0,x,\tau) = \begin{cases} P(x,\beta\bar{V},\tau), & \text{if } x > \bar{x}^s(0,\tau) \\ \beta(\bar{V}-x), & \text{if } x \le \bar{x}^s(0,\tau) \end{cases}$$

FIGURE 1: Ex Ante Expected Value as a Function of Deadlines (Known Initial Cost)

(A) $\beta = 0.3$: Latest Possible Deadline Optimal

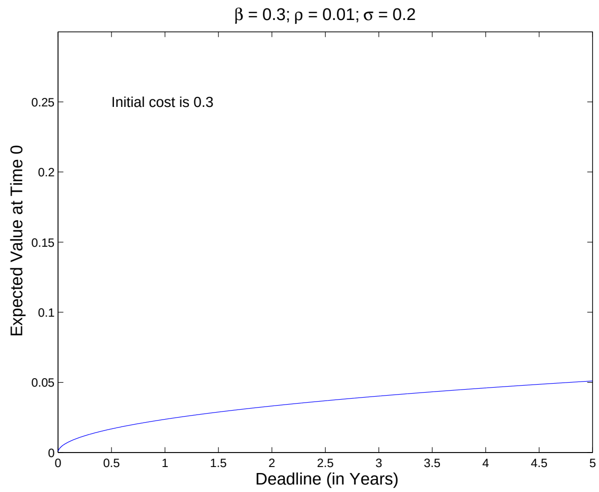

(B) $\beta = 0.6$: Immediate Deadline Optimal

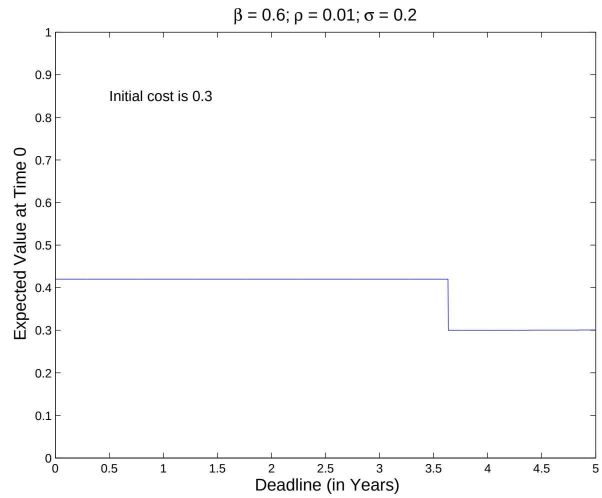

NOTE: On the horizontal axis is the deadline; i.e., $t = 0$ represents an immediate deadline, while $t = 2$ represents a deadline of two years. On the vertical axis is the *ex ante* expected value of the time 0 self.

Taking expectations over all possible initial values of x, we have $\bar{W}^s(\tau) = \int_x W^s(0, x, \tau) dF(x)$. That is, if the decision maker completes the task immediately at time 0, she obtains $\beta(V - x)$, while if she delays at time 0, then her value is the price of the corresponding option.

First suppose that the initial cost is known with certainty to be $x(0)$, and the completing the task yields a payment of $\bar{V} = 1$. The results are shown in Figure for the case of $\beta = 0.3$ and $\beta = 0.6$. As can be seen, when there is extreme present-bias (i.e., $\beta = 0.3$), the decision maker actually prefers the latest possible deadline (panel (A)). This is so because even an immediate deadline cannot compel the decision maker to complete the task immediately. Hence, given that she will delay, and because sophisticates behave as if they are exponential (Proposition 4), the latest possible deadline is optimal. In contrast, when present-bias is less severe, as long as the deadline is sufficiently tight, the decision maker will complete the task immediately, which is the preference of the *ex ante* self. However, if the deadline is too lax, then the decision maker will not complete the task immediately; hence the discontinuous drop in the expected payoff function in panel (B).

We next look at the case in which the initial cost is unknown, but drawn from a uniform distribution on $[0, 1]$. The results are shown in Figure 2. In panel (A), the decision maker is moderately present-biased, with $\beta = 0.7$, which leads her to prefer an immediate deadline over one even as long as five years into the future. In contrast, in panel (B), which depicts a mild present-bias of $\beta = 0.9$, the decision maker prefers the longest possible deadline. Further numerical analysis, not shown, suggests that the desire to self-impose a deadline is

FIGURE 2: Ex Ante Expected Value as a Function of Deadlines (Unknown Initial Cost)

(A) $\beta = 0.7$: Immediate Deadline Optimal (B) $\beta = 0.9$: Latest Possible Deadline Optimal

