Setbacks, Shutdowns, and Overruns*

Felix Feng Curtis Taylor Mark Westerfield Feifan Zhang

March 29, 2022

Abstract

We employ novel methods to investigate optimal project management in a setting plagued by unavoidable setbacks. The contractor can cover up delays from shirking either by making false claims of setbacks or by postponing the reports of real ones. The sponsor induces work and honest reporting via a soft deadline and a reward for completion. Late-stage setbacks trigger randomization between cancellation and extension. Thus the project may run far beyond its initial schedule, generating arbitrarily large overruns, and yet be canceled. Absent commitment to randomize, the sponsor grants the contractor more time to complete the project.

JEL Classifications: C72, D21, D86, M11
Key Words: Project Management, Dynamic Agency, Soft Deadline, Cost-Plus-Award-Fee Contract, Schedule Slippage.

---

*Feng (ffeng@uw.edu) and Westerfield (mwesterf@uw.edu) are from the University of Washington. Taylor (crtaylor@duke.edu) and Zhang (feifan.zhang@duke.edu) are from Duke University. We are grateful to Nageeb Ali, Arjada Bardhi, Dirk Bergemann, Ralph Boleslavsky, Alessandro Bonatti, Dino Gerardi, George Georgiadis, Brett Green, Bård Harstad, Zhiguo He, Ayca Kaya, Aaron Kolb, Fei Li, Moritz Meyer-ter-Vehn, Takuo Sugaya, and Mike Waldman for helpful comments and suggestions. We similarly thank seminar and conference participants at: Cornell's JGSM-SBE, the Econometric Society's 2021 NASM, The Finance Theory Group Fall 2021 conference, Hong Kong UST, The 2021 Midwest Theory conference, the spring 2021 meeting of the NBER Organizational Economics Working Group, the 2021 SFS Cavalcade, and the 2021/22 Webinar Series on Credence Goods and Markets for Expert Services. The usual disclaimer applies.

> ***Hofstadter's Law:*** *It always takes longer than you expect, even when you take into account Hofstadter's Law.*
>
> —Douglas Hofstadter,
> *Gödel, Escher, Bach: An Eternal Golden Braid* (1979)

Broad and expanding swaths of the modern economy are dedicated to the planning and execution of projects, "temporary endeavor[s] undertaken to create a unique product, service or result ... The development of software for an improved business process, the construction of a building or bridge, the relief effort after a natural disaster, the expansion of sales into a new geographic market all are projects."[1] Given the current and growing significance of this mode of production, it is important to understand its intrinsic characteristics and – in particular – how best to improve its efficacy. Indeed, the annals of project management are rife with jobs that ran notoriously over time and over budget, some of which were ultimately canceled by their sponsors resulting in little if any residual value.

For example, what might be "the most highly publicized software failure in history" (Goldstein, 2005) is the FBI's contracting debacle with SAIC to develop a virtual-case-file-(VCF) system. Irigoyen (2017) summarizes the VCF project failure as going through "significant management and implementation problems and cost overruns, which culminated in the cancellation of the project in 2005, with little to show for the USD170 million investment." The FBI Director at the time, Robert S. Mueller, III, testified before a Congressional subcommittee that he was disheartened by "the setbacks which have plagued this program..."[2]

The FBI is hardly alone in its project management woes. For instance, "According to a 2017 report from the Project Management Institute, 14 percent of IT projects fail. However, that number only represents the total failures. Of the projects that didn't fail outright, 31 percent didn't meet their goals, 43 percent exceeded their initial budgets, and 49 percent were late" (Greene, 2019). The same pattern exists in large scale construction and industrial manufacturing, as illustrated by the high profile cases listed in Table 1

According to Lineberger and Hussain (2016), "The combined cost overrun for Major Defense Acquisition programs in 2015 was $468 billion ... with an average schedule

[1]Excerpted from What is Project Management? (Project Management Institute, 2020)

[2]https://archives.fbi.gov/archives/news/testimony/fbis-virtual-case-file-system.

Table 1

Project	Time Overrun	Cost Overrun	Source
Boston's Big Dig	9 years	190%	Haynes (2007)
Sydney Opera House	10 years	1400%	Wild (2015)
Boeing 787 DreamLiner	3 years	200%	Shenhar et al. (2016)
Berlin's Brandenberg Airport	10 years	300%	Brandt (2020)

delay of 29.5 months." Importantly, "setbacks are a near-universal, and universally costly, experience ... large capital projects are typically 20 months late, and 80% over the original authorized budget" (Billante, 2017). More, setbacks can cause a project to be canceled leaving the sponsor with huge bills and often nothing else. A prime recent example is South Carolina's V.C. Summer nuclear power plant construction project, canceled in 2017 after a series of major setbacks and cost overruns, saddling taxpayers with a bill of \$9 billion and "nothing to show for it" (Lacy, 2019).

In this paper we argue that project setbacks, overruns, and cancellations are not always the product of incompetence or inattention, but – at least to some degree – are unavoidable consequences of optimal project governance in the face of agency frictions. In particular, we introduce a model of project development in which setbacks arise naturally as part of the production process. Examples include discovering: adverse site conditions (construction), a design feature doesn't work as intended (manufacturing), or incompatibility of certain off-the-shelf subroutines (software engineering). Due to unforeseeable contingencies such as these, the amount of time and resources required to complete the project are necessarily uncertain.

In our model, as in practice, the sponsor (the principal) must hire a contractor (the agent) to run the project on her behalf. Both parties are risk-neutral, but the agent is protected by limited liability. Setbacks arrive randomly according to a Poisson process with known intensity. There is a flow cost of running the project, and the project is completed whenever a span of time $\bar{X}$ passes without the arrival of a setback. The first-best policy in this environment is straightforward. The project should be started and run until completed if and only if the value of the finished project to the sponsor exceeds the flow cost of operation times the expected duration.

The contractual setting we investigate is marked by both hidden actions and hidden states. The principal is unable to observe the progress of the project or the occurrence of setbacks herself, and must rely on unsubstantiated reports from the agent. However, delivery of the completed and working project is verifiable – the principal can use the software, fly the plane, or occupy the building once it is complete. Because the principal cannot observe the status of the unfinished project, the agent may surreptitiously divert the flow of operating capital to garner private benefits instead of advancing the project. The combination of hidden actions and hidden states gives the agent broad scope for committing moral hazard. Specifically, he may cover up the interruption of progress associated with resource diversion either by submitting false reports of setbacks or delaying the reports of real ones. Thus, the principal's problem is to write a contract, contingent only on the passage of time and project delivery, that induces the agent to work efficiently and report honestly.

The crucial incentive constraint is what we label the *No-Postponed-Setbacks* (NPS) condition. This constraint requires that whenever a setback occurs, the agent prefers to report it immediately rather than divert resources for any length of time and report it later. We show that (NPS) always binds under an optimal incentive scheme. This has several important implications. First, it implies that the agent also prefers not to cover up resource diversion with claims of false setbacks; that is, binding (NPS) is necessary and sufficient for incentive compatibility. Second, it allows us to fully characterize the optimal contract which closely resembles a *cost-plus-award-fee contract*.[3]

Under the optimal contract, the principal gives the agent a soft deadline or time budget for delivering the completed project and commits to pay the flow cost of operation. The contract ultimately ends for one of two reasons, either because the agent delivers the completed project or because it is canceled by the principal. If the agent delivers the completed project, then he is paid a *linear* reward consisting of a fixed fee plus an incentive award proportional to the time remaining before the soft deadline is exhausted. If a setback is reported when the time remaining on the soft deadline is greater than $\bar{X}$ (the required uninterrupted development time), then the principal takes no action. But if a setback is reported when the time remaining is less than $\bar{X}$, then project completion before the clock runs out is impossible. At this point, the optimal contract calls for a random termination procedure under which the project

[3] "A cost-plus-award-fee contract is a cost-reimbursement contract that provides for a fee consisting of (a)a base amount (which may be zero) fixed at inception of the contract and (b) an award amount, based upon a judgmental evaluation by the Government, sufficient to provide motivation for excellence in contract performance." (The U.S. General Service Administration FAR 16.401). We comment on this implementation more formally in the discussion following Proposition 1 below.

is either canceled with a terminal payment of zero to the agent or the deadline is extended to $\bar{X}$. All subsequent reports of setbacks are treated similarly. Thus, while the optimal incentive contract induces the agent to work diligently and report honestly, it may, nevertheless, result in the type of unfortunate outcomes observed in our leading examples. In other words, schedule and cost overruns, and even cancellations that yield no useful output, are *features* of an optimal contract. These features obtain because randomization is necessary: if the contract had a deterministic deadline, then a late-stage setback would render project completion impossible and the agent would *shirk out the clock*. On the other hand, if there was some sequence of reports that enabled the project to run indefinitely, then the agent would make those reports and shirk forever. The only solution is random termination, which yields the possibility of both overruns and inefficient cancellations.

Our analysis utilizes novel methods that do not involve the usual dynamic programming and differential equation techniques to characterize the principal's value function. In particular, we identify two fundamental martingales and invoke the optional stopping theorem. The principal's expected payoff equals the probability of project completion times the first-best value of the project net of expected agency rents. The probability of project completion is increasing, concave, and approaches 1 as the length of the soft deadline, S, tends to infinity. On the other hand, agency rents increase linearly in S. Hence, there exists a unique optimal initial time budget S^* to assign to the agent at project inception.

Interestingly, the principal's value function is a concave polynomial with kinks at $S = n\bar{X}$. These kinks imply that the optimal initial time budget S^* is an integer multiple of $\bar{X}$ for a non-negligible set of parameters – there are *focal contract lengths* that are multiples of the best-case duration. We use these observations to identify the conditions under which $S^* = \bar{X}$ is optimal; this is a *short-leash* contract in which the soft deadline equals the expected duration of the project and every reported setback results in cancellation with positive probability. Although a short-leash contract has an expected duration of $\bar{X}$, the support of the stopping time is unbounded due to the possibility of multiple project extensions. Hence, even when the principal commits to keep the agent on a short leash, arbitrarily large cost and schedule overruns occur with positive probability. Importantly, every optimal contract possesses a short-leash phase that is triggered whenever a setback occurs sufficiently late in the schedule.

After fully characterizing the optimal contract of our baseline model, we go on to consider two generalizations. First, the optimal contract of the baseline model requires the principal to commit to randomized extension or cancellation with explicit probabilities that are a function of the time remaining when a setback occurs. However, absent such commitment, the principal would strictly prefer to keep the agent working on the project by extending it rather than canceling it. So, we vary the baseline model by relaxing the assumption of full commitment and investigate a setting in which randomization is feasible but not verifiable. In this case, randomization by the principal is incentive compatible if and only if the agent, upon receiving an extension, himself randomizes between continuing to work and shirking away the granted time. While relaxing commitment is clearly harmful to the principal, we use the martingale methods for characterizing the value function to show that she optimally grants the agent a *longer* initial schedule in this setting. Intuitively, she does so in order to raise the likelihood that the project will be completed before the problematic short-leash phase of the schedule is ever reached. Put differently, granting large S^* not only increases the chance the agent will successfully complete the project, it decreases the chance that lack of commitment will ever come into play. Under this version of the contract, the principal commits to let the agent run the project for an initial interval of time after which his continued employment may be regarded as *at-will.*

In the second generalization, we consider a setback environment that encompasses numerous possible variations of the baseline model. These include: multi-stage projects where setbacks occur within stages, fractional setbacks where the magnitude of each mishap is a random variable between 0 and 1, and setbacks of bounded scope where the amount of progress lost is capped. In this general environment, we show that the agency frictions associated with limited liability and private information regarding the state of progress are necessarily controlled by optimal contracts with a similar structure. Instead of a time budget that counts down deterministically absent a setback, the principal offers an expected time budget that counts down on average. However, the principal's incentive devices remain control of time (i.e. termination) and a payment upon completion.

