Procrastination, Self-Imposed Deadlines and Other Commitment Devices*

Kyle Hyndman
University of Texas at Dallas†

Alberto Bisin
New York University and NBER‡

July 23, 2021

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) exponential decision makers, (2) naïve hyperbolic discounters and (3) sophisticated hyperbolic discounters. Absent deadlines, we show that sophisticated hyperbolic decision makers behave *as if* they were time consistent but instead have a smaller reward for completing the task, while naïfs never complete the task. In the presence of deadlines, sophisticated decision makers may, counter-intuitively, have a threshold which is decreasing as they approach the deadline. An extensive numerical study shows that, unlike exponential or naïfs who always prefer longer deadlines, sophisticated decision makers will often self-impose a binding deadline as a form of commitment, while naïve decision makers will not, and we show how this varies with changes in underlying cost, preference and self-control parameters.

*We are grateful to the editor, Nicholas Yannelis, and an anonymous referee for extremely helpful comments, which have led to a much improved paper.

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

‡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`.

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 naïve 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 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.[2]

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.[3] 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. They term this limiting case, "instantaneous gratification".

In Section 3, we consider the case in which the decision makers do not face any deadline – that is, we consider the infinite horizon model. In this case, we are able to provide a clean characterization of behavior. First, exponential decision makers have a constant threshold, above which they wait and below which they complete the task. Second, naïve hyperbolic decision makers, in the instantaneous gratification limit, *never* complete the task unless their cost is 0. On the other hand, sophisticated hyperbolic decision makers, in the instantaneous gratification limit, behave *as if* they are exponential, but with a lower value of completing the task.

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

[2] Other types of commitment include e.g., 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)).

[3] To be sure, there is a long history of theoretical work on present bias including, among others and beyond those already mentioned, Strotz (1956), Laibson (1997), O'Donoghue and Rabin (1999a,b) and Vieille and Weibull (2009). More recently, Calcott and Petkov (2021) have studied over-consumption due to present-bias, generalized to *any* present-biased time preferences and arbitrary time horizons. They show that the social value of commitment devices remains in the general structure that they consider.

In Section 4, we turn our attention to the finite horizon problem where the decision makers face a fixed deadline, T. Analytical results are more difficult because the problem becomes non-stationary. However, we show: (i) a decision maker with 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; and (ii) in the instantaneous gratification limit, 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. For sophisticated hyperbolic decision makers, behavior is much more complex and we show by simulations that, in the instantaneous gratification limit, the decision maker actually has a monotonically *decreasing* threshold. Moreover, in the instant before the deadline, the decision maker *never* completes the task. This is because, if a sophisticated decision maker, who is guaranteed to lose control in the next instant, knows that she will complete the task in the next instant, then it never makes sense to complete the task now.[4]

Given our examination 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.

For sophisticated hyperbolic decision makers, we rely on numerical results to gain insights into their willingness to set deadlines. We show that sophisticated hyperbolic decision makers may choose to self-impose a binding deadline $D < T$. Moreover, the tension in this decision is, generally, between an immediate deadline or no deadline at all. Because the threshold is relative flat (compared to an exponential decision maker) or decreasing if the probability of maintaining control is small, most task completions occur either right away or at the deadline, with relatively few task completions occurring in the interim. By imposing an immediate deadline, the sophisticated decision maker forces all tasks to be completed at time 0. Because the decision maker is impatient this has some value. However, it comes at the cost of destroying some option value that a decision maker with a relatively high initial cost can complete the task at all. Whichever effect dominates, generally, determines whether

[4]Although not central to the main paper, it is interesting to examine the question of self-imposed deadlines through other lenses. In Appendix A briefly do this for two other models. The first is based on misperceptions – though it is different from recent work by Gabaix and Laibson (2017), while the second is based on temptations à la Gul and Pesendorfer (2001, 2004). In the former case, we show that a sophisticated decision maker may commit to a deadline, while in the latter case, we show, surprisingly, that a decision maker with Gul and Pesendorfer (2001, 2004) preferences will never self-impose a binding deadline.

an immediate deadline or no deadline at all is optimal.

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 $\beta - \delta$ 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 (2020) 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. Indeed, Bisin and Hyndman (2020) showed that, in the presence of multiple, repeated tasks (subjects had to do an identical task three times over a two-week period), subjects appeared able to activate internal self-control, thereby rendering deadlines superfluous. This paper also showed the importance of perseverance and over-confidence about task difficulty in explaining the behavior of subjects.

Demand for commitment has also been shown in other contexts. For example, Houser, Schunk, Winter, and Xiao (2018) 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).[5]

The question of commitment is an important policy question too. For example, in response to low retirement savings (possibly due to present-bias), many governments mandate contributions to retirement accounts. Andersen and Bhattacharya (2021) show that this can be welfare-improving even if sophisticated present-biased decision makers try to borrower against their (forced) retirement savings to increase consumption in the present.

As noted in the introduction, our theoretical framework is based upon Harris and Laibson (2013).[6] 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 naïve 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 behave closer to the time-consistent benchmark than do naïve decision makers. Heidhues and Strack (2019) have recently pointed out that it is challenging to distinguish sophisticated quasi-hyperbolic from exponential decision makers based solely on their observed task completion behavior. While this is true in our setting in the infinite horizon case, when the horizon is finite, there will be a mass of hyperbolic decision makers completing the task at the deadline, and this spike at the deadline will be larger the more likely it is that the present self loses control in an instant (i.e., the higher is λ in the model we outline below).

3 The Model: No Deadlines

A decision maker is faced with a task to complete. If the task is completed the decision maker will receive a reward of $V > 0$, otherwise no reward is received. The decision maker faces no deadline for the completion of the task. 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

[5] For a more detailed summary, the reader is referred to Bryan, Karlan, and Nelson (2010) and Ericson and Laibson (2019).

[6] Cao and Werning (2016) have generalized and extended Harris and Laibson (2013) in several useful ways. Unfortunately, we have not been able to apply their techniques to our problem.

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 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 $W(x) = \sup_{t \leq \tau < \infty} \mathbb{E}_x(e^{-\rho(\tau - t)} \max\{\bar{V} - x, 0\})$ denote the value of the task when the current cost is x, the current time is 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 a perpetual 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, $\bar{x}$, such that for $x(t) \leq \bar{x}$ the decision maker will complete the task at t, while when the opposite inequality holds, the agent prefers to wait. It is well known that in the continuation region $W(x)$ is the solution to the following free boundary problem:

$$\rho W(x) = \frac{1}{2}\sigma^2 x^2 W_{xx}(x) \tag{2}$$

$$W(\bar{x}) = \bar{V} - \bar{x} \tag{3}$$

$$W_x(\bar{x}) = -1 \tag{4}$$

In the language of Dixit and Pindyck (1994), (2) represents the first order condition, (3) the value matching condition, and (4) the smooth pasting condition at the boundary.