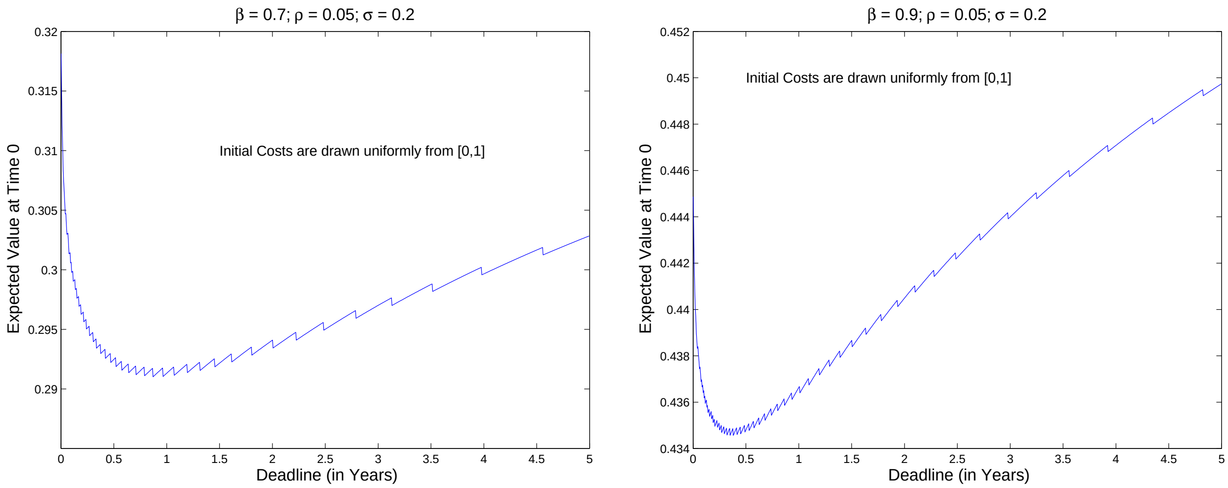

NOTE: On the horizontal axis is the deadline; i.e., $t = 0$ represents an immediate deadline, while $t = 2$ represents a deadline of two years. On the vertical axis is the *ex ante* expected value of the time 0 self.

decreasing in σ, which is intuitive since the option value to waiting is increasing in σ.

Although our results thus far suggest that either an immediate deadline (or at least one which ensures immediate task completion) or the latest possible deadline will be optimal, it turns out that intermediate deadlines may sometimes be optimal when the distribution of initial costs is discrete. The intuition is as follows. For a small change in the deadline, it may be that the set of initial cost realizations such that the decision maker completes the task immediately is unchanged. In this case, the first term in (18) will be 0, leaving only the positive second term. Hence, increasing the deadline is optimal. However, eventually the deadline will be increased sufficiently so that for one of the initial cost realizations will no longer complete the task immediately, leading to a discontinuous drop in the expected value function. From this point, as the deadline is increased further, the expected value function will, again, be increasing, until another initial cost realization no longer completes the task immediately, and so on. One example of an intermediate deadline is presented in Figure 3.

4 Commitment Via Other Methods

As we have seen, self-imposed deadlines are sometimes optimal. From an *ex ante* perspective, the sophisticated decision maker's future selves are not completing the task for high enough realisations of the cost. Therefore, by imposing a deadline, the *ex ante* self is increasing the thresholds used by the future selves, which is utility enhancing. However, as we have

Figure 3: Ex Ante Expected Value as a Function of Deadlines (Discrete Initial Cost Distribution)

