

Procrastination, Self-Imposed Deadlines and Other Commitment Devices

ALBERTO BISIN*

NEW YORK UNIVERSITY

KYLE HYNDMAN†

SOUTHERN METHODIST UNIVERSITY

MARCH 16, 2009

Preliminary and Incomplete

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 will then optimally wait 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. We show that sophisticated decision makers will often self-impose a deadline to ensure early completion of the task. Other forms of commitment are also discussed.

*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>.

†Department of Economics, Southern Methodist University, 3300 Dyer Street, Suite 301R, Dallas, TX 75275. E-mail: hyndman@smu.edu, url: <http://faculty.smu.edu/hyndman>.

1 INTRODUCTION

People seem to tend to procrastinate in many matters, including some with significant economic consequences. This behavior could be due to a form of time inconsistency of preferences (*e.g.*, a preference for the present as represented by hyperbolic or quasi-hyperbolic discounting; see Phelps and Pollak (1968), Laibson (1994, 1997) and O’Donahue and Rabin (1999a,b)). However, this behavior could also represent, an optimal decision to “wait for better times” in an environment in which acting requires the exertion of costs which evolve stochastically over time. To better distinguish between these two motivations for procrastination, in this paper we study the decision problem faced by a single decision maker who has to decide if and when complete a task before some fixed, final deadline that provides a delayed benefit but requires the immediate exertion of effort. We study the case where effort is costly and evolves according to a continuous time stochastic process, and the agents preferences are either time consistent or are characterized by quasi-hyperbolic discounting.

When agents have standard, time-consistent, preferences our model is formally equivalent to the exercise of an American put option. In particular, we show that the decision maker adopts a threshold rule. That is, at each time t , there is a threshold cost, such that if the realized cost of effort at time t is below the threshold, the agent completes the task; otherwise, she waits. The threshold cost increases as the agent gets closer to the final deadline. That is, for the same realized cost, the agent is less likely to delay the closer she is to the final deadline.

In order to study the behavior of decision makers with a present bias, we adapt the model of Harris and Laibson (2004) — an elegant generalization of the standard $\beta - \delta$ quasi-hyperbolic model of Phelps and Pollak (1968) and Laibson (1994) to continuous time. Decision makers can either be naïve or sophisticated (see, *e.g.*, O’Donahue and Rabin (1999b)). Naïve decision makers are unaware of their self-control problems, while sophisticates are aware that they have a tendency to procrastinate. We first show that naïve decision makers will never complete a task strictly before the final deadline. On the other hand, we show that a sophisticated decision maker 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 is increasing as she gets closer to the final deadline.

Therefore, the behavior of sophisticated quasi-hyperbolic and exponential agents are hardly distinguishable when their only choice regards when to exert effort and complete the task — they both have increasing thresholds, it is just that sophisticates’ thresholds are lower (*i.e.*, they delay more) than exponential agents. However, we expect sophisticated quasi-hyperbolic agents, recognizing their innate tendencies to procrastinate, to seek out commitment devices, such as the imposition of binding deadlines, to overcome the urge to delay exerting effort and completing the task. Indeed, we show that sophisticated decision

makers will sometimes, optimally, *self-impose* a deadline strictly more binding than the exogenous final deadline that we impose.

Our paper is related to Sáez-Martí and Sjögren (2008) who study how deadlines affect the timing of effort when agents get distracted (*i.e.*, face a higher cost of effort at certain times). While we focus on continuous time stopping problems, their model is in discrete time, and agents must exert effort in more than one time period. Therefore, it may be that agents start a task but fail to complete it. Another difference is that while they focus on incentives (*i.e.*, deadlines) imposed by principals, we examine whether there is a beneficial role for *self-imposed* deadlines. One surprising result from their analysis is that agents who are often distracted may outperform agents who are less frequently distracted. This follows because agents who are often distracted exert “precautionary” effort earlier (see ? for some experimental evidence on this point).

Our paper is also closely related to Hsiaw (2008) who studies the role of self-imposed, though non-binding, goals as a source of motivation for agents with a present bias. Like us, ? adopts the continuous time quasi-hyperbolic model of Harris and Laibson (2004), though the decision makers in her model face an infinite horizon stopping problem, while our agents face some finite deadline. In her model, goals serve as a reference point; in particular, agents get utility from consumption and from comparing outcomes to the self-set goal. Hsiaw (2008) shows that self-set goals can help overcome one’s tendency to procrastinate; however, if self-control problems are minimal (or absent), they can actually be detrimental.

The rest of the paper is organised as follows. The next section describes the decision problem faced by a single decision maker: at what point in time, at or before some final deadline, should the decision maker optimally complete a task that provides a delayed benefit but requires the immediate exertion of effort (where effort evolves according to a stochastic process). We then characterize optimal behaviour under various forms of time discounting. As mentioned above, we show that exponential discounters employ a threshold strategy, completing the task at time t if the cost is less than the time t threshold. Borrowing on results from the literature on American options, we show that the threshold is increasing in time, and also characterise other comparative statics. In particular, since the value of the option (to complete the task later) is increasing the further away from the final deadline, we are able to say that exponential discounters would never self-impose a deadline.