We can solve the problem analytically and show that (i) the threshold is the constant $\bar{x} = \frac{\alpha}{\alpha-1}\bar{V}$, where α is the negative root of the fundamental quadratic $\frac{\sigma^2}{2}\alpha(\alpha-1) - \rho = 0$ and $A = -{}^{\bar{x}^{1-\alpha}}/_{\alpha}$;[7] (ii) the value for $x > \bar{x}$ is $W(x) = Ax^{\alpha}$.

3.2 The Hyperbolic Case

We now move away from time consistent decision makers and move to a world of time inconsistent behaviour. We follow O'Donoghue and Rabin (1999b) 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 $t+s$

[7]With a little algebra, we obtain that $\alpha = {}^{1}/_{2} - {}^{\sqrt{8\rho+\sigma^2}}/_{2\sigma}$.

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 sophisticated 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 discount function of a decision maker born at time t can then be summarised as follows:

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

where $\rho > 0$ and $\beta \in (0,1)$ (a decision maker with $\beta = 1$ is exponential).

The length of time that the present self exercises control, τ_t, is stochastic and, as in Harris and Laibson (2013), is exponentially distributed with parameter $\lambda \geq 0$. In the limit for $\lambda \to \infty$ the present-self maintains control for a vanishingly small amount of time - the "instantaneous gratification" case.

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 $W^n(x)$ denote the value function of a naïf and $W^e(x)$ denote the value function of the exponential discounter. For any x, in an interval of length dt, the value function, $W^n(x)$, must solve:

$$W^n(x) = \max\left\{\bar{V}_\lambda - x\ ,\quad e^{-\rho dt}\mathbb{E}\left[e^{-\lambda dt}W^n(x+dx) + (1-e^{-\lambda dt})\beta W^e(x+dx)\right]\right\} \tag{5}$$

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.[8] 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

[8]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 $\lambda \to \infty$, $\bar{V}_\lambda \to \beta\bar{V}$.

of $W^e(x+dx)$.

Now suppose that we are in the region for which it is optimal to wait. Multiply both sides of (5) by $e^{\rho dt}$ and subtract $W^n(x)$ to obtain:

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

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

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

Dividing through by dt and taking the limit as $dt\to 0$, we arrive at the following expression for the first order condition:

$$\rho W^n(x)=\frac{1}{2}\sigma^2x^2W^n_{xx}(x)+\lambda[\beta W^e(x)-W^n(x)] \tag{6}$$

along with the value matching and smooth pasting conditions

$$W^n(\bar{x}^n) = \bar{V}_\lambda-\bar{x}^n \tag{7}$$

$$W^n_x(\bar{x}^n) = -1 \tag{8}$$

It follows that $\bar{x}_n\leq\bar{x}$ for all t and for all $\lambda>0$. This is so for two reasons. First, the reward that the hyperbolic discounter receives, $\bar{V}_\lambda$ is less than the reward that an otherwise identical exponential discounter would receive, $\bar{V}$. 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.

In the instantaneous gratification limit, as $\lambda\to\infty$, condition 6 takes the simple form $W^n(x)=\beta W^e(x)$ in the waiting region. Therefore,

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

and the decision maker will complete the task at time t provided that $x(t)\leq\beta[\bar{V}-W^e(x)]$.

In particular, as $\lambda \to \infty$, the naïve decision maker's value in the waiting region is $W^n(x) = \beta W^e(x) = \beta \max\{Ax^\alpha, \bar{V} - x\}$, where A and α are as defined in Section 4.1. In order to do the task, it must be that $\beta\bar{V} - x \geq \beta \max\{Ax^\alpha, \bar{V} - x\}$. First, observe that when the exponential decision maker waits, then so will the naïve decision maker. This is because $Ax^\alpha > \bar{V} - x > \beta\bar{V} - x$. Second, observe that if the exponential decision maker would complete the task, then for the naïve decision maker to complete the task would require $\beta\bar{V} - x \geq \beta(\bar{V} - x)$, which is only possible when $x = 0$. We have therefore proven the following:

Proposition 1. *In the limit as $\lambda \to \infty$, unless $x(t) = 0$, the naïve decision maker will never complete the task.*

3.2.2 Sophisticated Decision Makers

We turn now to sophisticated decision makers. In this case, 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 $W^s(x)$ ("s" for sophisticated) and a continuation value function $w^c(x)$ ("c" for continuation).

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

$$W^s(x) = \max\left\{\bar{V}_\lambda - x\ ,\quad e^{-\rho dt}\mathbb{E}\left[e^{-\lambda dt}W^s(x+dx) \right.\right.$$
$$\left.\left. +(1-e^{-\lambda dt})\beta w^c(x+dx)\right]\right\} \tag{9}$$

The first term inside the max expression indicates the value of completing the task now, while the second term represents the value of waiting. Running through the same steps as above, we can re-write (9) in the waiting region (i.e., where the second term exceeds the first) as:

$$\rho W^s(x) = \frac{1}{2}\sigma^2 x^2 W^s_{xx}(x) + \lambda[\beta w^c(x) - W^s(x)]. \tag{10}$$

Of course, optimality requires that at the threshold $\bar{x}^s$, where the decision-maker is indifferent between completing or waiting, the following value matchiung and smooth pasting conditions hold:

$$W^s(\bar{x}_s) \quad = \quad \bar{V}_\lambda - \bar{x}_s \tag{11}$$
$$W^s_x(\bar{x}_s) \quad = \quad -1. \tag{12}$$

In the instantaneous gratification case, as $\lambda \to \infty$, we have that $\beta w^c(x,t) = W^s(x,t)$ in the waiting region. This insight let's us characterize the solution. Indeed, we claim that the decision maker behaves *as if* she were exponential but with a value equal to $\beta\bar{V}$

for completing the task. That is, in the waiting region, $W^s(x) = Bx^\alpha$, where, as in the exponential case, α is the negative root of $\frac{\sigma^2}{2}\alpha(\alpha-1) - \rho = 0$, while $\bar{x}_s = \frac{\alpha}{\alpha-1}\beta\bar{V}$ and $B = -\frac{\bar{x}_s^{1-\alpha}}{\alpha}$. Although the calculations are messy, plugging in the values for $W^s(x)$ and $\bar{x}_s$, we can see that (9)–(12) are satisfied. We have therefore proven the following:

Proposition 2. *In the limit as $\lambda \to \infty$, and with no deadline, the sophisticated hyperbolic discounter behaves as if she were an exponential discounter facing a reward for completing the task equal to $\beta\bar{V}$.*

4 The Model: Finite Deadline

We now consider the problem in which the decision maker must complete the task by a finite deadline, $T < \infty$. That is, if the task is completed at time $t \leq T$, then the decision maker will receive a reward V at time $t + \Delta$, otherwise no reward is received. We maintain all previous assumptions.

4.1 The Exponential Case

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 $x(t)$ and the option expires at time 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, $\bar{x}(t)$, 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. Furthermore, $\bar{x}(t)$ is increasing in t. The continuation region $W(x,t)$ is the solution to the following free boundary problem:

$$\rho W(x,t) = \frac{1}{2}\sigma^2 x^2 W_{xx}(x,t) + W_t(x,t) \tag{13}$$
$$W(\bar{x}(t),t) = \bar{V} - \bar{x}(t) \tag{14}$$
$$W_x(\bar{x}(t),t) = -1 \tag{15}$$
$$W(x,T) = \max\{\bar{V} - x, 0\} \tag{16}$$

where, besides the first order, value matching, and smooth pasting conditions, (13-15), the terminal condition (16), describing optimal behaviour at the deadline, need be satisfied.