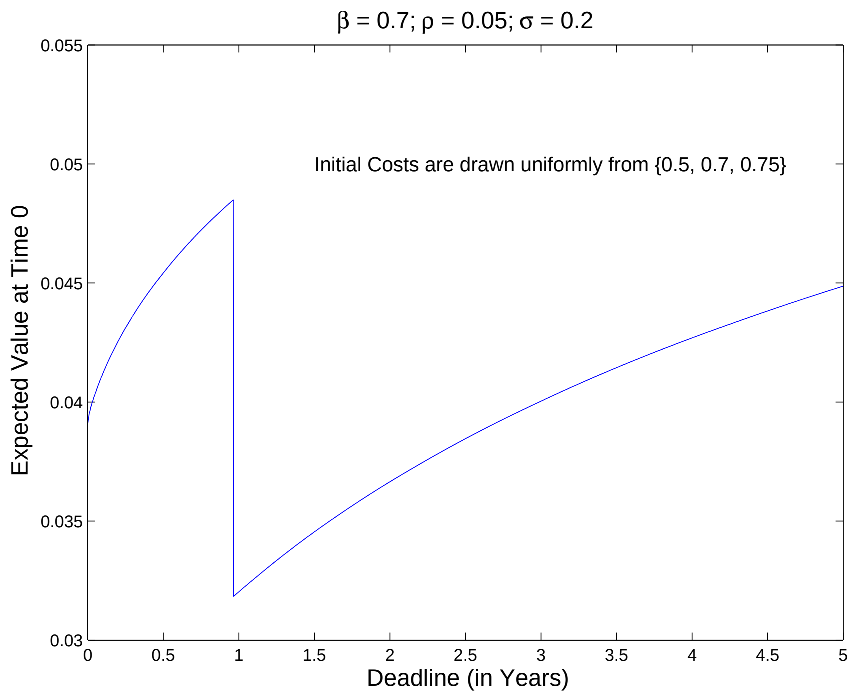

Note: On the horizontal axis is the deadline; i.e., $t = 0$ represents an immediate deadline, while $t = 2$ represents a deadline of two years. On the vertical axis is the *ex ante* expected value of the time 0 self.

said, deadlines are a blunt instrument and come at the cost of destroying the option value of completing the task at any time t beyond the deadline. It appears that, for moderate to high values of β, this cost dominates causing the *ex ante* self not to set a binding deadline. We now discuss a few alternative external commitment devices that increase the incentive to complete the task, but are different than a once-and-for-all deadline.

4.1 Making a Fixed Payment Conditional Upon Task Completion

Trope and Fishbach (2000) consider two slightly different commitment mechanisms for decision makers. In their first study, subjects were given a fixed payment (in terms of course grades) which was initially independent of whether or not they successfully completed a task (abstaining from glucose). However, the subjects were allowed to make all or part of the fixed payment conditional upon successfully completing the task. The authors found that subjects often do choose to make the fixed payment conditional. To see that this may be optimal formally, suppose that agents receive a fixed participation fee, v, independent of the completion of the task. Suppose also that we allow the agent to make the fee v conditional to the completion of the task. Under which conditions would the agent in fact choose to receive the fee conditionally to completion of the task? Let $W^s(x, t; \bar{V})$ denote the expected payoff of a sophisticated hyperbolic agent at time t when the payoff for completion is $\bar{V}$ and the cost of effort is x.[4] Let $\bar{x}^s(t; \bar{V})$ denote the associated optimal cutoff. The agent will make v conditional upon task completion when $\beta v + W^s(x, 0; \bar{V}) < W^s(x, 0; \bar{V} + v)$.

To see that this may work, suppose that the initial cost realisation, $x(0)$, is known with certainty and that $x(0) \in \big(\bar{x}^s(0, \bar{V}), \bar{x}^s(0, \bar{V} + v)\big)$. In this case, we have the following:

$$\beta v + W^s(x, 0; \bar{V}) = \beta v + \beta w^c(x(0), 0; \bar{V}) > \beta(v + \bar{V}) - x(0)$$

while,

$$W^s(x, 0; \bar{V} + v) = \beta(v + \bar{V} - x(0)).$$

Therefore, to the extent that $W^s(x, 0; \bar{V})$ is not *too much* greater than $\beta\bar{V} - x$, the sophisticated hyperbolic agent may prefer to make the fixed fee conditional upon task completion. In Figure 4 we provide simulation results showing that a sophisticated agent will sometimes prefer to make a fixed payment conditional on the completion of the task. For the parameters used in the simulation, when $\beta = 0.3$, for extremely short deadlines, the decision maker prefers to make the fixed payment conditional upon completion of the task, while for longer deadlines, taking the fixed payment up front is optimal, since with a longer deadline she is more likely to procrastinate and, hence, delay the time at which v is received. For $\beta \in \{0.4, 0.5, 0.6\}$, for sufficiently short deadlines, it does not matter whether the fixed payment is made conditional or not since, in either case, the sophisticated decision maker will immediately complete the task. For deadlines of intermediate length, making the fixed payment conditional is optimal because doing so ensures that the decision maker will immediately complete the task, whereas when v is taken unconditionally, she will delay task completion. For sufficiently long deadlines, even with v conditional upon task completion, the decision maker will delay. Therefore, it becomes optimal to take v up front. Finally, for $\beta \in \{0.7, 0.8, 0.9\}$, there is no difference between the two cases since in each case the decision maker immediately completes the task.

4.2 Imposing a Cost For Not Completing the Task

In another study, Trope and Fishbach (2000) let subjects impose a cost which is conditional upon *not* completing the task. We briefly demonstrate why subjects may find it optimal to impose such penalties and how it compares to the previous case.