After deriving our results for time-consistent agents, we turn our attention to agents with a present bias. In particular, we show that naïve hyperbolic agents will only complete a task at the final deadline and only if the cost of doing so is sufficiently low. On the other hand, sophisticated hyperbolic agents, like their exponential counterparts, adopt a threshold strategy — it is just that they adopt lower thresholds making delay more likely. We then show that a naïve hyperbolic agent will never self-impose a deadline, while it may be optimal

for a sophisticate to do so. We show that if the initial cost realization of task completion is known, then she may actually prefer to set a non-trivial deadline. Indeed, the deadline will be sufficiently strict so that she forces herself to complete the task immediately. When the decision maker faces some uncertainty about her initial cost of completing the task, we are unable to analytically solve for the optimal deadline. In Section 3, we conduct a numerical exercise in order to demonstrate that sometimes the decision maker will, in fact, self-impose a non-trivial deadline when there is uncertainty about her initial cost of task completion. Our numerical results would also seem to indicate that if the decision maker self-imposes a deadline, it is an immediate deadline.

In Section 4 we study commitment via methods other than self-imposed deadlines affects behaviour. Motivated by some experimental results by Trope and Fishbach (2000), we two such alternative commitment devices are: (i) Making a fixed payment conditional upon task completion and (ii) imposing a cost for not completing the task. In the former case, we show that, provided the final deadline is not too far away, a sophisticated hyperbolic discounter may prefer to make the fixed fee conditional upon task completion (doing so increases the likelihood that the task is ultimately completed). However, if the deadline is far away, then the decision maker prefers to take the payment up front since she knows that she is likely to delay excessively. In the latter case, when the time to the final deadline is sufficiently close, imposing a cost for not completing the task has a similar effect as does making a fixed payment conditional on completing the task. However, as the time to the final deadline increases, the expected burden of not completing the task diminishes and it has substantially less commitment power.

While deadlines in the set-up of Section 2 are extremely strict — if the task is not completed at or before the deadline, the entire value disappears — in reality, and also in the experiments of Ariely and Wertenbroch (2002) and the model of O’Donahue and Rabin (1999b), missed deadlines often lead to less severe penalties. Also in Section 4, we examine the case in which a penalty per time period that the task is not completed is imposed. Such an incentive scheme has its benefits and drawbacks. The benefit is that the weaker deadline does not completely destroy the option value of being able to wait for a lower cost of task completion. The main drawback is that the incentives for early task completion are also weaker. Like our previous results, we show that if β is sufficiently high, the sophisticated decision maker chooses a deadline (*i.e.*, a point in time in which the penalty kicks in) sufficiently strict so as to ensure immediate completion.

We end Section 4 with a final alternative commitment device: allowing the decision maker the opportunity to self-impose a deadline at any time, rather than only once and only before the decision maker can begin exerting effort. In this case, we show that, for β not too small, this restores full commitment, which means that the sophisticated decision maker

will behave *as if* she had exponential time preferences.

In Section 5 we return to the issue of whether a decision maker would ever self-impose a deadline, but instead we examine other models of decision making. Our first result here is to show that a decision maker with self-control problems à la Gul and Pesendorfer (2001, 2004) would never self-impose a non-trivial deadline. In this section we also discuss Brunnermeier, Papakonstantinou, and Parker (2007) which proposes a model of optimal expectations in which decision makers will both procrastinate and self-impose a deadline. In their model, decision makers consistently underestimate the amount of work required to complete a task, leading to less than optimal initial effort. However, this leads to an anticipatory utility effect: current felicity is boosted because the decision maker anticipates less work in the future. Finally, Section 6 provides some concluding remarks and discusses future directions.

2 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. This will not make a much difference for decision makers who exhibit exponential discounting; however, for hyperbolic discounters it creates a small time wedge between when costs are incurred and the rewards are received. Hence, such decision makers will exhibit a tendency to procrastinate.

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

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

where the parameter $\sigma > 0$ captures variance in the obvious way.¹ 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.

2.1 THE EXPONENTIAL CASE

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

¹An alternative formulation of the evolution of cost that is also intuitively appealing would be a mean reverting process: $dx = \eta(\bar{x} - x)dt + \sigma x \cdot dz$.

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, it is well known that the solution to this problem leads to a threshold function $\bar{x}(t)$, such that for $x(t) \leq \bar{x}(t)$ the decision maker will complete the task, while when the opposite inequality holds, the agent prefers to wait. Moreover, it is well known that in the continuation region $W(x, t)$ is the solution to the following free boundary problem:²

$$\rho W(x, t) = \frac{1}{2} \sigma^2 x^2 W_{xx}(x, t) + W_t(x, t) \quad (2)$$

$$W(\bar{x}(t), t) = \bar{V} - \bar{x}(t) \quad (3)$$

$$W_x(\bar{x}(t), t) = -1 \quad (4)$$

$$W(x, 0) = \max\{\bar{V} - x, 0\} \quad (5)$$

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.

2.1.1 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. The most important from our perspective being the properties of the equilibrium threshold function, $\bar{x}(t)$ and also the deadline. 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.

$\bar{x}'(t) \geq 0$. That is, the threshold cost realisation is increasing in time. That this is so is quite intuitive. As one approaches the deadline, there is less time remaining during which one can wait for lower cost realisations to materialise. More formally:

Proposition 1. *The optimal threshold, $\bar{x}(t)$ is increasing in t .*

Proof. 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

²Peskir and Shiryaev (2006) provides a thorough, yet advanced, treatment of the standard American put option analysis.

$W(x, t') - \max\{\bar{V} - x, 0\} \geq W(x, t) - \max\{\bar{V} - x, 0\} > 0$, which implies that $x > \bar{x}(t')$, which in turn implies that $\bar{x}(t)$ is increasing in t . \square

THE EFFECT OF ρ IS AMBIGUOUS. Observe that an increase in ρ has a negative effect on the option value of waiting, which all else equal, would increase the threshold cost realisation (see, in particular, Peskir and Shiryaev (2006)). However, since the reward, $\bar{V} = e^{-\rho\Delta}V$, it is also true that an increase in ρ lowers the reward from completing the task, and it is known that $\bar{x}(t)$ is increasing in the reward. Therefore, these two effects go in opposite directions and the total effect of an increase in ρ cannot be conclusively determined.