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 state a few results and also provide some intuition for them.

The main comparative statics are summarized as follows:

Proposition 3. *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 $\bar{x}(t)$

We formally prove (a) and provide an intuitive discussion for (b) and (c). For the interested reader, an intuitive discussion can be found in Hull (2005), while Peskir and Shiryaev (2006) contains an advanced treatment.

(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.

4.2 The Hyperbolic Case

We now turn to the case of time-inconsistent decision makers, facing a deadline at $T < \infty$.

4.2.1 Naïve Decision Makers

Following a similar approach as for the case of naïve decision makers with no deadline, the solution to the decision makers problem is given by the system:

$$\rho W^n(x, t) = \frac{1}{2}\sigma^2 x^2 W^n_{xx}(x, t) + W^n_t(x, t) + \lambda[\beta W^e(x, t) - W^n(x, t)] \tag{17}$$

along with the following optimality conditions:

$$W^n(\bar{x}_n(t), t) = \bar{V}_\lambda - \bar{x}_n(t) \quad (18)$$

$$W^n_x(\bar{x}_n(t), t) = -1 \quad (19)$$

$$W^n(x, T) = \max\{\bar{V}_\lambda - x, 0\} \quad (20)$$

In the instantaneous gratification limit, as $\lambda \to \infty$, we have that $W^n(x,t) = \beta W^e(x,t)$ 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 4. *In the limit as $\lambda \to \infty$, unless $x(t) = 0$, the naïve decision maker will complete the task only at time T and only if $x(T) \leq \beta\bar{V}$.*

4.2.2 Sophisticated Decision Makers

We turn now to sophisticated decision makers. Following a similar approach as for the case of sophisticated decision makers with no deadline, 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). We know that, if the decision maker would complete the task given cost x and time t, then $w^c(x,t) = \bar{V} - x$.

Unfortunately, compared to the other cases, we cannot guarantee the existence of a threshold cost $\bar{x}^s(t)$ below which it is optimal to complete the task, and above which it is optimal to wait. Regardless, consider $W^s(x,t)$. It may be expressed as:

$$W^s(x,t) = \max\left\{\bar{V}_\lambda - x\ , \quad e^{-\rho dt}\mathbb{E}\left[e^{-\lambda dt}W^s(x+dx, t+dt) + (1-e^{-\lambda dt})\beta w^c(x+dx, t+dt)\right]\right\} \quad (21)$$

The first term inside the max expression indicates the value of completing the task now, while the second term represents the value of waiting. Running through the same steps as above, we can re-write (21) in the waiting region (i.e., where the second term exceeds the first) as:

$$\rho W^s(x,t) = \frac{1}{2}\sigma^2 x^2 W^s_{xx}(x,t) + W^s_t(x,t) + \lambda[\beta w^c(x,t) - W^s(x,t)]. \tag{22}$$

Let $x^s_b(t)$ denote a boundary cost, that is a cost at which the decision-maker is indifferent between completing or waiting. Of course, optimality requires that at any such cost,

$$W^s(\bar{x}^s_b(t),t) = \bar{V}_\lambda - \bar{x}^s_b(t) \tag{23}$$

$$W^s_x(\bar{x}^s_b(t),t) = -1. \tag{24}$$

Furthermore, we know that at the deadline, it must be that:

$$W^s(x,T) = \max\{\bar{V}_\lambda - x, 0\}. \tag{25}$$

In the instantaneous gratification limit, as $\lambda \to \infty$, we have that $\beta w^c(x,t) = W^s(x,t)$ in the waiting region. However, unlike in the case of the naïve decision maker, this insight does not let us characterise the solution. It does, however, generate one key insight. First, observe that when $\lambda = \infty$, we know that the decision maker does have a threshold at the deadline – namely, $\bar{x}^s(T) = \beta\bar{V}$. Now consider the instant before the deadline, $t' = T - \epsilon$, and suppose that $x < \bar{x}^s(T)$. If the decision maker completes the task, her value is $\beta\bar{V} - x$, while if she waits, then she will earn approximately $e^{-\rho\epsilon}\beta(\bar{V} - x)$. This is because, in expectation, the cost will not change in the ϵ-interval that she waits between the present and the deadline, and the variance of cost is proportional to ϵ. Therefore, in the limit as $\epsilon \to 0$, we conclude that a sophisticated hyperbolic decision maker (with $\lambda = \infty$) will only complete the task an instant before the deadline if her cost is 0.

Therefore, reasoning by backward-induction from the final deadline, the decision maker may now complete the task when the cost is positive. For example, consider time $t'' = T - 2\epsilon$. If she completes the task at t'', then her earnings are $\beta\bar{V} - x(t'')$. On the other hand, if she waits, then she knows that she will not complete the task until the deadline. Therefore, she can expect to earn approximately $e^{-2\rho\epsilon}\beta(\bar{V} - x(t''))$. One can easily see that there is a strictly positive threshold for which the decision maker would complete the task. Thus, very close to the deadline, a sophisticated decision maker (with $\lambda = \infty$) may have a non-increasing threshold policy for completing the task.

A discrete version of this problem may have an even less regular solution. In this case, indeed, the sophisticated decision maker may not have a threshold policy. To see this, suppose that over the interval $[T - \tilde{\Delta}, T)$, $\tilde{\Delta} > 0$, the decision maker has a cutoff policy that is strictly decreasing. Now consider time $\tilde{t} = T - \tilde{\Delta} - \epsilon$. There are two regions of interest. First, consider $x(\tilde{t})$ close to but above the threshold for time $T - \tilde{\Delta}$. In this case, the previous logic suggests that the decision maker may complete the task for costs slightly higher than $\bar{x}(T-\tilde{\Delta})$. Second, consider a cost $x(\tilde{t}) \ll \bar{x}(T-\tilde{\Delta})$. In this case, the decision maker knows that if she waits an ϵ-instant, then she is guaranteed (as $\epsilon \to 0$) to complete the task at $T-\tilde{\Delta}$ because her cost will remain below $\bar{x}(T-\tilde{\Delta})$. Therefore, just as was the case in the instant before the final deadline, the sophisticated decision maker will also not complete the task for any small but positive cost.