The intuition can be seen in Figure 5. At the deadline, there is no difference between a

[4] This is exactly as defined by (12), we only make the dependence on $\bar{V}$ explicit.

FIGURE 4: Numerical Results; $r = 0.01$ & $\sigma = 0.2$; $x(0) = 0.3$, **known**. Comparison of a Fixed Payment vs. Making it Conditional

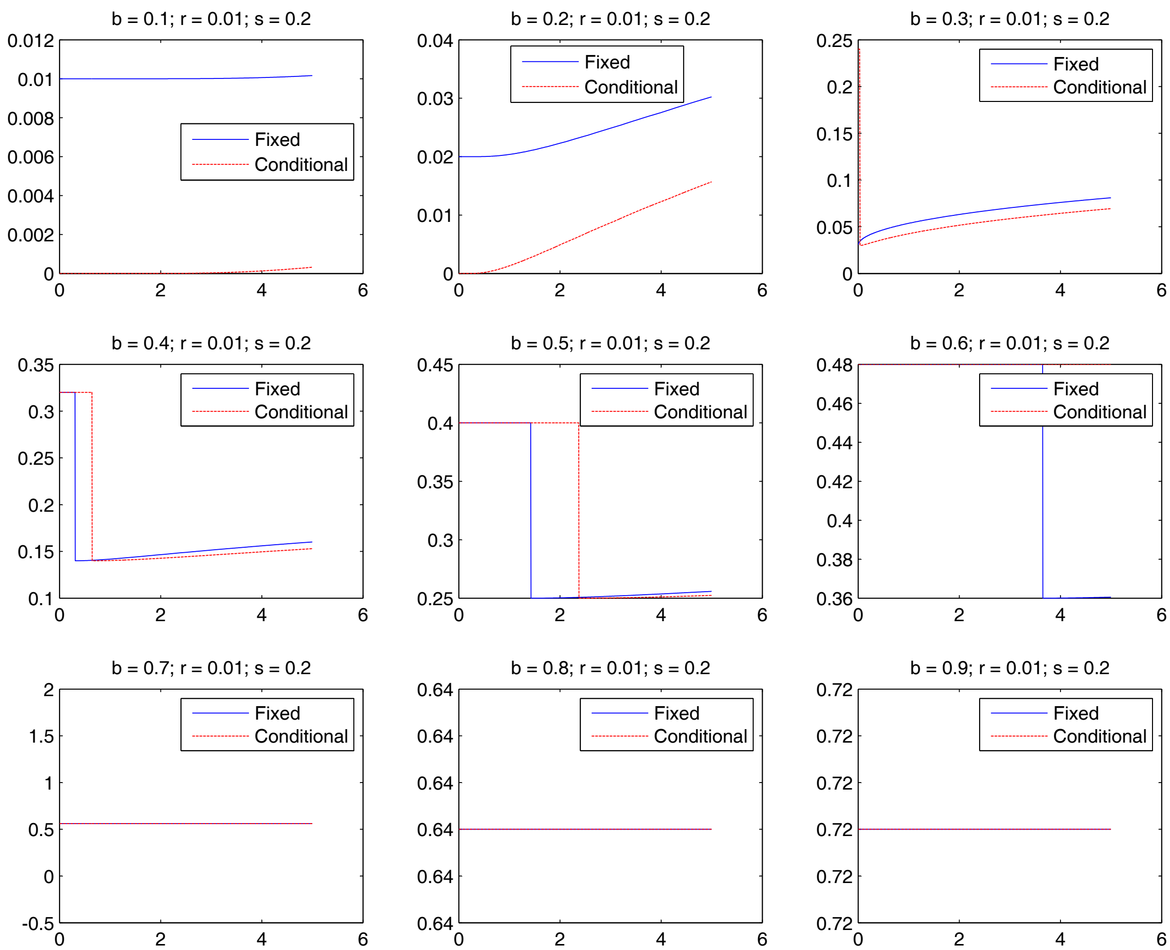

On the horizontal axis is the deadline; that is, $t = 0$ represents an immediate deadline, while $t = 2$ represents a deadline of two years. On the vertical axis is the *ex ante* expected value of the time 0 self.