$\frac{\partial \bar{x}}{\partial \sigma} < \mathbf{0}$. The higher is volatility, the more likely it is that there will be wide swings in the cost of task completion. Moreover, as is standard, the decision maker is shielded from cost increases (he can always *not* exercise the option) and benefits from cost decreases. We formally state, but do not prove, the following:

Proposition 2. *The optimal threshold, $\bar{x}(t)$ is decreasing in σ .*

2.1.2 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 $W^\tau(x, t; \omega_\tau)$ denote the value obtained by the decision maker at time t with initial cost realisation x , facing a deadline of τ and following the optimal policy. With this notation, we are now ready to demonstrate:

Proposition 3. *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. \square

2.2 THE HYPERBOLIC CASE

We now wish to move away from time consistent decision makers and move to a world of time inconsistent behaviour. Here, one may take at least two different approaches, depending on the level of sophistication one assumes of the decision makers. In the language of Rabin, decision makers may either be naïve or sophisticated. Naïfs suffer from time inconsistency,

but are themselves unaware of it. That is, in period t , she discounts future flows according to $\beta\gamma, \beta\gamma^2, \dots$; however, importantly, she assumes that *next* period she will behave as an exponential discounter would. This type of behaviour has been shown to create substantial amounts of procrastination. On the other hand, sophisticated decision makers recognise that they are time inconsistent. Therefore, the decision problem is often modelled 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 discounts according to $\beta\gamma$. Generally speaking, it is substantially more difficult to characterise optimal behaviour for sophisticates than for naïfs or even exponential discounters.

Our discussion of hyperbolic discounting will be divided into two parts: an analysis of naïfs and sophisticates. In both cases, we will adopt the framework of Harris and Laibson (2004) 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 $t + \mathfrak{J}$, and the future self, lasting from $t + \mathfrak{J}$ to ∞ . The length of time that the present self exercises control is then \mathfrak{J} , which is stochastic and, as in Harris and Laibson (2004), is exponentially distributed with parameter $\lambda \geq 0$. The discount function of a decision maker born at time 0 can then be summarised as follows:

$$D(t) = \begin{cases} \gamma^t & \text{if } t \in [0, \mathfrak{J}) \\ \beta\gamma^t & \text{if } t \in [\mathfrak{J}, \infty) \end{cases}$$

where $\gamma \in (0, 1)$ and $\beta \in (0, 1]$. For ease of notation, define $\rho = -\ln \gamma$.

2.2.1 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:

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

where $U(\lambda) = e^{-\lambda\Delta} \bar{V} + (1 - e^{-\lambda\Delta}) \beta \bar{V}$ is the expected discounted payment from completing the task immediately.³ If the decision maker transforms into her future self, which occurs

³Since 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 $\lambda \rightarrow \infty$, $U(\lambda) \rightarrow \beta \bar{V}$.

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:

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

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

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

along with the following optimality conditions:

$$W^n(\bar{x}^n(t), t) = U(\lambda) - \bar{x}^n(t) \quad (8)$$

$$W_x^n(\bar{x}^n(t), t) = -1 \quad (9)$$

$$W^n(x, T) = \max\{U(\lambda) - x, 0\} \quad (10)$$

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, $U(\lambda) \leq \bar{V}$. 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 $\lambda \rightarrow \infty$; 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 will complete the task at time t provided that $x(t) \leq \beta[\bar{V} - W^e(x, t)]$. Indeed,

we claim that the decision maker will never complete the task for $x(t) > 0$ and $t < T$. Take any $0 < x(t) \leq \bar{x}^e(t)$, so that the exponential discounter would rationally choose to complete the task. Therefore, 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, as $\lambda \rightarrow \infty$, it becomes ever less likely that the naïve hyperbolic discounter will complete the task at any time $t < T$ and will complete the task at time T if and only if $x(T) \leq \beta\bar{V}$. Formally, we have:

Proposition 4. *Unless $x(t) = 0$, the naïve decision maker will complete the task only at time T and only if $x(T) \leq \beta\bar{V}$.*

2.2.2 SOPHISTICATED DECISION MAKERS

We now turn our attention to sophisticated decision makers. The approach we take is similar to that of the previous subsection, though there are a few conceptual differences brought on by sophistication. In particular, the current self anticipates the actions that will be taken by her future self and incorporates these decisions into the current 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 with the previous cases, 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, it is not difficult to see that the continuation value function can be expressed the following partial differential equation:

$$\rho w^c(x, t) = \frac{1}{2}\sigma^2 x^2 w_{xx}^c(x, t) + w_t^c(x, t) \quad (11)$$

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:

$$W^s(x, t) = \max \left\{ U(\lambda) - x, \ e^{-\rho dt} \mathbb{E} \left[e^{-\lambda dt} W^s(x + dx, t + dt) + (1 - e^{-\lambda dt}) \beta w^c(x + dx, t + dt) \right] \right\} \quad (12)$$

Running through the same steps as for the naïve hyperbolic discounter, we are able to re-write (12) in the waiting region as:

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

Of course, optimality requires the following extra conditions:

$$W^s(\bar{x}^s(t), t) = U(\lambda) - \bar{x}^s(t) \quad (14)$$

$$W_x^s(\bar{x}^s(t), t) = -1 \quad (15)$$

$$W^s(x, T) = \max\{U(\lambda) - x, 0\} \quad (16)$$

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 $\lambda \rightarrow \infty$, 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: $\rho W^s(x, t) = \frac{1}{2}\sigma^2 x^2 W_{xx}^s(x, t) + W_t^s(x, t)$, 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 5. *In the limit as $\lambda \rightarrow \infty$, the sophisticated hyperbolic discounter behaves as if he were an exponential discounter in which the reward for completing the task is $\beta\bar{V}$.*