To illustrate these issues it is useful to turn to simulations. In Figure 1, we show the numerical solution for the optimal policy function. Panel (a) depicts the case where the grid of possible costs is quite fine (i.e., ${}^{c_{i+1}}/_{c_i} = (e^{\sigma\sqrt{dt}})^{1/40}$), while panel (b) depicts the case of a coarse cost grid (i.e., ${}^{c_{i+1}}/_{c_i} = e^{\sigma\sqrt{dt}}$). The black shaded areas represent those costs for which the sophisticated decision maker would complete the task at that time. As can be seen, the optimal policy is, decidedly, not a threshold. That said, as the cost grid becomes finer, the black shaded areas where the decision maker completes the task becomes more dense. Therefore, as $dt \to 0$ and ${}^{c_{i+1}}/_{c_i} \to 1$, the decision maker's policy would appear to be a threshold strategy – though, interestingly, one which is decreasing as t approaches the deadline.[9]

Given the differences in the decision rule of a sophisticated decision maker and an exponential decision maker, it is instructive to examine the likelihood that decision makers of various types (i.e., exponential, naïve, and sophisticated hyperbolic) would complete the task. To this end, we simulated 1,000,000 sample paths of cost realizations using the same underlying parameters as in Figure 1. Figure 2 plots the average cumulative probability of task completion for these three different decision makers. It is useful to present results in terms of the probability that a hyperbolic decision maker maintans control in an instant dt, which we denote Prob(Control). In fact, Prob(Control) $= 1 - e^{-\lambda dt}$ and is inversely related to λ: the instantaneous gratification case, where $\lambda \to \infty$ is represented by Prob(Control) $\to 0$. The panels differ in that Prob(Control) varies. Across all parameterizations of the probability of maintaining control, we see that the naïve decision maker's likelihood of completing the task is very flat on $t \in [0, 1)$ and then we see a significant jump at the deadline, $t = 1$. This is because if the cost is low enough, the naïve decision maker will complete the task immediately. However, if not, because she believe that her future self (when she loses control) is exponential, the decision maker prefers to wait until the deadline.

On the other hand, the sophisticated hyperbolic decision maker's cumulative likelihood of completing the task is increasing and, for Prob(Control) small, generally concave, owing

[9]Indeed, our numerical work suggests that this is particularly a problem as $\lambda \to \infty$ At lower levels of λ, the decision maker is "closer" to exponential, that is, she suffers less from the problems of discretization and so does have a threshold. However, the threshold is not necessarily monotonic.

Figure 1: Optimal Policy of Sophisticated Decision-Maker ($\lambda = \infty$; i.e., Prob(Control)= 0)

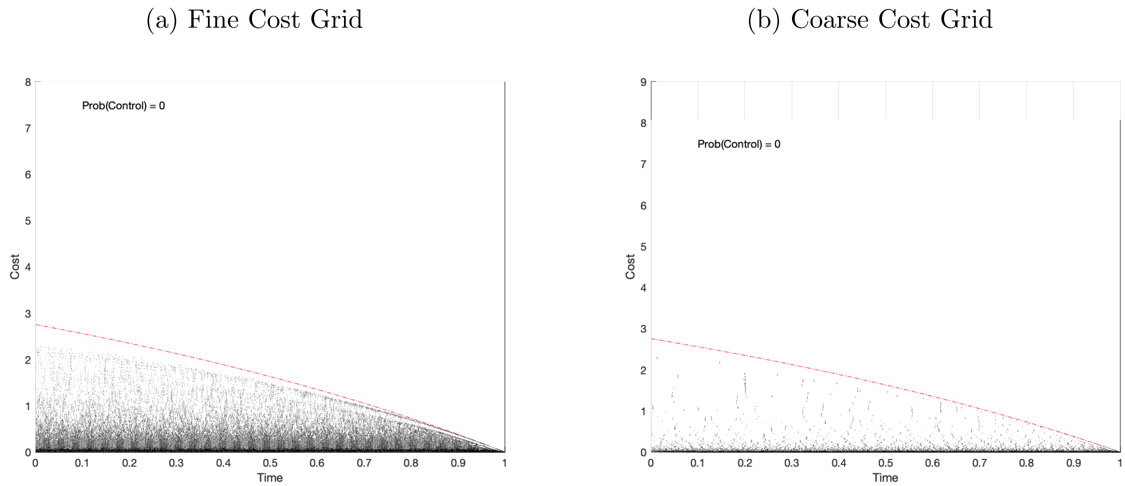

Note 1: On the horizontal axis is time, with $t = 0$ representing the start of the task and $t = 1$ representing the deadline. We set $\beta = 0.8$, $\sigma = 0.5$, $\rho = 0.1$ and $V = 10$. The policy was solved via backward induction, using the binomial method of Cox, Ross, and Rubinstein (1979) with 1440 discrete time intervals to approximate the geometric Brownian motion cost process.
Note 2: The black shaded areas represent those cost realizations for which the decision maker would complete the task at that time.
Note 3: The dashed, red line represents the line of indifference between completing the task "now" and waiting until the deadline to complete the task for sure at the same cost.

to the monotonically decreasing threshold. Although masked in the figure by the line for the naïve decision maker, there is also a large jump in task completions at the deadline for the sophisticated hyperbolic decision maker.

Panel (d), which considers the case of Prob(Control) = 0.8, deserves special attention in that it presents a counter-intuitive result in that there are more time 0 completions by both naïve and sophisticated decision makers than by an exponential decision maker. The intuition for this is as follows. For an exponential decision maker (or, equivalently, a hyperbolic decision maker who never loses control, Prob(Control) = 1), the threshold is strictly increasing: as the deadline approaches, there is less time to wait for a more favorable cost realization and so the decision maker completes the task at ever increasing costs. When Prob(Control) < 1, so that a hyperbolic decision maker sometimes loses control, this creates a temptation to wait. For Prob(Control) close to 1, the threshold is still monotone increasing but is flatter than that of an exponential decision maker. Now consider a sophisticated hyperbolic decision maker and an otherwise identical exponential decision maker at time t and suppose that, in both cases, there current cost exceeds their respective thresholds by the same amount. Because the hyperbolic decision maker's threshold is flatter, this means that she can expect to wait longer before her cost enters the completion region and, therefore,