1 Related Literature

The literature on the optimal provision of incentives in dynamic environments is extensive and active. Pioneering articles responsible for moving it forward at various stages include Spear and Srivastava (1987), Phelan and Townsend (1991), Quadrini (2004), Clementi and Hopenhayn (2006), DeMarzo and Sannikov (2006), Biais, Mariotti, Plantin, and Rochet (2007), Sannikov (2008), and Williams (2011).

Bergemann and Hege (1998) investigate venture capital financing in a discreet-time model where the arrival of revenues depends on whether the project is good or bad and whether the entrepreneur (agent) works or shirks. The dynamic agency costs may be high and lead to an inefficient early termination of the project. Toxvaerd (2006) considers a setting in which a finite number of observable arrivals are needed in order to complete a project. In his setting, the agent is risk averse and the optimal contract trades off optimal risk-sharing for incentive provision, but does not involve deadlines or inefficient termination. Biais et al. (2010) analyze a model in which large observable losses may arrive via a Poisson process, and an agent must exert hidden effort in order to minimize the likelihood of their arrival. In contrast with these papers, the Poisson shocks in our model are privately observed by the agent and arise as an unavoidable consequence of the production process – that is they are *discoveries*. The potential for their occurrence essentially gives the agent cover to commit moral hazard; i.e., to make plausible excuses for why project completion has been delayed.[4]

Most closely related are four recent papers (one publication and three working papers) that explore the optimal deadline for a project in the context of dynamic agency. The published article is Green and Taylor (2016) who study a setting where a project must have two Poisson breakthroughs in order to be completed. The agent hired to run the project privately observes the occurrence of the first breakthrough, or what the authors call *progress*. As in our setting, the agent can surreptitiously divert the principal's flow of investment in the project for private benefit, which delays progress. However, once the agent reports a breakthrough, there is no turning back, which limits his scope for further manipulation. The optimal contract in Green and Taylor (2016) features an intermediate deadline such that the agent is subjected to termination at a constant rate if he has not reported progress by that point. If the agent reports progress during this probationary phase, then he is given a relatively short deterministic amount of time to finish.

[4]Other papers that share some features with the environment we study include: Lewis (2012), Mason and Välimäki (2015), Rahmani, Roels, and Karmarkar (2017), Vasama (2017), and Hoffmann, Inderst, and Opp (2020).

We consider a richer and very different environment where progress corresponds to a continuum of states and in which a potentially infinite number of setbacks may occur en route to project completion. Thus our agent may repeatedly report the occurrence of false setbacks or repeatedly postpone reporting the occurrence of real ones, or any combination of these. The type of moral hazard problem this generates is not addressed in Green and Taylor (2016) or – to our knowledge – anywhere else in the literature. Specifically, setbacks are not a consequence of the agent shirking, but are naturally discovered while he works on the project. Because their occurrence is privately observed, the agent can make plausible excuses for overruns in schedule and budget. The agent in our model is not disciplined for failing to report progress, but for reporting the occurrence of late-stage setbacks, and the probability of termination is not constant, but depends on the state of the project when a setback is announced. There is no analogue of a short-leash contract or of partial setbacks in Green and Taylor (2016), and they do not consider relaxing commitment.

In an insightful working paper, Madsen (2021) studies how an organization should optimally manage a project of uncertain scope when advised by an expert with private information about the project's state who prefers to prolong his employment. In this model, a project turns from "good" to "bad" stochastically over time. The agent is a "advisor", who possesses private information regarding whether the project quality has changed and must be incentivized to report this. Mayer (2021) presents a dynamic contracting model in which a project succeeds if it survives until the completion date. While the project is in operation, the agent exerts unobservable precautionary effort in order to reduce the arrival rate of a privately observed failure shock that will kill the project before it reaches completion. As in Madsen (2021), the principal must provide incentives for the agent to report that the project has gone bad and should be terminated. Yet a third recent paper featuring a single privately observed transition is Curello and Sinander (2021). Similar in spirit but opposite in application to Madsen (2021), in this model, a technological breakthrough occurs exogenously at some random time witnessed only by the agent. The principal would like to adopt the innovation as soon as possible, but the agent prefers the *status quo* technology. Hence, the agent must be incentivized, through non-monetary means, to disclose the arrival of the innovation.

Our setting clearly differs from those studied in these three papers along a number of salient dimensions. Rather than the arrival of a single transition, our agent may observe numerous setbacks, none of which render project completion infeasible. Indeed, it is common knowledge from the outset that finishing the project is efficient. Our agent is not an advisor hired to monitor whether project quality has changed – His expertise resides in the ability to operate the project itself. This provides him with an informational advantage that the principal manages through implementation of a dynamic delivery-contingent contract.

2 The Model and The First-Best

2.1 The Project

A risk-neutral principal (she) hires a risk-neutral agent (he) over an infinite horizon to work on a project. The principal has deep pockets, and the agent has no wealth and is thus protected by limited liability. The project requires accumulated progress $\bar{X}$ before it is completed; $\bar{X}$ is the project's *scope*. As the agent works on the project, *progress* X_t accumulates deterministically. However, *setbacks* occur, following a Poisson process with arrival rate λ, which is the setback *frequency*. A setback at t resets progress from X_t to 0. When progress reaches $\bar{X}$, the project is complete and results in a monetary payoff of R to the principal. While the project is in operation, the principal must pay a flow cost of c to keep it running.

Two points are worth highlighting. First, for simplicity we assume that an incomplete project has no value to the principal. Second, setbacks result naturally as a result of unforeseeable contingencies, and, in particular, setbacks are not due to the negligence or indolence of the agent.[5]

The potential for moral hazard in this setting stems from the ability of the agent to surreptitiously divert the resource flow c to his own private benefit and cover the resulting cessation in progress by misinforming the principal about the occurrence or

[5] In other words, we model setbacks as discoveries resulting from working on the project, not from shirking. Unforeseeable contingencies are discovered that make the required time and resources uncertain. One interpretation is that there is a path to complete the project, and its length must be discovered through trial and error. Another is that there are many ex-ante equivalent paths by which the project may be completed, and each fails with probability $1 - e^{-\lambda \bar{X}}$. If, instead, setbacks occurred as a consequence of shirking, then the first-best could be implemented by setting a hard deadline of $\bar{X}$ and paying the agent a fixed award upon project delivery. In any case, we investigate the possibility of partial setbacks in Section 6.

timing of setbacks. Formally, the project's true progress follows

$$dX_t = a_t(dt - X_t dN_t), \tag{1}$$

where $a_t \in \{0, 1\}$ denotes the agent's private action. $a_t = 0$ represents shirking, which corresponds to diversion of the resource stream c, while $a_t = 1$ represents working, which corresponds to using the funds to develop the project. Shirking yields the agent a private flow benefit of b:

Assumption 1 *$b < c$, so shirking (or diversion) is socially inefficient.*

Whenever the agent shirks, progress on the project remains constant; i.e., setbacks are discovered only if the agent is working. Both the principal and agent are perfectly patient and possess outside options of zero.[6]

2.2 The First-Best

If the agent's actions are publicly observable, then the principal can induce his compliance without incurring additional cost. Clearly, if it is worth starting the project in the first place, then it is worth running it until it is eventually completed.

Suppose that the project is operated until completed and let F^{FB} be the value to the principal at inception. Because the time between setbacks is exponentially distributed with intensity λ, and the principal pays a flow cost c, we have the recursive relationship:

$$F^{\text{FB}} = \int_0^{\bar{X}} \lambda e^{-\lambda X}(F^{\text{FB}} - cX)dX + e^{-\lambda \bar{X}}(R - c\bar{X}), \tag{2}$$

where the integral in this expression corresponds to the possibility that a setback occurs before the project is finished, resetting progress X to 0, at which point the project must re-start. Integrating and solving yields

$$F^{\text{FB}} = R - \frac{c}{\lambda}\left(e^{\lambda \bar{X}} - 1\right). \tag{3}$$

[6] Our results hold if the principal and agent share a subjective discount rate, $r > 0$. The optimal contract is a time budget (soft deadline) with random extensions and termination, just as with $r = 0$. In fact, our economy with $r = 0$, including the principal's payoffs and policies, is attained as the limit of economies as $r \to 0$. We outline the solution to the contracting problem with discounting in Appendix E.

The first-best value is easily interpreted. Because the project is operated until it is complete, the principal eventually obtains R. Her expected cost when initiating the project is the flow cost c times the project's total expected duration $\frac{1}{\lambda}\left(e^{\lambda\bar{X}}-1\right)$. It is straightforward to verify that expected duration increases in $\bar{X}$ and λ and that $\lim_{\lambda\to 0}\frac{1}{\lambda}\left(e^{\lambda\bar{X}}-1\right)=\bar{X}$. Thus, a project with larger scope or higher frequency of setbacks has a longer expected duration, while a project for which setbacks never occur has a deterministic duration of $\bar{X}$.

It follows immediately that the first-best policy is to start the project and run it until completed if and only if the right side of (3) is positive. However, in the second-best, incentivizing the agent involves paying him rents, so a somewhat stronger assumption on the gross value of the project to the principal is required:

Assumption 2

$$R-\frac{c+b}{\lambda}\left(e^{\lambda\bar{X}}-1\right)>0. \tag{4}$$

As we show in Corollary 4 below, this condition is both necessary and sufficient for the principal to be willing to hire the agent to run the project. Interestingly, although Assumption 2 implies that the principal is willing to incur the flow cost $c+b$ until the project is eventually completed, this is not the outcome implemented by an optimally designed contract.

3 Unobservable Progress and Incentive Compatibility

3.1 Contracts and Reports

The agent's expertise and decision to work on the project allow him to privately observe its state at each instant. The principal, however, must rely on status reports from the agent. We assume:

Assumption 3 *The principal cannot observe the agent's choice of action $a_t\in\{0,1\}$, the state of the project X_t, or the occurrence of setbacks. The principal can observe project completion only upon delivery, which is contractually verifiable.*

This assumption allows the agent extensive latitude to commit malfeasance without detection. For instance, he could shirk for some time and then falsely claim a setback to cover up the lack of progress; or, following a real setback, the agent could shirk for a spell before reporting it.

However, the fact that the principal knows how long it takes to complete the project in the absence of a setback and that a completed project is verifiable does place some discipline on the agent's actions and reports. We assume

Assumption 4 *If the agent is verifiably detected misallocating resources or lying about the project's progress, then he is terminated at that point without severance.*

In other words, if the agent deviates from the principal's recommended actions, (shirks or lies), then the outcome generated must be consistent with *some* feasible path under the recommended actions. For example, the agent cannot shirk for an $\bar{X}$-interval of time without reporting a false setback or he will be fired for not delivering the completed project.