penalty for not completing the task and an equivalent bonus conditional upon completing the task. Therefore, the threshold in each case will be $\beta(V + v)$. Notice that in each case, therefore, the threshold is higher than in the baseline. That is, decision makers will be more likely to complete the task. However, away from the deadline, the two cases begin to differ. When there is an additional reward for completing the task, by completing the task the decision maker will always get the bonus. Therefore, even very far away from the deadline, making a fixed payment conditional upon task completion still has a substantial effect on the threshold. In contrast, when there is a cost for not completing the task, the further away the decision maker is from the deadline the weaker is the effect of this cost. This is so for two reasons. First, the further away from the deadline, the lower is the expected

Figure 5: Comparing a Cost of Not Completing a Task With a Reward For Completing

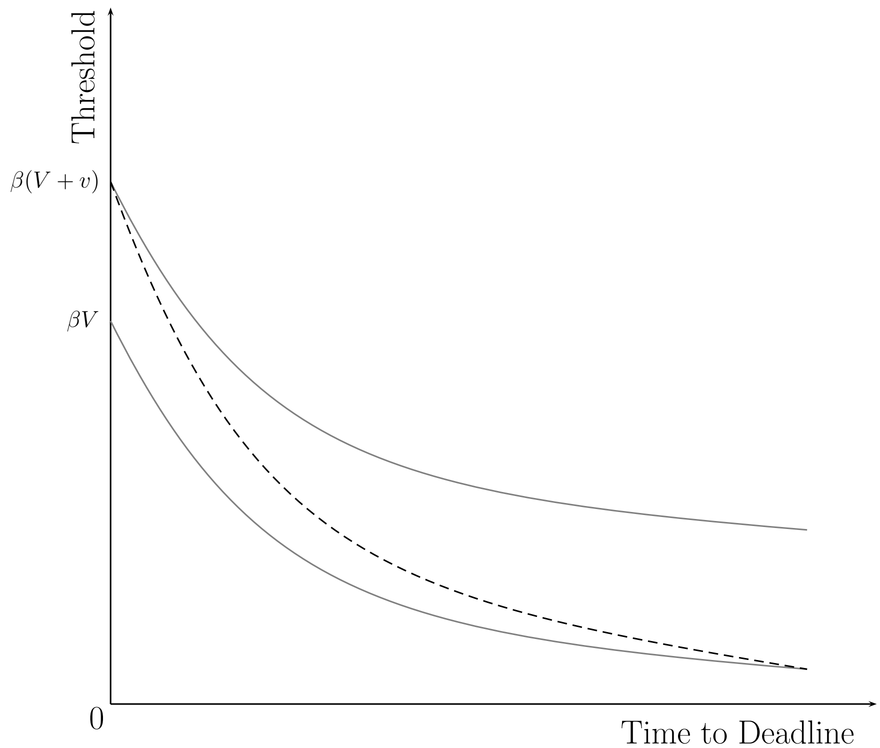

The lower gray line is the threshold for task completion for the baseline case in which there is a reward for completion of the task $V > 0$. The upper gray line represents the optimal threshold when an additional reward $v > 0$ is given for completing the task. Finally, the dashed, black line represents the case in which the reward for task completion is V, but a cost $c = \beta v$ is imposed if the task is not completed.

probability that the task *will not* be completed strictly before the deadline. Second, the further away from the deadline, the smaller is the present value of the cost of not completing the task. In the limit as the time to deadline approaches infinity, the decision maker will never face the cost of not completing the task before the deadline. Hence, the threshold approaches the baseline case in which no such penalty exists. Therefore, while imposing a penalty conditional upon not completing the task does impose some form of commitment, we would expect it to be weaker than the ability to convert an equivalent fixed payment to a reward which is conditional upon task completion.

5 Self-Imposed Deadlines in Other Models

Until now we have looked at the demand for commitment through the lens of present-biased time preferences, and showed that awareness of one's present bias is a necessary condition for a decision maker to self-impose a deadline. We now briefly discuss two other models and their predictions regarding whether or not such decision makers would self-impose deadlines.

5.1 Temptation & Self-Control

Miao (2008) adopts the temptation and self-control model of Gul and Pesendorfer (2001, 2004) to study the optimal exercise of options in a discrete time, infinite horizon setting. When the cost of exercising the option is immediate, while the benefit is delayed, Miao shows that agents are tempted to delay. One can go beyond this and ask whether the agent would bind herself by setting a deadline. In Appendix A, we consider the finite time version of Miao's model and prove that the value function is increasing in the time to complete the task. Therefore, an agent with Gul-Pesendforfer preferences will not self-impose deadlines.