Figure 2: Simulation of Cumulative Probability of Task Completion

(a) Prob(Control) ≈ 0

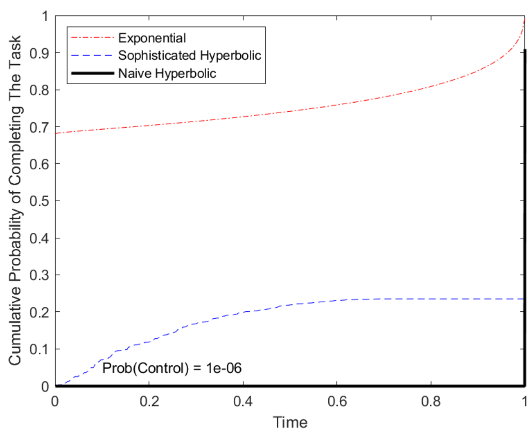

(b) Prob(Control) $= 0.1$

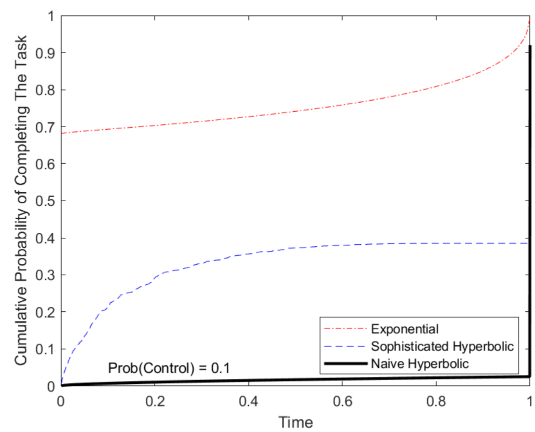

(c) Prob(Control) $= 0.2$

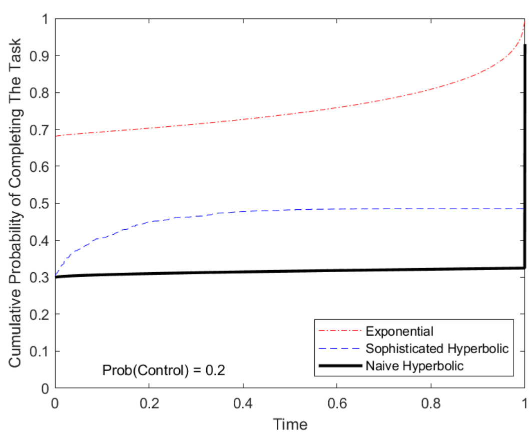

(d) Prob(Control) $= 0.8$

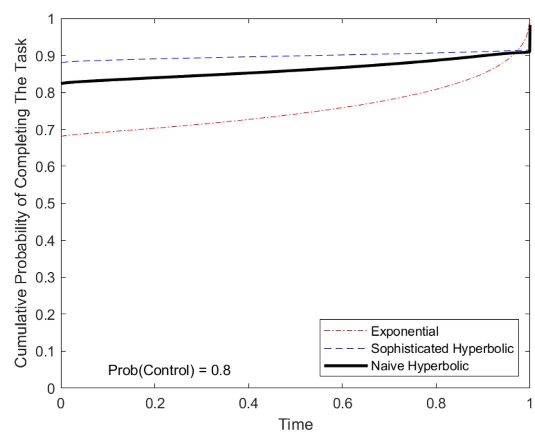

Note 1: On the horizontal axis is time, with $t = 0$ representing the start of the task and $t = 1$ representing the deadline. We set $\beta = 0.8$, $\sigma = 0.5$, $\rho = 0.1$ and $V = 10$. The policies were solved via backward induction, using the binomial method of Cox, Ross, and Rubinstein (1979) with 1440 discrete time intervals to approximate the geometric Brownian motion cost process. In each case, the initial cost was equally likely to be any value on the grid between 1 and 10.
Note 2: For each of the cases, to ensure comparability, the same random numbers were used. The figure depicts the average over 1,000,000 sample paths.

the value from waiting is lower. Because of this, the hyperbolic decision maker actually completes the task at a higher cost than the exponential decision maker. Note that this logic applies when sufficiently far from the deadline. Near the deadline, the hyperbolic decision maker prefers to wait because she is confident that she will complete the task at the deadline, thereby deferring the costs to her future self.

4.3 The Optimal Deadline

We now turn to studying optimal deadlines for exponential and hyperbolic decision makers.

We show first that an exponential decision maker prefers a later deadline to an earlier one; that is, in other words, she will never self-impose a deadline. 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 5. *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. □

A similar result holds for a naïve hyperbolic decision maker with $\lambda < \infty$. It is in fact an immediate consequence of Proposition 4 that, in the instant gratification case, a naïve hyperbolic decision maker will never self-impose a binding deadline, believing herself to be exponential at every future instant:

Proposition 6. *In the limit as $\lambda \to \infty$, a naïve hyperbolic decision maker prefers later to earlier deadlines.*

While we cannot solve analytically the optimal deadline problem for a sophisticated hyperbolic decision maker we illustrate here the results of the intensive numerical analysis we conducted. In particular, we chose a broad set of parameters to solve for the ex-ante expected payoff of both a naïve hyperbolic and a sophisticated hyperbolic decision maker. We varied the amount of time the decision maker had to complete the task on the interval $t = [0, 1]$ with time intervals $dt = 1/1440$. Specifically, the parameters of interest were $\{\beta, \rho, \sigma, p_c, c_l, c_u, \Delta\}$, where $p_c := e^{-\lambda dt}$ is the probability the decision maker maintains control in a given instant; c_l (resp. c_u) represents the lower (resp. upper) bound on the interval of possible initial costs; and Δ is the time between task completion and payment.[10] Our interest in this exercise is to determine how the underlying parameters influence hyperbolic decision makers' willingness to set a deadline.

We first note that in our numerical study, naïve hyperbolic decision maker's *never* set a binding deadline. Therefore, the discussion that follows focuses entirely on sophisticated hyperbolic decision makers. Second, note that if a decision maker imposes a binding deadline, then in the overwhelming majority of cases, the decision maker imposes the tightest possible deadline.[11] As we saw in Figure 2, for present-biased decision makers, most task completions occur either very early on or at the deadline. Therefore, the main thing that intermediate deadlines do is move the mass of task completions that would have occurred at ultimate deadline earlier to the self-imposed earlier deadline. Unless ρ is very high, this is of limited value.

Given the multidimensional nature of our numerical exercise, we present the data in Figure 3. In each subfigure we plot the frequency of choosing a binding deadline against the probability of maintaining control in an instant, dt focusing on one of the parameters of interest. In each figure, we plot one or two representative scatter plots, depending on whether the relationship is quite similar across different parameter values (as it is for β, and the initial cost bounds) or different (as it is for σ, ρ and Δ). As can be seen, and not surprisingly, the frequency of imposing a deadline is generally (weakly) decreasing in the probability of maintaining control.[12] We can also see that as σ increases, the frequency with which deadlines are set decreases. This is intuitive because as costs become more variable, the likelihood of entering the region where the task is completed increases. In contrast, as ρ increases – meaning that the decision maker is more impatient – the frequency of imposing a deadline increases. This too is intuitive because by imposing a deadline, it moves forward the time at which the decision maker completes the task, which is valuable because the decision maker is impatient. Lastly, we see that as the length of time the decision maker must wait increases, so too does the frequency of setting a deadline. This is because the longer time the

[10]Our parameters were, $\beta \in \{0.6, 0.7, 0.8, 0.9\}$, $\sigma \in \{0.1, 0.2, 0.3\}$, $\rho \in \{0.02, 0.05, 0.1, 0.15, 0.2\}$, $c_l \in \{0.5, 1, 1.5\}$, $c_u \in \{10, 11.67, 13.33\}$, $\Delta \in \{dt, 5dt\}$ and $p_c \in \{0.000001, 0.001, 0.01, 0.1, 0.5, 0.8, 0.9, 0.95, 0.99, 0.999, 0.9999\}$ and our numerical exercise was balanced, leading to 11,800 different parameter combinations.

[11]Indeed, 42.96% of the trials result in a deadline on the subinterval $t \in [0, 0.05]$, 0.94% of the trials result in a deadline on the subinterval $t \in (0.05, 0.95]$ and 56.1% of trials involve a deadline on the subinterval $t \in (0.95, 1]$, where 1 represents the case of no deadline. Importantly, of these 56.1% remaining, nearly all involve no deadline.