The agent makes a report of the project's current state, $\hat{X}_t$. Given the project's true evolution (1), reporting the path of $\hat{X}$ implicitly reports actions ($\hat{a}$) and setbacks ($\hat{N}$), with

$$d\hat{X}_t = \hat{a}_t(dt - \hat{X}_t d\hat{N}_t). \tag{5}$$

In fact, as long as the agent implicitly reports working, he needs only report the occurrence of setbacks with the understanding that "no news is good news" regarding progress.

The principal possesses two instruments for providing incentives so that the agent faithfully runs the project and honestly reports progress and setbacks. She can cancel the project prior to completion (i.e., *fire* the agent), or she can provide the agent with a reward when the completed project is delivered. We also allow the principal to provide the agent with rewards based on reported project status; however, we will show that because both parties are risk-neutral and are equally patient, it is without loss of generality to backload all monetary payments into a single reward granted upon successful completion.

Definition 1 (Contract) *Denote the probability space as* $(\Omega, \mathcal{F}, P)$*, and the filtration as* $\{\mathcal{F}_t\}_{t\geq 0}$ *generated by the history of reports* $\{\hat{X}_t\}_{t\geq 0}$*. Contingent on the filtration, a contract specifies a stopping time* τ *when the contract is terminated (by* *completion or cancellation), a terminal reward K_τ to the agent, and cumulative intermediate rewards $\{C\}_{t\geq 0}$. All quantities are assumed to be integrable and measurable under the usual conditions.*

Contracts are characterized using the agent's continuation utility as the state variable. Given a contract, the agent chooses actions $\{a_t\}_{t\geq 0}$ and reports $\{\hat{X}\}_{t\geq 0}$. His continuation utility is the expected value of the reward from project completion plus private benefits from any shirking:

$$W_t^{a,X} = E_{a,X}\left[\int_t^\tau b(1-a_s)ds + \int_t^\tau dC_s + K_\tau \middle| \mathcal{F}_t\right]. \tag{6}$$

The principal's objective function F_t is the expected value of the benefit from a completed project net of the expected operating cost and the expected reward to the agent:

$$F_t^{a,X} = E_{a,X}\left[-\int_t^\tau c ds + R_\tau - \int_t^\tau dC_s - K_\tau \middle| \mathcal{F}_t\right], \tag{7}$$

where $R_\tau = R$ if the project is completed and 0 if it is not.

Before we characterize general incentive compatibility, we can simplify the contracting space:

Lemma 1 (High Action and Prizes) *The principal will always choose to implement the high action ($a_t = 1$). The principal will pay the agent only upon successful completion of the project ($K_\tau > 0$ iff success; $dC_t = 0$).*

The first result holds because it is always more efficient to award the agent intermediate consumption than to implement inefficient shirking. The second result holds because both the principal and agent are equally patient and so payments can always be delayed.

A contract is incentive compatible if the agent chooses the high action and accurately reports the status of the project:

Definition 2 (Incentive Compatibility) *A contract is incentive compatible if the agent maximizes his objective (6) by choosing $a_t = 1$ and $\hat{X}_t = X_t$ for all $t \geq 0$.*

Then, in an incentive compatible contract, the agent's continuation utility is the expected value of the terminal prize. The principal's utility is the expected payoff of the project minus the running cost and expected prize.

A contract is optimal if it maximizes the principal's objective function within the class of feasible, incentive compatible contracts:

Definition 3 (Optimal Contract) *A contract is optimal if it maximizes the principal's objective function (7) over the set of contracts that 1) are incentive compatible, 2) grant the agent his initial level of utility W_0, and 3) honor limited liability, $W_t \geq 0$.*

3.2 Incentive Compatibility

In this subsection, we introduce a necessary incentive constraint, the *No-Postponed-Setbacks* (NPS) condition. This constraint provides the necessary incentives for the agent to report any setbacks immediately, rather than delaying the report and shirking in the meantime. Later, we will show that this constraint is also sufficient to prevent any other deviation.

We now summarize the evolution of the agent's continuation utility, W:

Lemma 2 (Incentive Compatibility) *Given any contract and any sequence of the agent's choices, there exists a predictable, finite, non-negative process J_t ($0 \leq t \leq \tau$) such that W_t evolves according to*

$$dW_t = J_t \left(\lambda dt - dN_t\right). \tag{8}$$

Between setbacks, J is deterministic. A necessary condition for incentive compatibility is that between setbacks, we have:

$$J_{t+\delta} \geq J_t + b\delta + \int_0^\delta \lambda J_{t+s} ds, \quad \forall \delta \in (0, \bar{X} - X). \tag{NPS}$$

The contract is terminated if $W_t = 0$.

The agent's continuation utility under an incentive compatible contract is a martingale. Hence, it drifts up deterministically at rate $\lambda J_t dt$ as the agent accumulates progress toward project completion and earning the prize, but it jumps down by J_t whenever there is a setback that wipes out the accumulated progress. To understand the (NPS) incentive constraint, suppose the state of progress is $X_t \in (0, \bar{X})$ when a setback occurs. Consider two possible paths the agent might take at this point:

- [Work] The agent reports the setback immediately, and then works as desired.

- [Shirk] The agent delays reporting the setback and shirks for time $\delta \leq \bar{X} - X_t$. Then, he reports a (bigger) setback and works as desired.

A critical feature of the shirk path is that after the postponed setback is finally reported, the agent has dissipated his persistent private information about the status of the project. The agent and the principal both believe that X_t is 0 and have the same information about the project and contract going forward. Thus, the agent's continuation utility and the principal's beliefs about it coincide.

Now, we compare the two paths, with working first. Since working is optimal, the agent's continuation utility is a martingale, and we have[7]

$$E\left[W_\tau\right] - W_{t-} = -J_t \ . \tag{9}$$

The only difference between the agent's expected utility when the project ends and his utility at $t-$ is the jump down from reporting the setback and realizing the loss in progress.

Next, we consider shirking. In this case, the change in continuation utility is

$$E\left[W_\tau\right] - W_{t-} = \int_0^\delta \lambda J_{t+s} ds - J_{t+\delta} \ , \tag{10}$$

with an additional private benefit due to shirking of $b\delta$. The first term accounts for the upward drift in the principal's beliefs about the agent's continuation utility as he (falsely) reports progress while shirking; and the second term captures the jump down in the principal's beliefs about the agent's continuation utility when he finally stops shirking and reports a larger setback than actually occurred. Adding the private benefit $b\delta$ to (10), comparing to (9), and re-arranging, we obtain (NPS). This constraint simply says that the value from the work path is at least as high as the value from the shirk path; i.e., the agent prefers to face the music immediately rather than to postpone reporting a setback.

The (NPS) constraint requires that the agent's loss of utility between setbacks is at least equal to the time he could have spent shirking between them. Thus, there is a round trip penalty imposed on the agent between any two truthfully reported setbacks. We call this a "round trip" because the agent goes from $X = 0$ through

[7] We adopt the standard convention of indexing the value of a process immediately prior to a jump with $t-$.

some path and back to $X = 0$. To see the penalty, imagine that the agent starts at time t with $X_t = 0$ and works until $t + \delta$ and $X_{t+\delta} = \delta$ when he receives a setback. With truthful reporting, (NPS) implies that the agent's continuation utility is

$$W_{t+\delta} = W_t + \int_0^\delta \lambda J_{t+s} ds - J_{t+\delta} \leq W_t - J_t - b\delta = W_t - b\delta, \tag{11}$$

where we have used in the final step that a setback at $X = 0$ has no effect and no penalty ($J_t = 0$ if $X_t = 0$). If the (NPS) constraint binds, then the agent's utility-drop between setbacks grows linearly with time: after each setback, the agent is re-started with a continuation utility that is lower by the amount of time he could have spent shirking.

3.3 Termination and Randomization

We now consider how the agent is terminated when the project is still incomplete.

First, termination is required. Imagine not; then, there is some path of X that would result in the project being funded without end. However, the agent could simply mimic that path with his reports while shirking, and thus obtain infinite utility. The agent would prefer this to any incentive compatible path.

Put differently, the NPS round trip penalty implies that between any two setbacks, the agent loses continuation value at least proportional to the elapsed time. However, termination must occur if $W_t = 0$ because, given limited liability, termination is the only way for the principal to deliver $W_t = 0$. Because the agent's initial utility W_0 is finite, the agent must also eventually run out of time.

Second, termination is random and not deterministic. We reason based on the NPS round trip utility penalty (11) and the fact that the agent has limited liability. Imagine that a setback occurs at t resulting in $W_t \in (0, b\bar{X})$. In this case, if the agent continues to work and makes progress $\delta > \frac{W_t}{b}$ before suffering another setback, then the drop in his continuation utility required by the (NPS) constraint would result in $W_{t+\delta} < 0$, which is not feasible. The agent would prefer to shirk rather than to report the second setback. What can the principal do about this? One option is simply to terminate the contract at t and give the agent a severance payment of W_t. However, allowing the agent to shirk or giving the agent a severance payment is never optimal (Lemma 1).

Instead, there is a better alternative: the principal can use randomization to either fire the agent without severance (generating $W_t = 0$) or increase W_t enough to restore incentive compatibility. Randomization preserves the agent's expected continuation utility – and thus the principal's expected payout to the agent – but (unlike paying severance) it allows for a positive probability that the project will be completed.

To preserve consistency and incentive compatibility, the agent must be randomly assigned a utility equal to 0 or greater than $b\bar{X}$. We assume now (and verify in Proposition 3) that the principal's value function is concave so that she wishes to use the least disperse randomization procedure possible.

Lemma 3 *If $W_t < b\bar{X}$ following a reported setback, then incentive compatibility and concavity of the principal's value function require that the agent is assigned utility of $b\bar{X}$ with probability $p = \frac{W_t}{b\bar{X}}$ or utility 0 with probability $1 - p$.*

4 Optimal Contract: A Time Budget

4.1 The Principal's Problem

The concept of a *time budget* plays a crucial role in the implementation of an optimal contract. A time budget is a stochastic deadline that counts down deterministically absent a setback and may jump (due to extension or termination) with the report of a setback. Formally we have the following:

Definition 4 (Time Budget) *A time budget is a non-negative process S_t satisfying $dS_t = -dt$ absent a reported setback. The principal initially grants the agent S_0 and then cancels the project iff $S_t = 0$ and the project has not succeeded. Thus, a time budget creates a random stopping time τ when the contract is terminated (on the events of project completion or cancellation).*

Our first result is that the (NPS) constraint is binding and sufficient, and thus the optimal contract can be implemented with a time budget and prize-upon-completion. In other words, the agent's loss of utility between setbacks evolves linearly with time, and this is enough to generate full effort and prevent any mis-reporting by the agent. Intuitively, the best the agent can do by lying to the principal is to gain time to divert resources from the project (shirk), and reducing the agent's expected prize by the amount he could have diverted is enough to deter such malfeasance. In turn, this means that the optimal contract can be implemented as a time budget:

Proposition 1 *The optimal contract has the following properties:*

i. The (NPS) constraint binds, and the agent's utility penalty for reporting a setback J_t is a function of X_t only:

$$J_t = J(X_t) = \frac{b}{\lambda}\left(e^{\lambda X_t} - 1\right). \tag{12}$$

ii. The contract can be implemented with a time budget which is set such that $bS_0 = W_0$ and the agent is terminated if $S_t = 0$ and the project is not delivered. If $S_{t-} < \bar{X}$ and a setback is reported, then S_t is set to either 0 with probability $1-p$ or $\bar{X}$ with probability p where

$$p = \frac{S_{t-}}{\bar{X}}. \tag{13}$$

These probabilities imply that $S_t + t$ is a martingale.

iii. The agent's continuation utility under the optimal contract is

$$W_t = bS_t + J(X_t). \tag{14}$$

If the agent completes the project at time τ, he receives a reward of

$$K_\tau = bS_\tau + \frac{b}{\lambda}\left(e^{\lambda \bar{X}} - 1\right). \tag{15}$$

This result says that it is optimal for the principal to assign the agent an amount of time S_0 to complete the project. The agent works on the project and makes continuous progress reports including the occurrence of any setbacks. If $S_t < \bar{X}$ remains on the clock and a setback is reported, then there is not enough time remaining to complete the project. At this point, the randomization procedure is invoked in which the project is either canceled with probability $1 - \frac{S_{t-}}{\bar{X}}$, or the schedule is extended to $S_t = \bar{X}$ with the complementary probability. Any subsequent setbacks are treated analogously, until the project is ultimately either canceled or completed.