2.2.3 THE OPTIMAL DEADLINE

It is obvious that the naïve decision maker will set the latest possible deadline since she believes as though she will behave like an exponential decision maker in the future, and it was shown in Proposition 3 that exponential decision makers prefer later deadlines. However, this is not necessarily the case for sophisticated decision makers. They would like their future selves to behave like exponential decision makers but know that they will not. Therefore, since the task is simply a stopping time problem, 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 6. *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 $\beta(V - x(0))$. 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 5, 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 $\tau \rightarrow \infty$, $W(x, 0; \tau) \rightarrow Ax^\alpha$ 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 $x \leq \hat{x}$. \square

The intuition for this would seem to be as follows: since there is a “discontinuous” benefit to completing the task immediately (which arises due to sophistication), the decision maker would like to commit her time 0 self to immediately complete the task. On the other hand, there is still the option value of waiting to complete the task, and this value is increasing in the time to the deadline. Therefore, if it is known that the initial cost will be quite low the discontinuous benefit of commitment outweighs the benefit of holding an option to complete the task at later date, while if it is known that the initial cost will be quite high, then a commitment to immediately complete the task becomes more costly relative to the benefit of holding the option to complete the task at a future date when, possibly, the cost of doing so is lower.

Of course, if there is a significant delay between when the decision maker chooses a deadline and when she may actually begin working on the task, the assumption that she knows $x(0)$ with certainty becomes strained. Therefore, suppose that $x(0)$ is a random variable, the realisation of which is drawn from the 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. Therefore, by integrating over the support of F , an intermediate deadline may become optimal.

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

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

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:

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

There are two opposing effects at play in the above equation. If τ is increased by a small amount, then the threshold for immediate task completion is lowered. This leads to a utility loss of $(1 - \beta)\bar{x}^s(\tau)$. On the other hand, by increasing the deadline, the decision maker is giving herself more opportunities to actually complete the task at some point in the future. Thus she experiences a gain, which is represented by the integral. Therefore, without further assumptions, it is not possible to analytically characterise the stationary points of (18). In order to gain further intuition, we solve for the optimal deadlines numerically in the next section.

3 A NUMERICAL EXERCISE TO FIND OPTIMAL SELF-IMPOSED DEADLINES

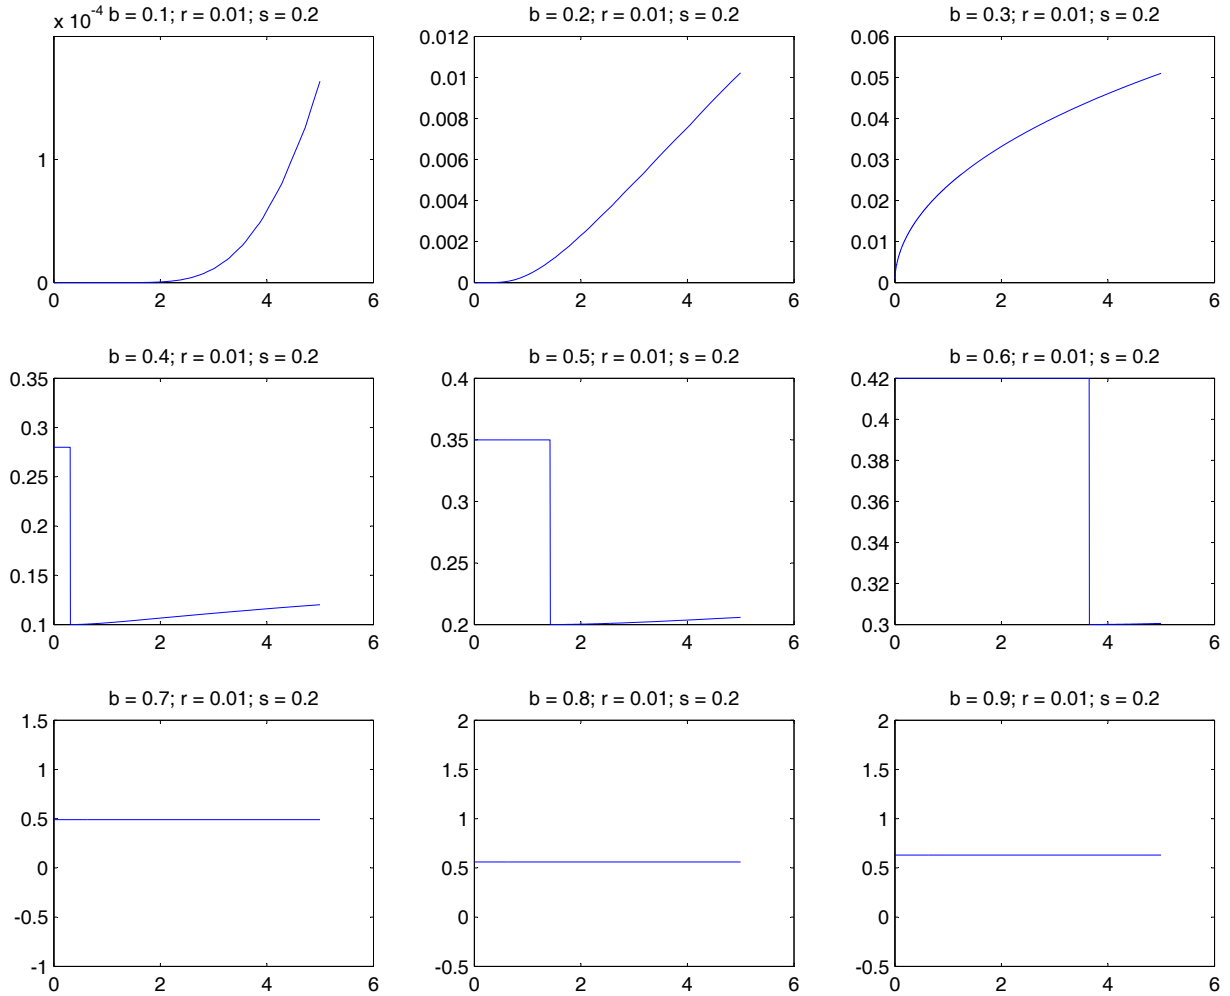
Given the inability above to formally show whether a sophisticated hyperbolic discounter will set a deadline when there is uncertainty over the initial cost realisation, we seek to answer this question numerically. Our task is greatly facilitated by Proposition 5, since it tells us that we need to solve 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:

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

Then, 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 $\beta(\bar{V} - x)$, while if she delays at time 0, then her value is the price of the corresponding option. The results of this exercise are shown in Figures 1 – 4. In each figure r and σ are fixed and β varies from 0.1 to 0.9. On the horizontal axis is the deadline, τ , where $\tau = 0$ represents an immediate deadline and $\tau = 5$ is a deadline of 5 years into the future.

In Figure 1, we conducted the simulation exercise for the case in which the initial cost, $x(0)$, is known with certainty and fixed at $x(0) = 0.3$. This set of simulations nicely illustrates

FIGURE 1: Simulation Results; $r = 0.01$ & $\sigma = 0.2$; $x(0) = 0.3$, **known**



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.


Proposition 6 and also provides further intuition. When β is very low, initially the decision maker prefers an infinite deadline. This follows because an immediate deadline would induce a threshold of $\beta\bar{V} < x(0)$, so the *ex ante* self cannot force his future self to complete the task — hence he prefers no deadline at all. However, as β increases, so that the decision maker’s self control problems become less severe, we see that the he prefers to impose a deadline such that the he immediately completes the task. As can be seen, in general, there exists a threshold deadline τ^* , depending upon the parameters, such that if $\tau \leq \tau^*$, the task will be completed immediately, leading to a utility of $\beta(\bar{V} - x(0))$ for the *ex ante* self. In contrast, for $\tau > \tau^*$, $\bar{x}^s(0, \tau) > 0.3$, the task is not immediately completed and the *ex ante* self experiences a discontinuous drop in his expected value. This follows because $\beta(\bar{V} - x(0)) > \beta\bar{V} - x(0) \approx W^s(x(0), 0, \tau^* + \epsilon)$. Of course, for $\tau > \tau^*$, the *ex ante* expected utility (*i.e.*, $W^s(x(0), 0, \tau)$) is increasing in τ . Finally notice that, since $\bar{x}^s(0, \tau)$ is increasing in β , we also have that τ^* is increasing in β .

Next turn to Figures 2 – 4. These figures provide numerical results for the more realistic case in which $x(0)$ is not known with certainty; in particular $x(0)$ was drawn from the uniform distribution with support $[0, 1]$. As we have discussed above, when $x(0)$ is not known, there is some scope for intermediate deadlines which do not necessarily guarantee immediate task completion. The goal of this exercise is to determine how large that scope is. Notice that the figures appear discontinuous; this is due to the finite grid that was used for the numerical exercise. A finer grid will produce more accurate graphs but at the cost of increased computation time.

Notice that in all cases an immediate deadline provides for a local maximum, with the *ex ante* expected value decreasing for very short deadlines. That this is so can be seen by examining once again (18). When $\tau = 0$, Carr, Jarrow, and Myneni (1992) allows us to conclude that $\frac{\partial \bar{x}^s}{\partial \tau} = -\infty$. Therefore, by increasing the deadline even a little bit, the *ex ante* self is substantially lowering the chance that her future self will immediately complete the task. Therefore, she loses the discontinuous benefit of immediate task completion. However, as τ increases further, $\bar{x}^s(\tau)$ becomes flatter and so the loss in the discontinuous benefit of immediate completion gets mitigated by the increase in option value — the second term in (18). Eventually, especially when β is relatively large this latter effect begins to dominate and $\bar{W}^s(\tau)$ becomes increasing in τ .

It would seem, therefore, that the result on the optimal deadline is the same whether or not the initial cost realisation $x(0)$ is known with certainty. That is, either the decision maker prefers an immediate deadline or no deadline at all. Moreover, the higher is β the more likely is it that the decision maker prefers no deadline at all — a result that makes sense since for β large the decision maker’s self control problems are not too severe. With respect to the particular examples we have solved, one can see in Figure 3 that for $\beta = 0.9$,

FIGURE 2: Numerical Results; $r = 0.05$ & $\sigma = 0.1$




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.

the decision maker prefers no deadline to an immediate deadline.