[12]However, for β, there is a slight non-monotonicity at low probabilities of maintaining control.

decision maker must wait for payment, the greater the chance that the decision maker will lose control, which increases the temptation to delay. Thus, an increase in Δ is closely related to a reduction in p_c. The interested reader can see a complete set of tables in Appendix B which shows the relationship between deadlines and the probability of maintaining control for all of the individual parameter values that we considered.

While Figure 3 shows how the decision to set a deadline varies with parameters, it is also of interest to see how beneficial are deadlines. To this end, we calculate the percentage increase in the ex ante expected payoff when the decision maker sets the best deadline relative to the case of not setting any deadline. In Figure 4, we report histograms of the percentage increase in expected payoff (conditional on a deadline being optimal). As can be see, there is a wide range of benefits with the 10^{th}, 25^{th}, 50^{th}, 75^{th} and 90^{th} percentiles being 0.33%, 1.45%, 3.46%, 7.13% and 10.31%, respectively.

In Table 1 we report regression results for the decision to set a deadline (column: "Set DL") as well as the benefit from setting a deadline (other columns) to see how the parameters effect the decision to set a deadline and the benefit from a deadline while controlling for the other variables. For the decision to set a deadline, conducted a simple logistic regression where the dependent variable takes value 1 if the decision maker would set a deadline and 0 otherwise.[13] The reported coefficients are marginal effects. Table 1 confirms some findings from Figure 3 above. In particular, the likelihood of setting a deadline is increasing in ρ and Δ (the time to wait for payment) and is decreasing in the probability of maintaining control and in σ. We also see that the likelihood of setting a deadline is decreasing in both the lower and upper bound of the initial cost distribution. One somewhat counterintuitive result is the positive marginal effect on β. Closer inspection reveals that the relationship is actually non-monotonic. For low values of β, the likelihood of setting a deadline is increasing in β. However, for high enough values of β it is decreasing.

The two other columns of the table report results of OLS regressions on the percentage increase in expected payoff from choosing the optimal deadline, with one including all data ("Uncond. Benefit") and the other conditioning on those instances in which a deadline is optimal ("Cond. Benefit"). As can be seen, as the probability of maintaining control increases, the benefit from a deadline decreases. Deadlines are also less valuable as σ or c_u increase and they are more valuable as Δ, ρ or c_l increase. Lastly, when it comes to the present-bias parameter, β, we see that the relationship is non-monotonic, first increasing and then decreasing. Thus, deadlines are both most frequent and most beneficial at intermediate

[13] To be sure, we do not mean to imply that the decision to set a deadline is probabilistic. We are simply using this framework to represent the marginal effects of a multidimensional relationship. In the model, we include both Prob(Control) and Prob(Control^2 to capture a potential non-linear relationship in the probability of maintaining control.