The payment to the agent for delivering the completed project at time τ consists of a fixed reward $\frac{b}{\lambda}\left(e^{\lambda \bar{X}} - 1\right)$ plus a *bonus* that is proportional to the remaining time on the schedule bS_τ. The bonus term is the inverse of the NPS constraint: because the agent's utility declines between setbacks, he must receive an incentive payment if the project succeeds before he reports another setback. Note that if the agent ever receives an extension (setting the time budget back to $S_t = \bar{X}$) then the project cannot be completed early, and the agent will just receive the fixed reward.

The fixed payment, which is equal to $J(\bar{X})$, is calibrated so that the agent is indifferent between working on the project when $X = 0$ and shirking; i.e., it is the minimal reward for project completion that will induce the agent to start working with no progress in hand. As the agent works and X increases, the likelihood of project completion without a setback rises, and the agent's corresponding utility from the partially completed project, $J(X)$, accordingly grows exponentially. Hence, if a setback occurs at progress level X_t, the agent's expected utility drops by $J(X_t)$, but if he manages to push the project from 0 to $\bar{X}$ without incurring a setback, then he receives the fixed payment of $J(\bar{X})$ plus any bonus he is due for early completion.

As noted in the introduction, the implementation of the optimal incentive mechanism characterized in Proposition 1 is a cost-plus-award-fee contract. In particular, the principal commits: (i) to cover the operating cost of the project $c\tau$, (ii) to pay a fixed fee $J(\bar{X})$ upon project completion, and (iii) to pay an incentive award bS_τ for early delivery.

It is worth noting that optimal incentives can be implemented with less stringent reporting requirements. Rather than requiring continuous progress reports, at project inception the principal can announce a *soft deadline* $T = S_0 - \bar{X}$ and then commit to fund the project until this date *no-questions-asked*. If the agent delivers the completed project at $\tau \leq T$, he receives K_τ as given in the proposition. Once the soft deadline has passed, the principal requires setbacks to be reported, and she follows the random termination procedure specified in Proposition 1 from that point on.[8]

[8] The general structure outlined in Proposition 1 remains qualitatively intact under various alternative model specifications. For example, suppose for any setback that occurs, there is a probability that it is an un-fixable dead end. Then, the optimal contract is unchanged except that the agent receives a severance payment equaling bS_{t-} for the report of a fatal setback. In Sections 5 and 6 we discuss additional modifications that carry useful economic insights: lack of commitment and partial setbacks.

4.2 The Initial Value of the Project

Given the form of the optimal contract as a time budget and terminal reward, we can derive the principal's value function $F(S, X)$. We are most interested in her valuation of a given time budget when starting from scratch: $F(S, 0)$.[9] As we will see, the value function is not everywhere differentiable with respect to S; instead the value function has kinks at multiples of $\bar{X}$, and the optimal time-budget has *focal lengths* at those points. Instead of the usual techniques involving PDEs and dynamic programming, we will use two martingales and the optional stopping theorem. This allows for a more fundamental understanding of the contract and its characteristics.

In order to simplify the principal's problem, we will define two useful auxiliary functions, with slight abuses of notation.[10] The first function, $\pi(S)$, is the probability that the agent is eventually successful in completing the project. The second function, $\sigma(S)$, is the expected remaining time until contract termination (with either success or failure). Both functions assume that the principal starts with $S_t = S$ and $X_t = 0$. Since the project can be completed early, the expected time to completion is weakly less than the time budget ($\sigma(S) \leq S$). We have

$$\pi(S) = \mathrm{E}_t \left[1_{X_\tau = \bar{X}} | S_t = S, X_t = 0\right] \tag{16}$$

$$\sigma(S) = \mathrm{E}_t \left[\tau - t | S_t = S, X_t = 0\right], \tag{17}$$

where τ is the contract stopping time (Definitions 1 and 4). These two functions capture the loss to the principal from the second-best contract. In the first-best, the agent runs the project as long as necessary to complete it. In the second-best, the principal imposes a stochastic time limit (the time budget) which reduces both the probability of success and the time allowed.

We can now significantly simplify the principal's problem. From the principal's and agent's payoffs (6 and 7), and using the auxiliary functions we have just defined, we have

$$F(S, X = 0) = \pi(S)R - W(S, X = 0) - c\,\sigma(S), \tag{18}$$

[9]For $X > 0$ we have $F(S, X) = \int_0^{\bar{X}-X} \lambda e^{-\lambda t} \left(F(S-t, 0) - ct\right) dt + e^{-\lambda(\bar{X}-X)}(R - c(\bar{X} - X))$.

[10]These methods are described without discounting in the text; however the same procedure can be followed with discounting, as we outline in Appendix E.

where $W(S, X=0)$ is the agent's continuation value. From the incentive compatibility condition (14), we have $W(S, X=0) = bS$, leaving us with

$$F(S, X=0) = \pi(S)R - bS - c\,\sigma(S). \tag{19}$$

This gives us a very intuitive representation of the principal's value of the project. It is the expected reward, minus the expected agency rents granted the agent by the time budget, and minus the expected direct running cost.

We continue by applying the optional stopping theorem[11] to two salient martingales in order to relate the probability of success to the expected time remaining. First, the agent's continuation utility $W_t = bS_t + J(X_t)$ is a martingale that can only stop at two boundaries, project success and failure. Second, the randomization probabilities (13) imply that $S_t + t$ is a martingale, so the optional stopping theorem shows that $S_0 + 0 = \mathrm{E}\left[S_\tau + \tau\right]$. Thus we have

$$bS_0 = W_0 = \mathrm{E}\left[W_\tau\right] = \mathrm{E}\left[bS_\tau + J(X_\tau)\right] = bS_0 + \mathrm{E}\left[-b\tau + J(X_\tau)\right] \tag{20}$$

$$0 = -b\sigma(S_0) + \mathrm{E}\left[J(X_\tau)\right] \tag{21}$$

$$= -b\sigma(S_0) + \pi(S_0)\mathrm{E}\left[J(\bar{X})\right] + (1-\pi(S_0))\mathrm{E}\left[J(0)\right] \tag{22}$$

$$= -b\sigma(S_0) + \pi(S_0)J(\bar{X}), \tag{23}$$

where the first line follows from the optional stopping theorem, the second line follows from the definition of $\sigma(S)$, the third line from the optional stopping theorem, and the fourth line from $J(0) = 0$ (12). Re-arranging and generalizing to any time with $X_t = 0$, we have

$$\pi(S) = \frac{b\sigma(S)}{J(\bar{X})} = \frac{\lambda\,\sigma(S)}{e^{\lambda\bar{X}} - 1}. \tag{24}$$

The intuition for this result comes from the martingale property of the agent's continuation utility. The agent's utility $W_t = bS_t + J(X_t)$ counts up as progress is obtained and down as time passes. Those two changes must cancel out on average to maintain incentives. Thus, the passage of time is exactly matched by an increase in

[11] As a reminder, Doob's optional stopping theorem shows that the expectation of a martingale at a stopping time is equal to the current value of the martingale. Our setting fits the version of this result given in Theorem 5.3.1 of Cohen and Elliott (2015).

the probability that the agent receives the constant part of his reward, $J(\bar{X})$ – and the average time to completion must be proportional to the probability of success.

When S is short, the project is very likely to be canceled and both $\pi(S)$ and $\sigma(S)$ are relatively small. On the other hand, a long time budget implies a small probability of cancelation, and in the limit as $S \to \infty$, the project is never canceled so that the expected time to termination corresponds to the expected time to completion; i.e., $\lim_{S\to\infty} \pi(S) = 1$. Plugging (24) into (19), we can simplify the principal's value function further:

Proposition 2 *The principal's initial valuation of a given time budget S is*

$$F(S,0) = \left(\frac{\lambda R}{e^{\lambda \bar{X}} - 1} - c\right)\sigma(S) - bS \tag{25}$$

$$= \pi(S)\left(R - \frac{c}{\lambda}\left(e^{\lambda \bar{X}} - 1\right)\right) - bS. \tag{26}$$

This value function is concave and hump-shaped in S.

This result has an intuitive interpretation. Plugging in the first-best value (3 and 26), we can write

$$F(S,0) = \pi(S)F^{\mathrm{FB}} - bS, \tag{27}$$

Written this way, we see that asymmetric information harms the principal for two related reasons. First, $\pi(S)$ is the probability that the agent eventually delivers a completed project starting with an initial time budget of S – the probability that the project is not inefficiently canceled. The second way in which the principal is harmed is that she has to pay the expected agency rent of $W_t = bS_t$.

Intuitively, the larger the time budget S, the more likely it is that the agent will complete the project; i.e., $\pi(S)$ is increasing and $\lim_{S\to\infty} \pi(S) = 1$. But, of course the principal will not commit to pay the agent an unboundedly large rent to obtain a payoff bounded by F^{FB}. In particular, she faces a tradeoff when setting the initial time budget between higher probability of project completion, $\pi(S)$, and paying higher agency rents, bS. This tradeoff manifests in the hump shape of the value function $F(S,0)$. At low levels of S both the principal and agent prefer a larger time budget. However, as S grows, diminishing marginal returns to the probability of project completion $\pi(S)$ are eventually dominated by the linear agency cost bS, and $F(S,0)$ peaks at some critical value S^* beyond which it decreases.

4.3 The Value Function Characterization

We now characterize the expected duration of the contract $\sigma(S)$ and use that to analyze the principal's value function.

In the randomization region ($S_t \le \bar{X}$), $X_t = 0$ triggers immediate randomization and implies that the project is never completed with extra time remaining. The randomization probabilities (13) imply that the time budget is also the expected time remaining, $\sigma(S) = S$. This feature yields the following result:

Lemma 4 *For $S \le \bar{X}$, we have $\sigma(S) = S$ and*

$$F(S,0) = S\left(\tfrac{R}{J(\bar{X})} - c - b\right) = S\left(\frac{\lambda R}{e^{\lambda \bar{X}} - 1} - c - b\right). \tag{28}$$

It is also possible to solve explicitly for $\sigma(S)$ if $S > \bar{X}$ by applying an iterative procedure that is detailed in the appendix. Using (25), we can then obtain the principal's value function in closed form as a piecewise polynomial in S. The important properties of $\sigma(S)$ are summarized as follows:

Proposition 3 (Value Function Properties) *$\sigma(S)$ is continuous in S. For all $n \ge 1$, we have*

(i) For $S \in ((n-1)\bar{X}, n\bar{X}]$, $\sigma(S)$ is a concave, increasing polynomial of order n.