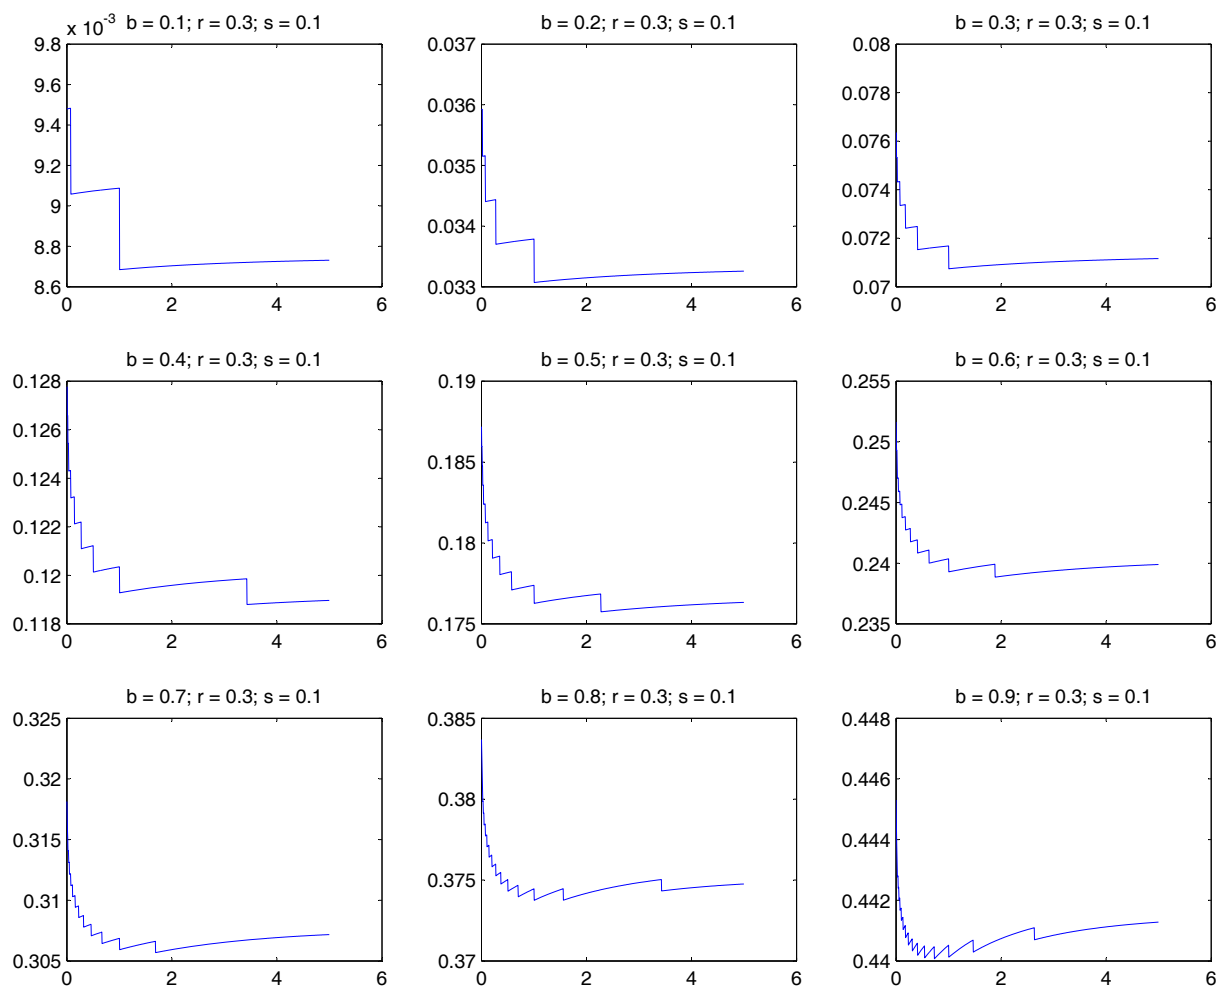
Another interesting comparative static concerns σ . In comparing Figures 2 and 3, it appears that as σ rises the desire to impose deadlines declines. The intuition for this is as follows: If σ increases then, by Proposition 2, the option value of waiting also increases since it is now more likely that a lower cost realization will occur in the future. However, imposing a deadline is a blunt instrument which kills the option value of waiting; therefore, deadlines must be strictly less attractive.

FIGURE 3: Numerical Results; $r = 0.05$ & $\sigma = 0.2$



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.

FIGURE 4: Numerical Results; $r = 0.3$ & $\sigma = 0.1$



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.

4 COMMITMENT VIA OTHER METHODS

As we have seen, self-imposed deadlines are sometimes optimal, in particular for $\beta \ll 1$? 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 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, they allow 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 so 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 $x(0) \in (\bar{x}^s(0, \bar{V}), \bar{x}^s(0, \bar{V} + v))$. 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,

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

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


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 5 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), rather than allowing subjects to make a fixed payment conditional on task completion, the authors instead 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 6. At the deadline, there is no difference between a 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 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

FIGURE 5: 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.

FIGURE 6: 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.

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.

4.3 IMPOSING A PENALTY PER TIME PERIOD THAT THE TASK IS NOT COMPLETED.

In the experiments of Ariely and Wertenbroch (2002) and in much of the theoretical literature (*e.g.*, O’Donahue and Rabin (1999b)), rather than imposing a strict deadline beyond which the task cannot be completed, a weaker deadline is imposed. In particular, if the task has not been completed by a certain time period, say t^* , a per unit of time penalty is imposed. Therefore, if the task is completed at time $t > t^*$, the reward is only $V - (t - t^*)\mu$, where $\mu > 0$ is the per period penalty. Comparing such a scheme with the strict deadlines that have been our focus thus far, a penalty for late completion provides only weaker incentives for early task completion, but it also does not destroy the option value of waiting. Therefore, in some circumstances it may be a useful commitment device for sophisticated decision makers. The appropriate question is then, what time t^* , after which a penalty for late completion kicks in, would a decision maker set?