Figure 3: Frequency of Deadline Choice Given Probability of Maintaining Control

(a) $\beta = 0.7$

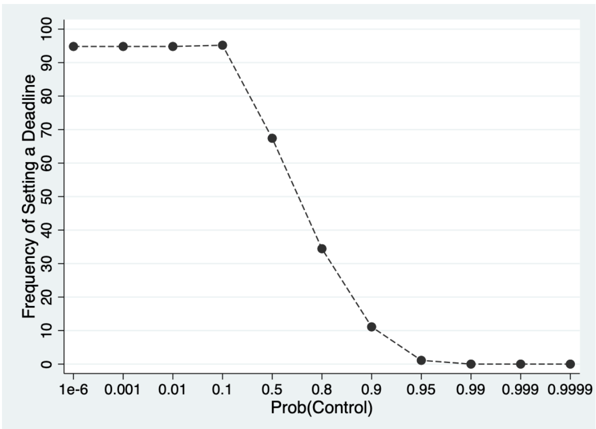

(b) $\sigma \in \{0.1, 0.3\}$

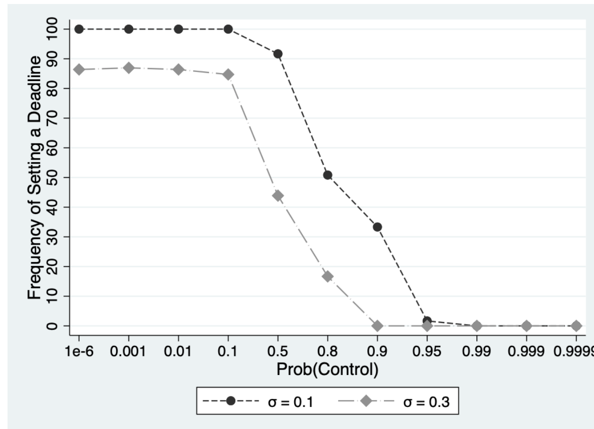

(c) $\rho \in \{0.02, 0.1\}$

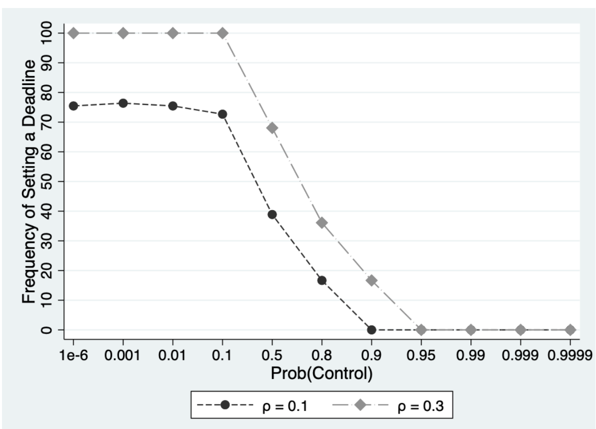

(d) Time To Wait For Payment, $\Delta \in \{dt, 5dt\}$

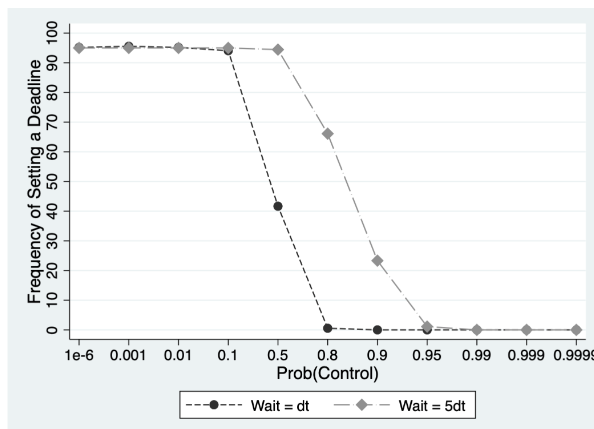

(e) Lower Bound on Initial Cost = 1

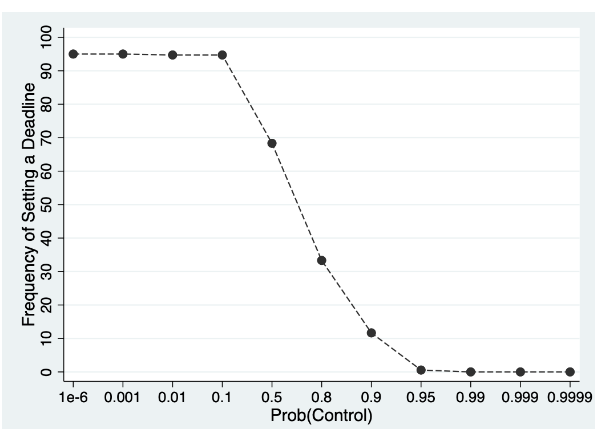

(f) Upper Bound on Initial Cost = 10

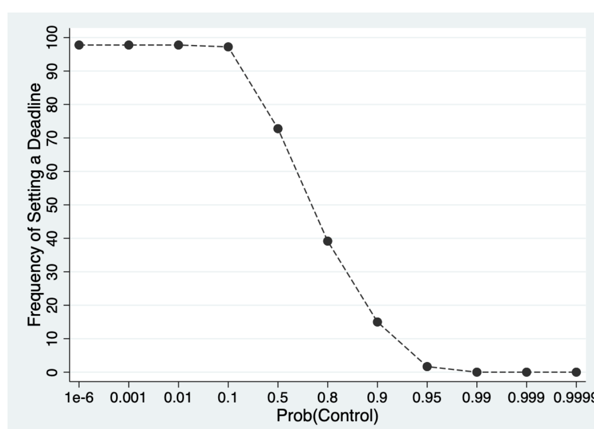

Figure 4: Percentage Increase in Ex Ante Expected Payoff From Imposing Deadline (Conditional on Deadline Optimal)

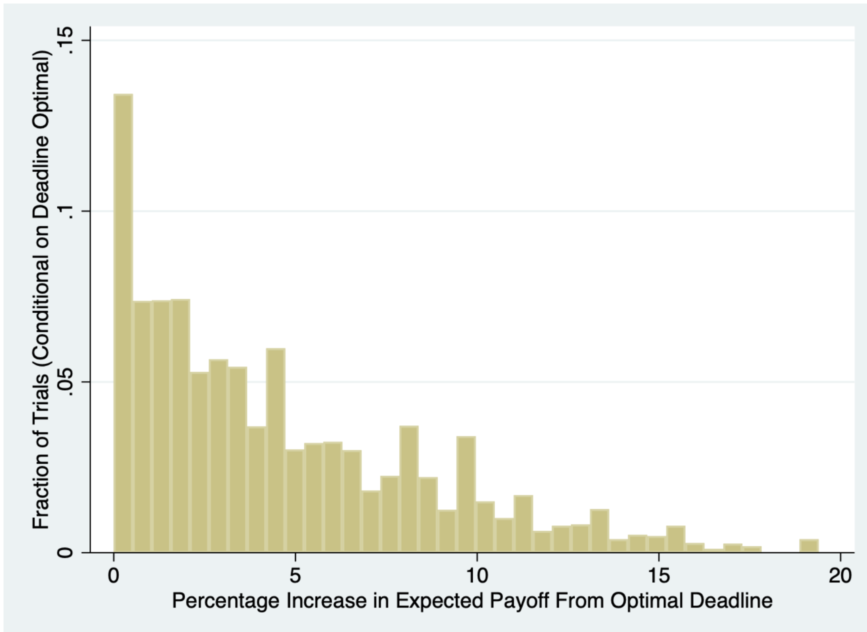

Table 1: Marginal Effects of Logit Regression on Decision to Impose Deadline

Parameter	Set DL		Uncond. Benefit		Cond. Benefit	
Prob(Control)	-0.533***	(0.009)	-7.152***	(0.258)	-2.756***	(0.420)
Prob(Control)^2			3.077***	(0.267)	-6.311***	(0.548)
$\mathbf{1}[\Delta = 5dt]$	0.130***	(0.003)	0.495***	(0.042)	1.315***	(0.063)
β	0.041***	(0.016)	7.643**	(3.010)	15.494***	(4.424)
β^2.			-9.494***	(2.003)	-20.261***	(2.943)
σ	-0.849***	(0.026)	-4.200***	(0.254)	-10.256***	(0.374)
ρ	0.883***	(0.030)	17.540***	(0.315)	38.399***	(0.467)
c_l	-0.014***	(0.004)	0.783***	(0.049)	1.762***	(0.072)
c_u	-0.012***	(0.001)	-0.024	(0.015)	-0.086***	(0.022)
$\mathbf{1}$[set DL]			1.279***	(0.085)		
Constant			1.982*	(1.128)	1.858	(1.656)
Observations	11880		11880		5331	
	Logit, Marginal Effects		OLS		OLS	

levels of present-bias.

5 Conclusions

In this paper we examined the behavior 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. Exponential decision makers correctly value the option value of waiting to complete the task and follow a threshold rule which is increasing as time approaches the deadline. The behavior of hyperbolic decision makers is more challenging to characterize and is closely connected to the parameter λ, which captures the likelihood that the present-self will maintain control in an instant. As $\lambda \to \infty$, the present-self is in control only for an instant. In this case, we are able to show that naive hyperbolic decision makers *never* complete the task strictly before the deadline. For sophisticated hyperbolic decision makers, who are aware of their self-control problems, matters are more complex and analytical solutions are difficult to derive. However, we argue that the threshold of such decision makers may actually be monotonically *decreasing* and that, in the instant before the final deadline, the decision maker will never complete the task, knowing that she will complete it at the deadline.

While we know that exponential decision makers will never impose a deadline, in a numerical study, we also demonstrate that naïve hyperbolic decision makers will, similarly, never self-impose a binding deadline. However, 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, which is behind the identification strategy of Bisin and Hyndman (2020) and Ariely and Wertenbroch (2002), among others. Both of these papers show a robust demand for commitment via self-imposed deadlines.[14] Our numerical

[14] Another distinguishing feature – and thus (partial) identification strategy – is that exponential decision makers will always have a threshold which is strictly increasing. On the other hand, hyperbolic decision makers with λ sufficiently large have either flat or decreasing. Additionally, there will be a discontinuous jump in completions at the deadline for hyperbolic decision makers, while for exponential decision makers, no such jump occurs at the deadline.

study confirmed several intuitions. First, deadlines are more common and more valuable the higher is the discount rate, ρ, and are less common and less valuable the higher is the variability of costs, σ. Second, deadlines are more common and more valuable the *lower* the chance of the decision maker maintaining control and the longer the decision maker must wait before being paid. However, our numerical exercise also showed a non-monotonicity in the present-bias parameter, β. In particular, deadlines are more common and more valuable when they are in an intermediate range. For β sufficiently low or sufficiently high, deadlines become less common and less beneficial to the sophisticated hyperbolic decision maker.

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.[15] This could be why Ariely and Wertenbroch (2002) found that evenly spaced deadlines were the most effective. However, Bisin and Hyndman (2020) call into question the robustness of this result as they show that subjects performed the worst under exogenous and evenly spaced deadlines.

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 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.