(ii) $\lim_{S\to\infty} \frac{\partial}{\partial S}\sigma(S) = 0$ and $\lim_{S\to\infty} \sigma(S) = \frac{1}{\lambda}\left(e^{\lambda\bar{X}} - 1\right)$.

(iii) $\lim_{S\uparrow n\bar{X}} \frac{\partial}{\partial S}\sigma(S) > \lim_{S\downarrow n\bar{X}} \frac{\partial}{\partial S}\sigma(S) > 0$.

The first observation is a straightforward implication of the iterative procedure: each round adds a higher order term but the function is always concave and increasing. Combining this observation with Lemma 4 we see that

$$\sigma(S) \begin{cases} = S, & \text{if } S \le \bar{X} \\ < S, & \text{if } S > \bar{X}. \end{cases}$$

In other words, the expected duration of the project at inception, $\sigma(S_0)$, is weakly less than the initial time budget, S_0, allotted to the agent. $\sigma(S_0) < S_0$ means that the principal builds some slack or *slippage* time into the contract: she initially gives the agent more expected time to complete the project than its actual expected duration. This implies that whenever the project is extended such that $\tau > S_0$, a schedule and cost overrun must have happened. That is, the project runs longer than its initial expected duration of $\sigma(S_0)$ and costs more than the initial estimate of $c\sigma(S_0) + bS_0$.

The second observation in Proposition 3 is a restatement of the fact that the marginal benefit from increasing S becomes arbitrarily small as $\pi(S)$ goes to 1. This along with the fact that the marginal cost remains constant at b implies that there exists a unique S^* at which $F(S, 0)$ is maximized.

The third observation identifies kinks in the value function at positive integer multiples of $\bar{X}$. One implication is that it is optimal for the principal to assign an initial time budget equal $n\bar{X}$ for a non-negligible set of parameters. In other words, the time budget has *focal lengths* that are integer multiples of the minimal time to completion. In practice, project managers often report three estimates for completion time: Best-case, worst-case, and most-likely. A particularly salient situation is when the principal holds the agent to the best-case scenario in expectation, $S^* = \bar{X}$, and responds to any reported setback by either canceling the project or resetting the clock. We analyze such a contract in the following subsection.

4.4 A Short-Leash Contract

Here, we use the first kink in the principal's value function to derive the conditions under which it is optimal for her to set an initial time budget of $S_0 = S^* = \bar{X}$. We call these *short-leash* contracts. Applying the iterative procedure in the appendix gives

$$\lim_{S \uparrow \bar{X}} \frac{\partial}{\partial S}\sigma(S) = 1 > 1 - e^{-\lambda \bar{X}} = \lim_{S \downarrow \bar{X}} \frac{\partial}{\partial S}\sigma(S). \tag{29}$$

Using this observation to evaluate the derivative of the value function in (25) directly yields the following result:

Corollary 1 *The optimal contract involves the minimal initial time budget of* $S_0 =$

$S^* = \bar{X}$ *and a fixed prize of* $\frac{b}{\lambda}\left(e^{\lambda\bar{X}} - 1\right)$ *iff*

$$\frac{c+b}{\lambda}\left(e^{\lambda\bar{X}} - 1\right) < R < \frac{c + \frac{b}{1-e^{-\lambda\bar{X}}}}{\lambda}\left(e^{\lambda\bar{X}} - 1\right). \tag{30}$$

The first inequality is a restatement of Assumption 2 (that the project is contractually feasible), and it ensures that $F(S, 0)$ is increasing for $S < \bar{X}$. The second inequality then ensures that $F(S, 0)$ is decreasing for $S > \bar{X}$. Hence, when (30) holds, the kink in the value function at $S = \bar{X}$ corresponds to the peak and it is optimal for the principal to set an initial time budget of $S^* = \bar{X}$. In other words, she should keep the agent on a short leash, granting him in expectation only the minimal amount of time necessary to complete the project, requiring him to report every setback, and canceling the project with positive probability each time one is reported.

The key parameters in (30) are b, the agent's per-period benefit from shirking, and λ, the expected frequency of setbacks. If b is too large, then the left inequality in (30) fails and moral hazard precludes the project from ever getting off the ground (i.e., Assumption 2 is violated). On the other hand, if b is too small, then the right inequality fails. In this case, moral hazard is less concerning, and the principal prefers to give the agent more than the minimal initial time to complete the project.

Intuitively, a short-leash contract is optimal if λ is sufficiently small. To see this, we calculate the limit as $\lambda \to 0$ in (30) to obtain

$$(c+b)\bar{X} < R < \infty,$$

which holds by Assumption 2. Hence, when the expected frequency of setbacks is small enough, the principal allows no slack in the schedule, committing to only the minimal rent of $b\bar{X}$ necessary to induce the agent to work on the project. This makes sense – when setbacks are relatively unlikely, the project is completed in the minimal time, $\bar{X}$, with high probability. In other words, cancellation is of little concern, and the principal prefers to promise the agent only the minimal rent necessary.

At inception of a short-leash project, the expected duration is $\sigma(S = \bar{X}) = \bar{X}$, and the expected cost to the principal is $(c+b)\bar{X}$. However, if the agent reports a setback with S_t left on the schedule, then he is granted an extension of $\bar{X} - S_{t-}$ (i.e., the clock is reset) with probability $\frac{S_{t-}}{\bar{X}}$. Hence, the support of the stopping time τ is unbounded, implying that the project may run arbitrarily long, incur arbitrarily large costs, and yet may still be canceled.

In fact, it is possible to determine explicitly the probabilities of extensions, cancellations, and overruns, and we do so for a short-leash contract. Figure 1 plots the probabilities of these events. Define $\mu \equiv \lambda \bar{X}$ to be the expected number of setbacks experienced while the project is in operation. Then,

1. $P_{\mathrm{OT}}(\mu) = e^{-\mu}$ is the probability that the project is completed on time (left panel, red dotted curve). The value of this function decreases from 1 to 0 because as the expected number of setbacks rises, the probability that none occur falls. When setbacks are a virtual certainty, the project cannot be completed on time.

2. $P_{\mathrm{EC}}(\mu)$ is the probability that the project is canceled early, before its initial expected duration of $\bar{X}$ (left panel, blue solid curve).[12] This function increases from 0 because as the expected number of setbacks rises, it becomes ever more likely that the project will not survive the requisite randomizations before time $\bar{X}$ has elapsed. Indeed, $\lim_{\mu\to\infty} P_{\mathrm{EC}}(\mu) = 1$ because a steady stream of setbacks must result in early project cancellation for any $\bar{X} > 0$.

3. $P_{\mathrm{OR}}(\mu) = 1 - P_{\mathrm{EC}}(\mu) - P_{\mathrm{OT}}(\mu)$ is the probability of an overrun, $\Pr\{\tau > \bar{X}\}$: the probability that the project ends, either from completion or cancellation, after the initial expected duration $\bar{X}$ (right panel, black solid line). It is low for small values of μ because the project will most likely be completed on time. It rises until achieving a maximum of approximately 0.39 when $\mu = 3.34$ and then decreases as the probability of early cancellation becomes ever more likely. We break this into two sub-possibilities: overruns for which the project is eventually completed P_{ORC}, and overruns followed by eventual cancellation P_{ORF}.[13]

4.5 The Value of Randomization

We can now address the value to the principal of the reporting process: why not simply assign a deadline of S_0 and stick with it? The answer is that the value of

[12] This can be obtained analytically from $P_{\mathrm{EC}}(\mu) = p(x = 1; \mu)$, where $p(x; \mu)$ is the solution to the second-order ODE $p''(x; \mu) + \mu p'(x; \mu) + \mu p(x; \mu) = \mu$ with boundary conditions $p(x = 0; \mu) = p'(x = 0; \mu) = 0$.

[13] A project that is completed is either completed on time or completed after an overrun, so we can solve from the probability of success of any kind, $\pi(\bar{X}) = \frac{\lambda \bar{X}}{e^{\lambda \bar{X}} - 1} = \frac{\mu}{e^{\mu} - 1} = P_{OT}(\mu) + P_{ORC}(\mu)$. Then, $P_{ORF} = P_{OR} - P_{ORC}$ because an overrun will result in either cancellation or completion.

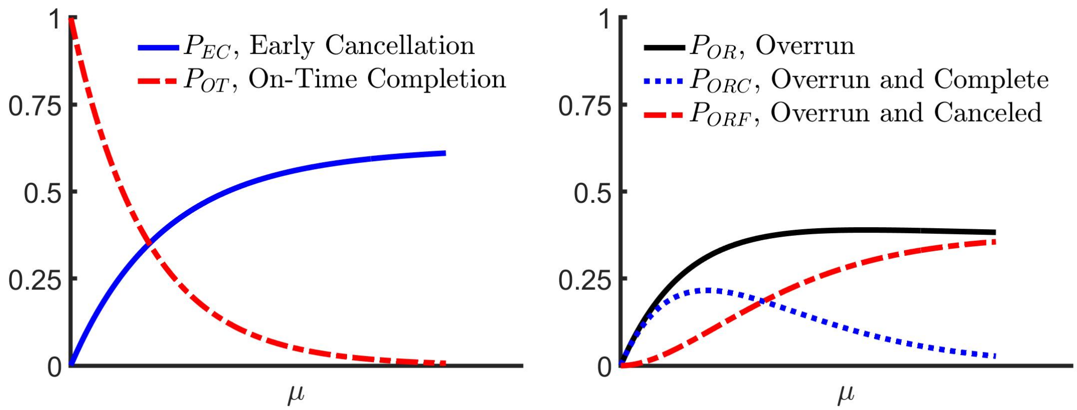

Figure 1: **Probability of an Overrun Under a Short-leash Contract**
The left panel of this figure plots the probability of early cancellation (P_{EC}, blue solid line) and the probability of on-time completion (P_{OT}, red dash-dot line). The right panel of this figure plots the probability of an overrun (P_{OR}) $\mu \equiv \lambda \bar{X}$ is the expected number of setbacks experienced while the project is in operation.

reporting derives directly from application of the randomization procedure once the contract enters the short-leash phase, $S_t \leq \bar{X}$. Assume instead that the principal does offer a fixed deadline of S_0 and the prize (15) upon completion, and then does not solicit or review reports. The agent will use high effort until experiencing a setback at $S_t < \bar{X}$ and will then shirk for the remainder of the contract.

Under the optimal contract, when $S_t \geq \bar{X}$, the time budget behaves just like a deadline, counting down naturally (i.e. $dS_t = -dt$), and the principal effectively ignores all reports. In fact, the (NPS) constraint binds, implying that the agent is always indifferent between working or shirking, and hence, willing to work for $S_t \geq \bar{X}$. However, if a setback occurs at $S_t < \bar{X}$ and the deadline is fixed, then it is impossible for the agent to complete the project and receive payment in the remaining time. Because he can only realize any positive utility through the final reward at the time of project completion if he chooses to work, he prefers to shirk out the clock and report a last-second setback to obtain bS_t. In other words, a fixed deadline and prize can be used to induce full effort until the short-leash phase of the contract. At that point, there isn't enough time left to complete the project if a setback occurs, and the agent postpones reporting the setback until the end of the contract and shirks. The reports are useful precisely because they enable the principal to create a soft deadline instead of a hard one: the principal stochastically extends the agent's time or terminates him, in place of inefficient shirking.