5.2 Misperceptions

In our model, there is a tension between preserving option value (with a later deadline) and the commitment value of an earlier deadline. Depending on this trade-off a sophisticated, but present-biased, decision maker may self-impose a deadline. We now show that, instead of present-bias, misperceptions about costs may also create scope for deadlines. Along the lines of Brunnermeier, Papakonstantinou, and Parker (2008), suppose that decision makers misperceive the costs and/or benefits of completing the task. For example, suppose that the actual stochastic process for cost contains a drift term; that is, costs evolve according to:

$$dx = \alpha x dt + \sigma x \cdot dz$$

where $\alpha > 0$. Under this altered model, let $\beta = 1$ and redefine the "current" and "future" selves as follows: the current self is aware of the drift term, while the future self believes that $\alpha = 0$. All other aspects of our model remain unchanged. In this model, the future self mistakenly believes that there is an equal chance that the cost will go up or down, when it is actually more likely that costs will increase because of the positive drift. Because of this mistaken belief, she will be tempted to delay — hoping for a lower cost in the future.

Now consider a decision maker at time 0 who is aware of her future self's misperception. Because $\alpha > 0$, as time passes, it becomes increasingly less likely that the task will be completed. Therefore, the option value of waiting is less valuable, while the commitment benefit of a deadline remains, making it more likely that the time 0 self will set a deadline.

6 Conclusions

In this paper we examined the behaviour of decision makers who must decide whether and, if so, when to complete a task before some final deadline. Decision makers could have standard exponential time preferences or have a present bias for immediate gratification along the lines of Harris and Laibson (2013). Moreover, those with a present bias may be naïve or sophisticated. A naïf is unaware of her self-control problems, while a sophisticate is aware of her self-control problems. Having characterized behavior conditional on a fixed deadline, we showed that exponential and sophisticated decision makers are almost equivalent; in particular, sophisticated decision makers behave as if they were exponential but had a lower benefit from completing the task. Naive decision makers, on the other hand, are extreme procrastinators, completing the task only at the deadline an only if costs are low enough.

The fact that sophisticates and exponential decision makers behave so similarly gives rise to an identification problem if an observer only sees the timing of task completions. However, our results also showed that while exponential (and naive) decision makers would never self-impose a binding deadline, sophisticated decision makers may, in fact, optimally self-impose a binding deadline. The reason for doing so is that it encourages some low initial cost types to immediately complete the task, which, from an ex ante perspective, leads to a discontinuous benefit. Depending on its severity, the deadline also preserves some option value for higher initial cost types to delay completion to a later time when, hopefully, costs are lower. Thus, the way to identify (at least some) sophisticated decision makers is through their demand for commitment.

While our model focused on a single task, it can be extended to multiple tasks. However, care must be taken on this front since, without further alterations, this extended model predicts identical thresholds for all tasks. That is, as soon as the decision maker completes one task, she will also complete the rest. Instead, it seems likely that fatigue might set in. To capture this, one could include a discrete increase in cost upon the completion of a single task. Therefore, once a decision maker completes one task, her cost will increase, forcing her to "relax" and wait for a lower cost to arise in the future.[5] This could be why Ariely and Wertenbroch (2002) found that evenly spaced deadlines were the most effective.

With somewhat greater difficulty, our model could also be extended to continuous tasks along the lines of either Ariely and Wertenbroch (2002), where the reward is increasing in effort, or Burger, Charness, and Lynham (2011), where the task takes a fixed amount of time to be completed, but that this time can be divided over disjoint intervals. This adds an

[5] In a discrete time model more suitable for analyzing experimental data, Bisin and Hyndman (2014) solves a multiple-task problem and allows for fatigue (i.e., cost increases after task completion) or learning by doing (i.e., cost decreases after task completion). The results are more supportive of fatigue.

additional layer of complication since it introduces another state variable — the amount of exertion required to complete the task — turning the decision problem into a control problem. However, we conjecture that such an extended model could be parameterized to reconcile the cycles found in Study 1 of Burger, Charness, and Lynham (2011). We also conjecture that, as their experimental results suggest, deadlines would be ineffective at increasing task completion. The reason is that the main benefit of deadlines is to encourage immediate task completion for at least some initial cost types. However, as soon as the task is not immediately completed, the decision maker prefers the latest possible deadline for maximum flexibility. Since the task in Burger, Charness, and Lynham (2011) could not possibly be finished immediately, this suggests that intermediate deadlines would be of limited value.

References

Ariely, D., and K. Wertenbroch (2002): "Procrastination, Deadlines, and Performance: Self-Control by Precommitment," *Psychological Science*, 13(3), 219–224.

Ashraf, N., D. Karlan, and W. Yin (2006): "Tying Odysseus to the Mast: Evidence from a Commitment Savings Product in the Philippines," *Quarterly Journal of Economics*, 121(2), 635–672.