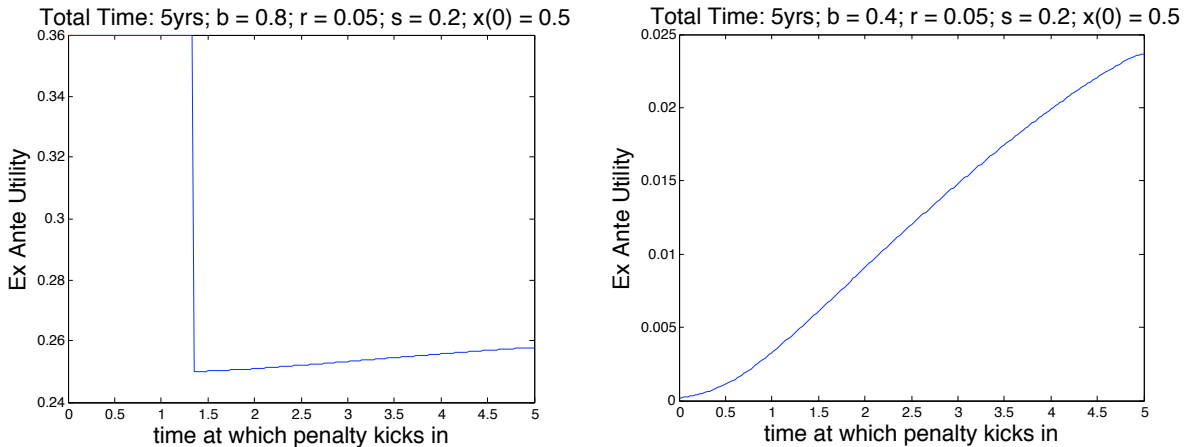
Figure 7 provides this answer for two cases. In each case the value for completing the task is $\bar{V} = 1$ and the penalty is 0.01 per unit of time.⁵ For simplicity, we also assumed that the initial cost, $x(0)$, of task completion is known with certainty. The result is similar to the case of a strict deadline: if β is sufficiently high, then the sophisticated decision maker prefers to choose a deadline sufficiently early so as to ensure immediate task completion. On the other hand, for β small, by setting an immediate deadline, the decision maker cannot compel her future self to complete the task immediately. In this circumstance, she prefers the latest possible deadline.

4.4 ALLOWING THE DECISION MAKER TO SELF-IMPOSE A DEADLINE AT *Any* TIME.

A fourth approach is as follows. Suppose that the agent can impose a deadline on her future selves at any time t . That is, at any time t , once she realises $x(t)$ she can impose a deadline strictly more binding than the current deadline. While this may seem like it complicates the

⁵The unit of time was chosen such that if the penalty kicks in at $t = 0$, then at $t = 5$, which is a “hard” deadline, there is no value remaining from completing the task.

FIGURE 7: Numerical Results: Imposing a Penalty Per Time Period That the Task is not Completed



model a great deal, we argue that it simply reduces the model, for β not too small, to the case in which the agent has *full commitment*.

Let $\bar{x}^s(t)$ be the cutoff of a sophisticated hyperbolic agent at time t and let $\bar{x}^e(t)$ be the cutoff of an exponential agent at time t . We know that $\bar{x}^s(t) < \bar{x}^e(t)$. Suppose now the cost at time t , $x(t) \in (\bar{x}^s(t), \bar{x}^e(t))$. In this case, a sophisticated hyperbolic agent does not want to complete the task but, importantly, she would prefer that her future self, at time $t + \epsilon$, does complete the task.⁶ Therefore, she will set a deadline such that, with high probability, $x(t + \epsilon) < \bar{x}^s(t + \epsilon) < \bar{x}^e(t + \epsilon)$. Indeed, since x follows a geometric Brownian motion as the length of the time interval, ϵ , shrinks so does the variance (which is proportional to σdt); therefore, as $\epsilon \rightarrow 0$, an immediate deadline means that the new threshold is $\bar{x}^s = \beta \bar{V}$. Provided that β is not too small, this will imply that $\bar{x}^s = \beta \bar{V} < \bar{x}^e(t)$.

This shows that a sophisticated hyperbolic agent who can impose deadlines at any moment on her future selves will in fact do so, for every path of x which enters in the region of completion, provided that her self-control problems are not too severe. Therefore, with respect to task completion, such agents will behave (in a formal sense) identically to exponential agents.

5 SELF-IMPOSED DEADLINES IN OTHER MODELS

Our analysis above suggests that self-imposed deadlines in the classic quasi-hyperbolic discounting framework are a relatively rare event. We now briefly discuss a few other models and

⁶This follows because the continuation value, w^c , is discounted exponentially.

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 various 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, and therefore procrastinate. One can easily go beyond his analysis and ask whether the agent would like to bind herself by setting a deadline. In Appendix A, we consider the finite time version of Miao’s model and prove in Proposition 7 that the value function is increasing in the time to complete the task. Therefore, an agent with Gul-Pesendorfer preferences will not self-impose deadlines.

5.2 OPTIMAL EXPECTATIONS.

Brunnermeier, Papakonstantinou, and Parker (2007) propose a model of optimal expectations in which decision makers will both procrastinate *and* self-impose binding deadlines. In their model, decision makers consistently under-estimate the amount of work required to complete a particular task, which leads to lower than optimal initial effort. However, the fact that the decision maker underestimates the required effort leads to an anticipatory utility effect: current felicity is boosted because he anticipates less work in the future. Thus these over-optimistic beliefs cause low initial effort, give an anticipatory boost to utility, but lead to extra future effort. The authors show that the *ex ante* utility benefits to over-optimism outweighs the *ex post* cost of poor planning; therefore, procrastination is, in some sense, optimal. Brunnermeier, Papakonstantinou, and Parker (2007) then show that agents will self-impose deadlines, which are less stringent than would be imposed by an outsider. It is this feature of their model which they claim is supported by the experimental evidence of Ariely and Wertenbroch (2002).