5 The Role of Commitment

A potential criticism of the optimal contract characterized in Proposition 1 is that it involves the principal committing to randomly cancel the project if a setback is reported when $S_t < \bar{X}$. In particular, if the agent will continue to work after receiving an extension, then the principal strictly prefers extending the project to canceling it, and – lacking commitment to randomize – she will grant an extension with probability 1. On the other hand, if the agent believes that the project will always be extended, then he will shirk and report false setbacks *ad infinitum*, leading the principal to prefer project cancelation. This is the familiar logic underpinning a mixed-strategy equilibrium. To relax the required level of commitment, we modify the baseline model in this section by assuming that randomization by either party is possible but not contractually verifiable. All other aspects of the environment remain intact.

Consider again the situation in which a setback results in $W_t \in (0, b\bar{X})$. As we noted in subsection 3.3, one way to give the agent his promised continuation utility is simply to cancel the project and make a severance payment. This option has value $-W_t$ to the principal. We showed that the principal could do better by randomizing between project cancellation without severance ($W_t = 0$) and the minimally feasible project extension ($W_t = b\bar{X}$). Indeed, because the agent continues to work if granted an extension, this randomization has strictly positive net value to the principal. However, as per the argument in the previous paragraph, this scenario can only be implemented when randomization is verifiable.

Nevertheless, it is possible to implement an outcome that delivers expected value of zero to the principal, which still dominates payment of severance. To see how, suppose that rather than working with probability 1 when granted an extension, the agent randomizes at that moment between continuing to work and shirking out the clock, two options over which he is indifferent. If the agent randomizes such that the principal's expected value from extending the project is 0, then she will be indifferent between canceling the project and extending it, and will be willing to randomize herself. In terms of payoffs, the only difference between this setting and the baseline model is that the occurrence of a setback inside the randomization region drops the principal's value function down to zero. The agent's continuation utility W_t is still a martingale and (NPS) still binds.

Proposition 4 (Non-verifiable Randomization) *If randomization is not verifi-*

able, then the optimal contract for the principal can still be implemented with a time budget and the same prize structure K_τ. If $S_{t-} < \bar{X}$ and a setback is reported, then the principal extends the schedule to $\bar{X}$ with probability $\frac{S_{t-}}{\bar{X}}$ and cancels the project with probability $1 - \frac{S_{t-}}{\bar{X}}$. Upon receiving an extension, the agent randomizes between shirking out the clock with probability $q = \frac{\hat{F}(\bar{X},0)}{\hat{F}(\bar{X},0)+c\bar{X}}$, and continuing to work with probability $1-q$, where

$$\hat{F}(S = \bar{X}, 0) \equiv Re^{-\lambda \bar{X}} - \frac{c+b}{\lambda}\left(1 - e^{-\lambda \bar{X}}\right). \tag{31}$$

Note that this version of the contract does not need to specify specific probabilities of cancellation or extension – it needs only to specify that the principal has the right to cancel or extend the project *at will*. This is somewhat more general than the contract characterized in Proposition 1 where the exact probabilities of cancellation and extension were necessarily an explicit part of the agreement.

The inability to commit to explicit probabilities harms the principal. In particular, $\hat{F}(S = \bar{X}, 0)$ is the value she derives from a short-leash contract when she cannot commit and

$$\hat{F}(S = \bar{X}, 0) = F(S = \bar{X}, 0)\left(\frac{1 - e^{-\lambda \bar{X}}}{\lambda \bar{X}}\right). \tag{32}$$

The fraction on the right side of this equation is less than 1 since $\lambda \bar{X} > 0$. A setback during the course of a short-leash contract under full commitment still leaves the principal with a positive expected payoff as shown in Corollary 4, whereas a setback in the randomization region absent commitment results in an expected payoff of zero. Although the inability to commit lowers the principal's value for the project, the above expression implies that Assumption 2 remains necessary and sufficient for feasibility of contracting when the principal cannot commit to randomize.

Note that while we specified that the agent randomizes between continuing to work and shirking out the clock after receiving an extension, there are other strategies the agent could pursue that are expected payoff equivalent for both parties. For example, while shirking the agent could mimic nature by randomly reporting setbacks at rate λ, continuing to shirk after each and every extension. Relative to shirking out the clock, this would increase the variance of τ, though not its expected value.

Perhaps surprisingly, the lack of commitment to randomize leads the principal to grant the agent *more* initial time:

Proposition 5 (Optimal Initial Time Budget) *Define S^* to be the principal's optimal initial time budget when the randomization is verifiable, and $\hat{S}^*$ to be the principal's optimal initial time budget when the randomization is not verifiable. Assume that the agent's initial outside option value is sufficiently low such that his participation constraint is met in both economies. Then, $S^* \leq \hat{S}^*$.*

Because the inability to commit to explicit randomization harms the principal (i.e. $\hat{F}(S,0) < F(S,0)$), a reasonable conjecture is that she would prefer to grant the agent less time. After all, for any given value of S, lack of commitment implies a lower probability of project completion and reduces the initial value of the project to the principal. It is, therefore, somewhat surprising that she responds by devoting more time and money to the less valuable project. In addition, the agent benefits from the principal's lack of commitment because his expected payoff is proportional to schedule length.

The intuition is actually straightforward. Lack of commitment power only harms the principal if a setback occurs in the short-leash randomization region, $S_t < \bar{X}$. By granting the agent a longer initial time budget, the principal raises the probability that the project will be completed before the lack of commitment becomes a problem. That is, she reduces the likelihood that her inability to commit will even come into play. In a sense, the principal doubles down on the part of the contract to which she can commit (the length of the schedule) in order to reduce the impact of the part to which she cannot (explicit probabilities of project cancellation and extension).

6 Partial Setbacks

Our baseline model assumes that setbacks wipe out progress completely. An important question in this regard is how general is our implementation using a time-budget and a prize-upon-completion? There are many ways to model *partial setbacks*. In this section we show that our contract form – a type of time-budget and reward – is part of any optimal contract across a wide range of setback-style models.

To accomplish this, we assume a general setback structure:

$$dX_t = a_t(dt - Y_t dN_t), \tag{33}$$

where $Y_t = Y(X_t, \theta_t) \in [0, X_t]$ is a function of the state of the project augmented by an additional random variable θ_t. With this framework, we can match several examples of partial setbacks:

- Let the project have two stages, 0 to $\tilde{X}$ and $\tilde{X}$ to $\bar{X}$. We set $Y_t = X_t$ if $X_t \leq \tilde{X}$ and $Y_t = X_t - \tilde{X}$ if $X_t \geq \tilde{X}$, meaning that setbacks wipe out progress within a stage only, but the principal cannot observe the stage or progress within stage.
- Let the project have setbacks of random size, so that $Y_t = \theta_t X_t$, where θ_t are i.i.d with $\theta \sim G(\theta)$.
- Let setbacks have a maximum size, so that $Y_t = \max(X_t, \bar{Y})$. Thus, setbacks can wipe out the early development of a project (the phase with $X_t \leq \bar{Y}$) but not the entire project if near completion.

We assume that the project is always worth continuing (analogously to Assumption 2), and that the principal can voluntarily re-set progress to zero[14] Then, the arguments in Lemma 1 remain valid in this setting:

Lemma 5 (High Action and Prizes) *The principal will always choose to implement the high action ($a_t = 1$). The principal will pay the agent only upon successful completion of the project ($K_\tau > 0$ iff success; $dC_t = 0$).*

This lemma is crucial because it means that the principal continues to use only two incentive devices in the more general setting: termination and a prize upon completion. Then, to show that something like a time budget is necessary, we need only to specify the termination time under an incentive compatible contract. To deliver utility of W_0 in an incentive compatible way *requires* terminating the agent after enough time has gone by. The intuition for this result is simple: if there were a path of the project that would enable the agent to continue operating for long enough, then the agent would simply report that path while shirking and gain more utility than the optimal contract promises.

To see the result, fix a contract and define τ^R to be the stopping time associated with a particular sequence of reports: the time the agent either successfully reports completion of the project or is terminated. Since the agent can only report paths

[14]Setting $X_t = 0$ is value destroying, but it is a convenient tool to allow the principal to re-start the contract in the proof of Lemma 5.

that are possible, the space of possible paths and the space of possible reports are the same. For any such contract that delivers utility W_0 to the agent, we must have $\mathrm{E}[\tau^R] \leq W_0/b$, otherwise the agent could make this sequence of reports while shirking to gain utility greater than W_0. Thus, while τ^R is not bounded, its expectation is bounded for incentive compatible contracts, and the agent either completes the project or is fired in finite time with probability 1.

Then, define

$$\sigma_t = \mathrm{E}_t[\tau] - t, \tag{34}$$

where τ is the stopping time associated with truthful reporting, which makes σ_t analogous to $\sigma(S)$ (17). Because $\mathrm{E}_t[\tau]$ is a conditional expectation and hence a martingale, $\mathrm{E}[d\sigma_t] = -dt$. More, σ_t is equal to zero whenever the project is completed or terminated. To proceed, we define the notion of an *expected* time budget:

Definition 5 (Expected Time Budget) *An expected time budget is a non-negative process σ_t satisfying $E[d\sigma_t] = -dt$. The principal initially grants the agent σ_0 and then, if the project does not succeed, cancels the project iff $\sigma_t = 0$. An expected time budget recovers the random stopping time τ from termination of the contract (on the events of project completion or cancellation).*

Now, the principal's problem can be simplified under partial setbacks in a manor analogous to that in the baseline section. Denote the probability of success by π_t:

$$\pi_t = \mathrm{E}_t\left[1_{X_\tau = \bar{X}}\right]. \tag{35}$$

Then, from the principal's and agent's payoffs (6 and 7), we have that the principal's value can be written

$$F_t = \pi_t R - W_t - c\sigma_t, \tag{36}$$

which is analogous to (18). This is sufficient to show:

Proposition 6 *Assume that setbacks follow the structure in (33), and that the contract offered is incentive compatible and delivers the agent an initial utility of W_0. Then, the contract terminates the agent when an expected time budget hits zero, unless the project is complete, in which case the agent is awarded a prize which equals* *his terminal utility. The principal's payoff is given by (36).*

This result generalizes the notion of time-budgets to models with partial setbacks: any optimal contract can be implemented with an expected time budget and a (path dependent) prize-upon-completion. The expected time budget captures termination, and the prize captures the agent's terminal utility, and those are the only two incentive devices that the principal uses (Lemma 5). We have not assumed that an optimal contract exists, and a full characterization of an optimal contract would require fixing a specific model of partial setbacks from the class described by (33). Instead, Proposition 6 tells us that regardless of the particular selection we fix, an optimal contract is always implemented with an expected time budget (soft deadline) and a reward for successful completion, just as in our baseline model.

7 Conclusion

At a very general level, projects are usually viewed as possessing three defining features, scope, schedule, and budget – the so-called "iron triangle." (Wyngaard et al., 2012) The scope of a project is the quality of the deliverable, be it a software application, a power plant, or a doctoral thesis; The schedule is the time allotted to production of the deliverable; and the budget is the monetary or other physical resources committed to it. However, because projects are, by definition, at least somewhat unique, their implementation typically involves considerable uncertainty. In this paper we held scope fixed, and presented a model of project implementation focusing on what appears to be the most common sources of project uncertainty, schedule setbacks and the concomitant cost overruns.