Bisin, A., and K. Hyndman (2014): "Present-Bias, Procrastination and Deadlines in a Field Experiment," NBER Working Paper 19874.

Brunnermeier, M., F. Papakonstantinou, and J. Parker (2008): "An Economic Model of the Planning Fallacy," Working Paper.

Bryan, G., D. Karlan, and S. Nelson (2010): "Commitment Devices," *Annual Review of Economics*, 2, 671–698.

Burger, N., G. Charness, and J. Lynham (2011): "Field and Online Experiments on Self-Control," *Journal of Economic Behavior & Organization*, 77, 393–404.

Carr, P., R. Jarrow, and R. Myneni (1992): "Alternative Characterizations of American Put Options," *Mathematical Finance*, 2, 78–106.

Cox, J. S., S. Ross, and M. Rubinstein (1979): "Option Pricing: A Simplified Approach," *Journal of Financial Economics*, 7, 229–263.

Dixit, A., and R. Pindyck (1994): *Investment Under Uncertainty*. Princeton University Press, Princeton, N.J.

Duflo, E., M. Kremer, and J. Robinson (2011): "Nudging Farmers to Use Fertilizer: Evidence From Kenya," *American Economic Review*, 101(6), 2350–2390.

Giné, X., D. Karlan, and J. Zinman (2010): "Put Your Money Where Your Butt Is: A Commitment Contract for Smoking Cessation," *American Economic Journal: Applied Economics*, 2(4), 213–235.

Grenadier, S. R., and N. Wang (2007): "Investment Under Uncertainty and Time-Inconsistent Preferences," *Journal of Financial Economics*, 84, 2–39.

Gul, F., and W. Pesendorfer (2001): "Temptation and Self-Control," *Econometrica*, 69, 1403–1436.

——— (2004): "Self-Control and the Theory of Consumption," *Econometrica*, 72, 119–158.

Harris, C., and D. Laibson (2013): "Instantaneous Gratification," *Quarterly Journal of Economics*, 128(1), 205–248.

Houser, D., D. Schunk, J. Winter, and E. Xiao (2010): "Temptation and Commitment in the Laboratory," working paper.

Hsiaw, A. (2013): "Goal-Setting and Self-Control," *Journal of Economic Theory*, 148(2), 601–626.

Hull, J. C. (2005): *Options, Futures & Other Derivatives*. Prentice Hall, Upper Saddle River, N.J., 6th edn.

Laibson, D. (1994): "Essays in Hyperbolic Discounting," Ph.D. thesis, Massachusetts Institute of Technology.

——— (1997): "Golden Eggs and Hyperbolic Discounting," *Quarterly Journal of Economics*, 113, 443–477.

Miao, J. (2008): "Option Exercise With Temptation," *Economic Theory*, 34, 473–501.

O'Donoghue, T., and M. Rabin (1999a): "Doing It Now or Later," *American Economic Review*, 89(1), 103–124.

——— (1999b): "Incentives for Procrastinators," *Quarterly Journal of Economics*, 114, 769–816.

Peskir, G., and A. Shiryaev (2006): *Optimal Stopping and Free Boundary Problems*. Birkhäuser Verlag, Basel.

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

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

Thaler, R. H., and S. Benartzi (2004): "Save More Tomorrow: Using Behavioral Economics to Increase Employee Saving," *Journal of Political Economy*, 112(1), S164–187.

Trope, Y., and A. Fishbach (2000): "Counteractive Self-Control in Overcoming Temptation," *Journal of Personality and Social Psychology*, 79(4), 493–506.

A Temptation and Self-Control

In this appendix we discuss the work of Miao (2008) and formally show that the a decision maker with Gul-Pesendorfer preferences will never self-impose a binding deadline. Recall that if the agent faces a choice set B_t when there are t periods remaining and W_t is the agent's intertemporal utility, then self-control preferences à la Gul-Pesendorfer are given by:

$$W_t(B_t) = \max_{c_t \in B_t} \{u(c_t) + \delta \mathbb{E}[W_{t-1}(B_{t-1})] + v_t(c_t)\} - \max_{c_t \in B_t} v_t(c_t)$$