5.3 MISPERCEPTIONS.

In our model of decision-making, because they believe (correctly in the case of exponential and sophisticated hyperbolic decision makers), that the option value of having extra time to complete the task outweighs the “commitment” gains from imposing a deadline, decision makers do not choose to self-impose deadlines very often. In order for decision makers to more often wish to self-impose deadlines, it must be that the option value is actually much weaker than in our model. Along the lines of Brunnermeier, Papakonstantinou, and Parker (2007), suppose that decision makers misperceive the costs and/or benefits of completing

the task. Recall (1):

$$dx = \sigma x \cdot dz$$

where we assumed that the cost of completing the task follows a geometric Brownian motion, without drift. Two seemingly plausible misperceptions could be the following. First, it may be that the actual stochastic process for cost contains a drift term so that (1) becomes:

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

where $\alpha > 0$ implies that the cost of completing the task is increasing over time. Under this altered model, we assume that $\beta = 1$ 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. Now, in this model, when a decision maker is faced with a high cost of task completion, she thinks that it was due to a high realization of dz , rather than due to positive drift. Therefore, she may not complete the task now because she (incorrectly) expects the cost to be lower in the future.

Now consider the decision maker at time 0 who is aware of her problems with misperception. Because $\alpha > 0$, as time passes, it becomes increasingly less likely that the decision maker will complete the task. Therefore, the option value of having extra time is much less valuable, while the commitment benefit of a deadline remains, making it much more likely that the time 0 self will set a deadline.

As an alternative to the drift term, one could assume instead that the stochastic process for costs actually follows a jump-diffusion. That is, there is some Poisson arrival rate at which the cost of completing the task “jumps” by a discrete, positive amount. In this model, with a similar reinterpretation of the current and future selves, the presence of misperceived jumps also mitigates the option value of having extra time to complete the task, making self-imposed deadlines much more likely.

6 CONCLUSIONS

In this paper we have 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 (strong) present bias for immediate gratification along the lines of Harris and Laibson (2004). 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. Our results have shown that neither exponential discounters nor naïfs would ever willingly self-impose a non-trivial deadline. The distinction, between these two decision makers lies, therefore in when these decision

makers will complete a task. While an exponential discounter has a threshold at every time t , and will complete the task if the cost of doing so is less than the threshold, a naïf will never complete the task strictly before the deadline. Therefore, such decision makers suffer from a very extreme form of procrastination.

In contrast, our model shows that sophisticates will sometimes self-impose a deadline. The reason for this is because, *ex ante* there is a discontinuous benefit to immediately completing the task. On the other hand, if the task is not completed immediately, then our results show that the sophisticated decision maker behaves like an exponential decision maker, but with a lower value of task completion. Consequently, if the task is not immediately completed, the sophisticated decision maker (just like an exponential discounter) prefers to have as much time as possible to complete the task. Thus there is a kind of bang-bang property of self-imposed deadlines. Either the sophisticated decision maker prefers an immediate deadline or no deadline at all. Somewhat surprisingly, the decision maker will never set an intermediate deadline.

While our model has focused on a single task, it can also be extended to multiple tasks. However, care must be taken on this front for it seems that without further alterations, this extended model would predict identical threshold 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 (perhaps deterministic) 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. 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 (*i.e.*, tasks that require the exertion of effort for some, possibly random, amount of time). This adds an 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, such an extended model could easily be parameterized to reconcile the cycles found in Study 1 of Burger, Charness, and Lynham (2009). In particular, instead of the geometric Browning motion used to model the evolution of opportunity costs, one could work with a mean reverting process with cycles.

REFERENCES

ARIELY, D., AND K. WERTENBROCH (2002): “Procrastination, Deadlines, and Performance: Self-Control by Precommitment,” *Psychological Science*, 13(3), 219–224.

- BRUNNERMEIER, M., F. PAPAKONSTANTINOY, AND J. PARKER (2007): “An Economic Model of the Planning Fallacy,” Working Paper.
- BURGER, N., G. CHARNES, AND J. LYNHAM (2009): “Three Field Experiments on Procrastination and Willpower,” Working Paper.
- 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.
- 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 (2004): “Instantaneous Gratification,” Working Paper.
- HSIAW, A. (2008): “Goal-Setting, Social Comparison, and Self-Control,” Working Paper.
- 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’DONAHUE, 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.

SÁEZ-MARTÍ, M., AND A. SJÖGREN (2008): “Deadlines and Distractions,” *Journal of Economic Theory*, 143, 153–176.

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 choose to 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 known and deterministic 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 7. $W_{t+1}(x) \geq W_t(x)$ 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 $t = 0, 1, \dots, n$. We now show that $W_{n+1}(x) \geq W_n(x)$. Observe that:

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

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

Of course, while Proposition 7 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 so far has always concerned itself with the limiting case of $\lambda \rightarrow \infty$. 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, this issue of the agent's willingness to self-impose a binding deadline arises again. Recall that for $\lambda = \infty$, so that the sophisticated agent knows that he will exercise control only for an infinitesimal length of time, self-imposed deadlines are a relatively rare occurrence.

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

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

The reason for this continuity in λ is as follows. Despite being sophisticated, and therefore knowing that his future self will not complete the task for high enough realisations of the cost, the decision maker is 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 $U(\lambda)$ is received after a length of time dt , rather than immediately). However, if the future self completes the task, there is a savings of $(1 - \beta)x$, because now the cost will also be discounted by β . 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, written as a function on λ , $\bar{x}(\lambda)$, is decreasing.

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.