Whether a project is under taken in-house (e.g., the Boeing Dreamliner) or outsourced (e.g., South Carolina's V.C. Summer nuclear plant or the FBI's virtual-case-file system), its progress will almost surely be hampered to some degree by agency frictions. To study this, we embedded a natural model of production with random setbacks into a dynamic agency environment and solved for the optimal contract from the principal's perspective. This analysis yielded a number of novel insights and conclusions. Among the most robust are: 1) an optimal contract can always be implemented with a time budget and a terminal payment corresponding to a cost-plus-award-fee contract; 2) penalties for reports of setbacks or delays are generally more severe the later they occur in project development; and 3) mishaps that are reported near the end of the allotted schedule either result in project cancelation or minimally feasible project extension.

There are numerous avenues available for future research. For example, expanding our current treatment to incorporate common strategies for dealing with the time pressure created by unanticipated setbacks seems promising. The completion of projects is frequently time-sensitive as noted by Lewis and Bajari (2011) who investigate the procurement of highway construction projects where completion delays can have large social costs. In this vein, exploring the possibility of speeding up production through *fast-tracking* (running several phases in parallel) or *crashing* (deploying more resources) to make up for unanticipated delays is a potentially important consideration. Finally, there is the question of scope itself. Throughout we supposed that the project was either incomplete (worth zero to the sponsor) or complete (worth a fixed amount). In reality, the ultimate quality of many projects varies along a continuum. Indeed, *scope creep* on the part of sponsors (demanding a higher quality deliverable than originally specified) is often cited as a contributing factor to project failure.[15] We leave these considerations and others for future work, having judged this particular project to be deliverable as complete.

[15] Ely and Szydlowski (2020) considers a dynamic moral hazard problem in which the principal uses scope creep to entice the agent to exert effort on a project that he would not have agreed to work on if he had known the full scope of the project at the outset.

Appendix – Proofs and Derivations

A Proof of Lemmas 1, 2, and 3

A.1 Proof of Lemma 1

First, the principal always induces the high action ($a_t = 1$). Imagine there is an interval of time in which the principal induces shirking. The project does not advance, nor is there a setback. The principal can award the agent intermediate consumption without paying the flow cost of the project during the shirking interval without changing the agent's continuation utility. This makes the agent indifferent and the principal better off, because $c > b$ implies that assigning the agent any positive amount of utility by allowing shirking is more costly for the principal than directly paying the agent.

Second, any contract with intermediate payment can be weakly improved by one without. Because the principal and the agent share the same discount rate (0), the principal can simply delay any intermediate payments until the end, leaving both participants indifferent.

Third, any contract with severance pay upon termination can be improved by one that pays only on the event of success. Notice that the principal can re-start any existing incentive compatible contract and both participants will have positive value going forward. Any contract that ends with a severance payment can be replaced with one that randomizes between re-starting the contract and termination with zero payment. The probability of re-starting the project can be set to make the agent indifferent to the randomization. The principal is better off because she receives zero (termination) or a positive value (re-start) instead of making a severance payment.

A.2 Proof of Lemma 2

First, by Lemma 1, $dC_t = 0$ for all $t < \tau$. Therefore (6) becomes

$$W_t^{a,X} = E_{a,X}\left[\int_0^\tau b(1-a_s)ds + K_\tau \,|\mathcal{F}_t\right],$$

where $\mathcal{F}_t$ is the filtration generated by the agent's report $\{\hat{X}_t\}_{t\in[0,\tau]}$. Note that W_t is an $\mathcal{F}_t$-martingale. Thus, by the martingale representation theorem for jump processes, there exists a $\mathcal{F}_t$-predictable, integrable process J such that

$$dW_t^{a,X} = J_t(\lambda dt - d\hat{N}_t). \tag{37}$$

If the contract is incentive compatible, the agent will exert high effort and report the setback N_t truthfully. In this case, the analysis following the statement of Lemma 2 applies, and the (NPS) condition

$$-J_t \geq b\delta + \int_0^\delta \lambda J_{t+s} ds - J_{t+\delta}, \quad \forall \delta \in (0, \bar{X} - X) \tag{38}$$

must hold as a necessary condition. Finally , J_t must be weakly positive, otherwise the agent would gain from falsely reporting a setback. $\square$

A.3 Proof of Lemma 3

Suppose $W_t < b\bar{X}$ following a reported setback. Then randomization between $W_{t'} = 0$ and $W_{T'} \geq b\bar{X}$ is necessary. To see why, suppose that immediately after randomization (including degenerate randomization) or at the beginning of a phase of Poisson termination, the agent's continuation utility is $W_{t'} \in (0, b\bar{X})$. Imagine the agent continues to work for $\delta \in (W_{t'}/b, \bar{X})$ and then suffers a setback which he truthfully reports. By the round-trip property of (NPS) we have

$$W_{t'+\delta} = W_{t'} + \int_0^\delta \lambda J_s ds - J_\delta \tag{39}$$

$$\leq W_{t'} - J_0 - b\delta = W_{t'} - b\delta < 0, \tag{40}$$

which violates limited liability. Finally, concavity of the principal's value function (established in Proposition 3) implies that randomization should involve minimal dispersion. Hence, to deliver W_t to the agent, he should receive $W_{t'} = 0$ with probability $(1 - p)$ and $W_{t'} = b\bar{X}$ with probability $p = \frac{W_t}{b\bar{X}}$. $\square$

B Proof of Proposition 1

The proof of Proposition 1 consists of four parts. Part 1 demonstrates that the (NPS) constraint can be re-written in a more tractable way. Part 2 shows that (NPS) binds with equality in an optimal contract. Part 3 shows that if (NPS) holds with equality, then the optimal contract can be implemented with a simple time budget. Part 4 shows that a simple time budget is sufficient to make truthful reporting optimal for the agent. We will also take as given that the principal's value function is concave with respect to W (or, S), which is verified in Proposition 2 and the analysis in Section 4.2.

We begin with the following lemma which is helpful several times:

Lemma 6 (Lower bound for the marginal value of agent's utility) *Let $F(W,X)$ be the principal's value function (7), then $F_W(W,X) \geq -1$.*

Proof: Imagine not: $F_W(W,X) < -1$. Then the principal would gain by giving the agent intermediate consumption. But this cannot be the case (Lemma 1).□

B.1 Re-writing (NPS)

An important fact is that J can only vary with time (or, the progress of X) between t and $t+\delta$. Because no setbacks are being reported, the passage of time is the only thing the principal can observe. Thus, we can write J as a function of current project progress and all of history prior to the previous setback, $J(X,\cdot)$. We will suppress the $(\cdot)$ notation for convenience. Then, (NPS) becomes

$$J(X_t) \leq -b\delta - \int_0^\delta \lambda J(X_t+s)ds + J(X_t+\delta),\ \forall X_t \in [0,\bar{X}) \text{ and } \delta > 0. \tag{41}$$

First, we reformulate (NPS) so that J can be written as the sum of its minimum, binding value and a term capturing the excess. If (NPS) binds everywhere, (41) holds with equality and we can take the derivative with respect to δ to obtain

$$0 = -b - \lambda J(X+\delta) + J'(X+\delta). \tag{42}$$

This has the (general) solution $J(X) = C_1 e^{\lambda X} - \frac{b}{\lambda}$, where C_1 is a constant. Because $J(X) \geq 0$, the minimum value of C_1 is $\frac{\beta}{\lambda}$. Thus, the minimum, binding value of $J(X)$ is

$$J^{min}(X) = \frac{b}{\lambda}e^{\lambda X} - \frac{b}{\lambda}. \tag{43}$$

Next, we define the functions $f(X)$ and $g(X)$ such that

$$f(X) = \frac{1}{b}\left[J(X) - J^{min}(X)\right] \tag{44}$$

$$g(X) = f(X) - \int_0^X \lambda f(u)du. \tag{45}$$

So, $f \geq 0$ captures the difference between J and its minimum, and g is a convenient summarizing function. Note that $f(0) = g(0)$, so $f(X)$ can be fully recovered from $g(X)$.

Second, we show that (NPS) is satisfied if and only if $g(X) \geq 0$ and f and g are weakly increasing. Substituting the definition of $f(X)$ into (41) yields

$$f(X) \leq -\int_0^\delta \lambda f(X+u)du + f(X+\delta),\ \forall X_t \in [0,\bar{X}) \text{ and } \delta > 0. \tag{46}$$

Then, from the definition of g, we have

$$\begin{aligned} g(X+\delta) - g(X) &= f(X+\delta) - \int_0^{X+\delta} \lambda f(u)du - f(X) + \int_0^X \lambda f(u)du && (47)\\ &= f(X+\delta) - \int_X^{X+\delta} \lambda f(u)du - f(X) && (48)\\ &\geq 0, && (49)\end{aligned}$$

where (46) shows that (48) is non-negative, so $g(X)$ is weakly increasing. More, $f(0) = g(0)$ and g weakly increasing imply $g(X) \geq 0$. Then, using the definition of g and $f \geq 0$, we see that we must also have f weakly increasing. Using the definitions of f and g, we see that f and g positive and weakly increasing are also sufficient to show (41). Thus, we can use f and g instead of (46) to characterize (NPS).

We can now write the change in the agent's continuation utility as X starts at $X_t = 0$ and progresses until a setback is experienced at $X = X_{t+s}$. Using $X_{t+u} = u$ and the f and g notation, we have

$$\begin{aligned} \int_0^s \lambda J(X_{t+u})du - J(X_{t+s}) &= \int_0^s be^{\lambda u}du - b\bar{X} - \frac{b}{\lambda}e^{\lambda s} + \frac{b}{\lambda} - bg(X_{t+s})\\ &= \frac{b}{\lambda}\left(e^{\lambda s} - 1\right) - bs - \frac{b}{\lambda}e^{\lambda s} + \frac{b}{\lambda} - bg(X_{t+s})\\ &= -bs - bg(X_{t+s}). && (50)\end{aligned}$$

Similarly, we can write down the prize the agent receives, given that the agent starts at time t with W_t and $X_t = 0$ and progresses to project completion:

$$W_t + \int_0^{\bar{X}} \lambda J(x)dx = W_t + \int_0^{\bar{X}} be^{\lambda x}dx - b\bar{X} + b\int_0^{\bar{X}} \lambda f(x)dx. \tag{51}$$

B.2 The Optimality of Binding the (NPS)

We show that the optimal contract has $f = g = 0$. To proceed, using (50 and 51) and that the probability that any particular try at the project is successful is $e^{-\lambda\bar{X}}$, we have that

$F(W, X=0)$ can be written as

$$F(W, X=0) = e^{-\lambda\bar{X}} \left[R - \frac{b}{\lambda}\left(e^{\lambda\bar{X}} - 1\right) + (b-c)\bar{X} - W - b\int_0^{\bar{X}} \lambda f(u)du \right] \tag{52}$$

$$+ \int_0^{\bar{X}} \lambda e^{-\lambda u} F(W - bu - bg(u), X=0)du, \tag{53}$$

with constraints $g'(u) \geq 0$ and $g(u) \geq 0$. We will use the Hamiltonian maximization method with $\zeta(u) = g'(u)$ as the control variable, $f(u)$ and $g(u)$ as the state variables, and $f'(u) = \lambda f(u) + \zeta(u)$ and $g'(u) = \zeta(u)$ as the laws of motion. The constraints are $\zeta(u) \geq 0$ and $g(u) \geq 0$. The objective function, ignoring constant terms, is

$$\max \int_0^{\bar{X}} \left[\lambda e^{-\lambda u} F\left(W - bu - bg(u), X=0\right) - b\lambda e^{-\lambda\bar{X}} f(u) \right] du. \tag{54}$$

Then, the Hamiltonian is

$$\mathcal{H} = e^{-\lambda u}\lambda F\left(W - bu - bg, X=0\right) - b\lambda e^{-\lambda\bar{X}} f \tag{55}$$
$$+ \gamma_1(\lambda f + \zeta) + \gamma_2\zeta + \eta_1(\zeta - 0) + \eta_2(g - 0).$$

The optimality conditions are

$$0 = \frac{\partial\mathcal{H}}{\partial\zeta} = \gamma_1 + \gamma_2 + \eta_1 \tag{56}$$

$$-\gamma_1' = \frac{\partial\mathcal{H}}{\partial f} = -b\lambda e^{-\lambda\bar{X}} + \gamma_1\lambda \tag{57}$$

$$-\gamma_2' = \frac{\partial\mathcal{H}}{\partial g} = -e^{-\lambda X} b\lambda F_W\left(W - bu - bg, X=0\right) + \eta_2. \tag{58}$$

We can solve for $\gamma_1(u)$ directly:

$$\gamma_1(u) = k_1 e^{-\lambda u} + b e^{-\lambda\bar{X}}, \tag{59}$$

for some constant k_1.