It would be interesting to study whether other forms of commitment are possible in this

[15] In a discrete time model more suitable for analyzing experimental data, Bisin and Hyndman (2020) 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.

basic framework. For example, Trope and Fishbach (2000) experimentally study whether subjects are willing to make a fixed payment conditional on task completion or are willing to impose a cost on themselves for not completing a task. We believe that it is highly likely that sophisticated decision makers as in our model might be willing to engage in such alternate forms of commitment, but we leave this question to future research. Although the quasi-hyperbolic model of decision-making has been well-studied in the literature, it is not the only way to study the broad topic of self-control problems. It is also interesting to study whether decision makers motivated not by present-bias but some other potential bias may be willing to commit themselves to a binding deadline. In Appendix A, we briefly consider decision makers who misperceive their true cost process as well as decision makers who suffer from temptation à la Gul and Pesendorfer (2001, 2004). We show that decision makers in the former case may self-impose a binding deadline, while they would not do so in the latter case. More research in this area would certainly be fruitful.

References

Andersen, T. M., and J. Bhattacharya (2021): "Why mandate young borrowers to contribute to their retirement accounts?," *Economic Theory*, 71, 115–149.

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 (2020): "Present-Bias, Procrastination and Deadlines in a Field Experiment," *Games and Economic Behavior*, 119, 339–357.

Brunnermeier, M., F. Papakonstantinou, and J. Parker (2017): "Optimal Time-Inconsistent Beliefs: Misplanning, Procrastination, and Commitment," *Management Science*, 63(5), 1318–1340.

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.

Calcott, P., and V. Petkov (2021): "Excessive consumption and present bias," *Economic Theory*, Forthcoming.

Cao, D., and I. Werning (2016): "Dynamic Savings Choice With Disagreements," NBER Working Paper 22007.

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.

Ericson, K. M., and D. Laibson (2019): "Intertemporal Choice," in *Handbook of Behavioral Economics*, ed. by B. D. Bernheim, S. DellaVigna, and D. Laibson, vol. 2, chap. 1, pp. 1–67. Elsevier.

Gabaix, X., and D. Laibson (2017): "Myopia and Discounting," NBER Working Paper #23254.

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.

Heidhues, P., and P. Strack (2019): "Identifying Present-Bias from the Timing of Choices," SSRN Working Paper 3386017.

Houser, D., D. Schunk, J. Winter, and E. Xiao (2018): "Temptation and Commitment in the Laboratory," *Games and Economic Behavior*, 107, 329–344.

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.

Ikeda, S., and T. Ojima (2021): "Tempting goods, self-control fatigue, and time preference in consumer dynamics," *Economic Theory*, Forthcoming.

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.

Vieille, N., and J. W. Weibull (2009): "Multiple solutions under quasi-exponential discounting," *Economic Theory*, 39, 513–526.

A 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.

A.1 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 (2017), 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 $\alpha > 0$.[16] 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.

A.2 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.[17]

[16] This is different from Gabaix and Laibson (2017) who also consider a model of misperceptions. In their model, the uncertainty about an event is increasing in the time to the event.

[17] Indeed, the literature on temptation and self-control is rich. See Ikeda and Ojima (2021) for a recent study, which also contains an up-to-date literature review on this interesting area.

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 what follows, 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.

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

Proof. The proof is by induction. 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 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 Tables: Frequency of Deadlines By Parameters

In the tables below, we provide more detailed results on the frequency with which deadlines are set. These tables complement and expand upon Figure 3.

Table 2: How Deadlines Vary With Time to Wait For Payment and Prob(Control)

	Time to Wait For Payment	
Prob(Control)	dt	$5dt$
1.00E-06	95.19	95.00
0.001	95.56	95.00
0.01	95.19	95.00
0.1	94.07	95.00
0.5	41.67	94.44
0.8	0.56	66.11
0.9	0.00	23.33
0.95	0.00	1.11
0.99	0.00	0.00
0.999	0.00	0.00
0.9999	0.00	0.00

Table 3: How Deadlines Vary With Cost Upper Bound and Prob(Control)

	Upper Bound of Initial Cost		
Prob(Control)	10.00	11.67	13.33
1.00E-06	97.78	94.44	93.06
0.001	97.78	94.72	93.33
0.01	97.78	94.44	93.06
0.1	97.22	93.61	92.78
0.5	72.78	66.67	64.72
0.8	39.17	32.50	28.33
0.9	15.00	10.00	10.00
0.95	1.67	0.00	0.00
0.99	0.00	0.00	0.00
0.999	0.00	0.00	0.00
0.9999	0.00	0.00	0.00

Table 4: How Deadlines Vary With Cost Lower Bound and Prob(Control)

	Lower Bound of Initial Cost		
Prob(Control)	0.5	1	1.5
1.00E-06	96.94	95.00	93.33
0.001	97.50	95.00	93.33
0.01	97.22	94.72	93.33
0.1	95.56	94.72	93.33
0.5	68.89	68.33	66.94
0.8	33.33	33.33	33.33
0.9	11.67	11.67	11.67
0.95	0.56	0.56	0.56
0.99	0.00	0.00	0.00
0.999	0.00	0.00	0.00
0.9999	0.00	0.00	0.00

Table 5: How Deadlines Vary With σ and Prob(Control)

	σ		
Prob(Control)	0.1	0.2	0.3
1.00E-06	100.00	98.89	86.39
0.001	100.00	98.89	86.94
0.01	100.00	98.89	86.39
0.1	100.00	98.89	84.72
0.5	91.67	68.61	43.89
0.8	50.83	32.50	16.67
0.9	33.33	1.67	0.00
0.95	1.67	0.00	0.00
0.99	0.00	0.00	0.00
0.999	0.00	0.00	0.00
0.9999	0.00	0.00	0.00

Table 6: How Deadlines Vary With β and Prob(Control)

	β			
Prob(Control)	0.6	0.7	0.8	0.9
1.00E-06	92.59	94.81	97.41	95.56
0.001	92.59	94.81	97.41	96.30
0.01	92.59	94.81	97.04	95.93
0.1	92.22	95.19	95.19	95.56
0.5	66.30	67.41	71.11	67.41
0.8	34.44	34.44	34.44	30.00
0.9	11.11	11.11	12.22	12.22
0.95	0.00	1.11	1.11	0.00
0.99	0.00	0.00	0.00	0.00
0.999	0.00	0.00	0.00	0.00
0.9999	0.00	0.00	0.00	0.00

Table 7: How Deadlines Vary With ρ and Prob(Control)

	ρ				
Prob(Control)	0.02	0.05	0.1	0.15	0.2
1.00E-06	75.46	100.00	100.00	100.00	100.00
0.001	76.39	100.00	100.00	100.00	100.00
0.01	75.46	100.00	100.00	100.00	100.00
0.1	72.69	100.00	100.00	100.00	100.00
0.5	38.89	66.67	68.06	81.94	84.72
0.8	16.67	22.22	36.11	41.67	50.00
0.9	0.00	5.56	16.67	16.67	19.44
0.95	0.00	0.00	0.00	0.00	2.78
0.99	0.00	0.00	0.00	0.00	0.00
0.999	0.00	0.00	0.00	0.00	0.00
0.9999	0.00	0.00	0.00	0.00	0.00