where $B_t = \{0, 1\}$ provided that for all periods $n > t$, $c_n = 0$; that is, the agent can either complete the task or wait, and once she has completed the task, the decision problem ends and the agent receives the appropriate payoffs. Let c_t^* denote the optimal choice, then an agent with such preferences suffers a utility loss due to temptation of $v_t(c_t^*) - \max_{c \in B_t} v_t(c)$. It is this utility loss due to temptation which causes procrastination; in particular, the temptation to delay exerting costly effort. Miao (2008) then specialises to stopping time problems and considers three cases: immediate costs, immediate rewards and both immediate costs and rewards, and the reader is referred to his paper (specifically, Section 3.1) for more details. The case that is relevant for us is that of immediate costs. While Miao considers an infinite horizon problem, his model is easily adapted to a finite horizon setting. We also adapt his model to make the reward from task completion known, but the cost of completion stochastic. In no way does this change the results. Denote the value function when there are t periods remaining as:

$$\begin{aligned} W_t(x) &= \max\{\delta V - (1+\gamma)x, \delta \int W_{t-1}(x')dF(x')\} - \gamma \max\{0, -x\} \\ &= \max\{\delta V - (1+\gamma)x, \delta \int W_{t-1}(x')dF(x')\}, \end{aligned}$$

where V is the benefit from completing the task, x is the stochastic realisation of the cost of task completion, and $F(\cdot)$ is the distribution function from which the cost of task completion, x is drawn. The second equality follows from the fact that $x \geq 0$; therefore, $\max\{0, -x\} = 0$. Of course, notice that $W_0(x) \equiv 0$. Since the benefit of completing the task is delayed by one period, we must discount the reward - hence the appearance of the term δV; on the other hand, the cost of task completion cost, denoted by x, is stochastic. Finally, γx is the cost of exercising self-control and immediately completing the task.

We claim the following:

Proposition 6. *$W_{t+1}(x) \geq W_t(x)$ for all t and x.*

Proof. The proof is by induction. Obviously, since $W_0(x) \equiv 0$, and $W_1(x) = \delta V - (1+\gamma)x$ for $x < \frac{\delta V}{1+\gamma}$ and zero otherwise, the result is true for $t = 0$. Next suppose that the result is true for all $t = 0, 1, \ldots, n$ We now show that $W_{n+1}(x) \geq W_n(x)$. Observe that:

$$\begin{aligned} W_{n+1}(x) &= \max\{\delta V - (1+\gamma)x, \delta \textstyle\int W_n(x')dF(x')\} \\ &\geq \max\{\delta V - (1+\gamma)x, \delta \textstyle\int W_{n-1}(x')dF(x')\} \\ &= W_n(x) \end{aligned}$$

where the inequality follows from our induction hypothesis that $W_n(x) \geq W_{n-1}(x)$. □

Of course, while Proposition 6 shows that a decision maker with such preferences prefers the latest possible deadline, that is not to say that she will not procrastinate. In particular, one can easily show that the threshold cost of task completion is decreasing in γ, which measures the cost of self-control.

B The Finite λ Case

Our discussion has concerned itself with the limiting case of $\lambda \to \infty$. However, one may also be interested in the finite λ case where the present self can expect to exercise control for a measurable time interval. In particular, the issue of the agent's willingness to self-impose a binding deadline arises again.

We claim that our model is continuous in λ in the following sense. Rewrite (12):

$$W^s(x,t) = \max\left\{\bar{V}_\lambda - x\ , \quad e^{-\rho dt}\mathbb{E}\left[e^{-\lambda dt}W^s(x+dx,t+dt) \right.\right. \\ \left.\left. +(1-e^{-\lambda dt})\beta w^c(x+dx,t+dt)\right]\right\} .$$

To see this, note that a sophisticated decision maker is still tempted to delay because with probability $1 - e^{-\lambda dt}$, the agent will relinquish control of the problem to his future self. Importantly, however, if he believes that the future self is highly likely to complete the task in the next instance, then waiting leads to a small delay cost (by waiting $\bar{V}_\lambda$ is received after a length of time dt, rather than immediately). However, if the future self completes the task, there is a savings of $(1-\beta)x$. Therefore, the higher is λ, the higher is the probability of the current self relinquishing control to his future self, and so the greater is the temptation to delay. That is, the threshold for task completion, $\bar{x}(\lambda)$, is decreasing in λ.

Consider next the incentive to impose a deadline. The time 0 self wants to impose a deadline because he knows that when it comes time to actually exert effort, his future selves will not do so optimally, but will instead procrastinate. Therefore, a deadline will increase the threshold and force his future selves to complete the task for higher cost realisations. As we have seen, imposing a deadline is costly because it destroys the option value of waiting for a lower cost realisation in the future. From the perspective of the time 0 self, since $\bar{x}(\lambda)$ is decreasing in λ, the option value of waiting is also decreasing in λ, while the benefit of completing the task sooner remains unchanged. Therefore, as λ increases, the sophisticated hyperbolic agent will be more likely to impose a binding deadline.