We now work through the various cases with respect to η_1 and η_2:

- Imagine that neither constraint binds for some X, so $\eta_1 = \eta_2 = 0$. Then, $\gamma_2 = -\gamma_1 = -k_1 e^{-\lambda u} + b e^{-\lambda\bar{X}}$ and so $-\gamma_2' = -\lambda k_1 e^{-\lambda u}$. Plugging that back in (58), we obtain $bF_W\left(W - bu - bg, X=0\right) = k_1$ which implies $u + g$ is a constant, which violates g weakly increasing.

- Imagine that $g' \geq 0$ binds but $g \geq 0$ does not, so $\eta_1 > 0$ and $\eta_2 = 0$. Then, $-\gamma_2' = \gamma_1' + \eta_1' = -\lambda k_1 e^{-\lambda u} + \eta_1'$. This is a valid differential equation solution.
- Imagine that $g' \geq 0$ does not bind but $g \geq 0$ does, so $\eta_2 > 0$ and $\eta_1 = 0$. This is the solution in which $g(u) = 0$ and is valid.

Thus, we need only to consider the case in which $g(u) = g(0)$ is constant and optimize over that constant.

To proceed, we differentiate the definition of g and solve for f to obtain

$$f(u) = e^{\lambda u}\left[g(0) + \int_0^u e^{-\lambda v} g'(v) dv\right]. \tag{60}$$

If g is constant, we have $f(u) = e^{\lambda u} g(0)$. Using the definition of $g(u)$ to show that $\int_0^u \lambda f(v) dv = g(\bar{X}) - f(\bar{X})$, the objective function for the principal can be written as

$$\max \int_0^{\bar{X}} \left[\lambda e^{-\lambda u} F\left(W - bu - bg(0), X = 0\right) du\right] + b(g(\bar{X}) - f(\bar{X})) \tag{61}$$

$$= \max \int_0^{\bar{X}} \left[\lambda e^{-\lambda u} F\left(W - bu - bg(0), X = 0\right) du\right] + bg(0)(1 - e^{\lambda \bar{X}}). \tag{62}$$

Given the concavity of F in W, the first-order condition for $g(0)$ is

$$-\int_0^{\bar{X}} \left[\lambda e^{-\lambda u} b F_W\left(W - bu - bg(0), X = 0\right) du\right] - b(e^{\lambda \bar{X}} - 1). \tag{63}$$

By Lemma 6, $F_W(\cdot, X = 0) > -1$, so this expression is negative and the optimal value of $g(0)$ is 0. Since g is constant, we also have $f = g = 0$, so

$$J(X) = J^{min} = \frac{b}{\lambda}\left(e^{\lambda X} - 1\right). \tag{64}$$

B.3 Implementation

The optimal contract can be implemented with a time budget with the following properties:

1. $S_0 = W_0/b$.
2. $dS_t = -dt - M_t dN_t$. $M_t = 0$ if $S_t \geq \bar{X}$. If $S_t < \bar{X}$, M_t is a binary random variable that takes value $\bar{X} - S_t$ with probability $p = S_t/\bar{X}$ and $-S_t$ with probability $1 - p$.
3. The contract ends at $t = \tau$ if $S_\tau = 0$.

Part I and II of the proof shows that if a setback is reported at any time t, the agent's continuation utility is $W_t = W_0 - bt$. Therefore, $W_t = b(S_0 - t) = bS_t$, where the first equality comes from Properties 1 above and the second equality utilizes Property 2, respectively. Randomization occurs when a setback is reported and $W_t < b\bar{X}$ where W_t jumps to $b\bar{X}$ with probability $p = W_t/b\bar{X}$ or 0 with probability $1-p$. Because $W_t = bS_t$ following any setback, randomization occurs if $S_t < \bar{X}$, and S_t jumps to $\bar{X}$ with probability $p = bS_{t-}/b\bar{X} = S_{t-}/\bar{X}$ or 0 with probability $1-p$ (Property 2 and 3). Finally, combining (50) and (51) and using the fact that $f = g = 0$, the prize for project completion is

$$K_\tau = W_0 - b\tau + \frac{b}{\lambda}\left(e^{\lambda\bar{X}} - 1\right) \tag{65}$$

$$= bS_\tau + \frac{b}{\lambda}\left(e^{\lambda\bar{X}} - 1\right). \tag{66}$$

B.4 Truthful Reporting

Next we show that the linear time-budget is sufficient to induce the agent to report truthfully and take the high action. The agent's objective is to solve (6), subject to the evolution of state variables:

$$dX_t = a_t(dt - X_t dN_t) \tag{67}$$

$$dS_t = -dt - M_t d\hat{N}_t, \tag{68}$$

where

$$M_t = \begin{cases} 0 & \text{if } S_{t-} \geq \bar{X} \\ \bar{X} - S_{t-} & \text{if } S_{t-} < \bar{X} \text{ , with probability } \frac{S_{t-}}{\bar{X}} \\ -S_{t-} & \text{if } S_{t-} < \bar{X} \text{ , with probability } 1 - \frac{S_{t-}}{\bar{X}}. \end{cases} \tag{69}$$

Let $U(X,S)$ denote the agent's value function. Then $U(X,S)$ satisfies

$$0 = \max\Big\{ \mathrm{E}U(X, S-M) - U(X,S), \\ U_X a + (1-a)b - U_S \\ + \lambda \max\left\{\mathrm{E}U(0, S-M) - U(X,S), U(0,S) - U(X,S)\right\} \Big\} \tag{70}$$

with the following boundary conditions,

$$U(\bar{X}, S) = bS + \frac{b}{\lambda}\left(e^{\lambda \bar{X}} - 1\right) \tag{71}$$

$$U(X, 0) = 0. \tag{72}$$

The first line of (70) represents the change of utility from reporting a false setback or not. The second line represents the flow utility from working or shirking, and the third line represents the change of utility from postponing the report of a true setback or not. The two boundary conditions come from project completion and termination, respectively, both of which are verifiable events for the principal.

The idea behind (70) is that the agent is making a branched choice. The first max is over whether to announce a false setback or not to do that and instead to see what happens over dt. If there is no false setback, the agent chooses to work or shirk, and whether to postpone the report of a setback when one occurs.

Outside the randomization region, $M = 0$ which implies $U(X, S - M) - U(X, S) = 0$. That is, the agent does not benefit from falsely reporting a setback. Furthermore, when $M = 0$, (70) implies

$$0 \geq \max_a \left[U_X a + (1-a)b - U_S + \lambda \max\left(U(0, S) - U(X, S)\right)\right], \tag{73}$$

where the equality is achieved when $a = 1$ and

$$U(X, S) = bS + \frac{b}{\lambda}\left(e^{\lambda X} - 1\right). \tag{74}$$

Moreover, (74) satisfies the boundary conditions (71) and (72). Therefore, the agent achieves the highest utility from exerting the high effort and reporting actual setbacks immediately outside the randomization region.

Inside the randomization region, M_t equals $\bar{X} - S_{t-}$ with probability $S_{t-}/\bar{X}$ and equals $-S_{t-}$ with probability $1 - S_{t-}/\bar{X}$. Then equation (70) becomes

$$\begin{aligned} 0 = \max \Bigg\{ & \left(1 - \frac{S}{\bar{X}}\right) U(X, 0) + \frac{S}{\bar{X}} U(X, \bar{X}) - U(X, S), \\ & U_X a + (1-a)b - U_S \\ & + \lambda \max\left\{\left(1 - \frac{S}{\bar{X}}\right) U(0, 0) + \frac{S}{\bar{X}} U(0, \bar{X}) - U(X, S), U(0, S) - U(X, S)\right\}\Bigg\}. \end{aligned} \tag{75}$$

With boundary conditions $U(X, 0) = 0$ and $U(X, \bar{X}) = b\bar{X} + \frac{b}{\lambda}\left(e^{\lambda X} - 1\right)$ (from (74) at

$S = \bar{X}$). Replacing $U(X,0)$ and $U(X,\bar{X})$ in (75) with these boundary conditions and substituting in $U(X,S) = bS + \frac{b}{\lambda}\left(e^{\lambda X} - 1\right)$ yields

$$0 = \max\left\{ \left(\frac{S}{\bar{X}} - 1\right)\frac{b}{\lambda}\left(e^{\lambda X} - 1\right), b\left(e^{\lambda X} - 1\right)(a-1)\right\}, \tag{76}$$

where we have used the fact that $S/\bar{X} < 1$ inside the randomization region. Equation (76) implies that the agent prefers working if $X > 0$ and is indifferent if $X = 0$. The agent strictly prefers not announcing a false setback, and is indifferent between postponing the report of a setback or not when one actually occurs. $\square$

C Proof of Proposition 3

Recall the definition of $\sigma(S)$ from (17). From the randomization for $S \in [0, \bar{X}]$, we have $\sigma(S) = \frac{S}{\bar{X}}\sigma(\bar{X})$ for $S \in [0, \bar{X}]$.

We will use an iterative procedure to find $\sigma(S) = \sigma_n(S)$ on interval $S \in (\bar{X} + n\bar{X}, \bar{X} + (n+1)\bar{X}]$ for $n \geq 0$. We start by observing that for any $S \geq \bar{X}$,

$$\sigma(S) = \int_0^{\bar{X}} \lambda e^{-\lambda t}(t + \sigma(S-t))dt + e^{-\lambda\bar{X}}\bar{X}. \tag{77}$$

Further define

$$\phi_n(\nu) = \int_{\nu - n\bar{X}}^{\bar{X}} \lambda e^{-\lambda t}\sigma_{n-1}(\nu - t)dt \tag{78}$$

$$\xi = \int_0^{\bar{X}} \lambda e^{-\lambda t}tdt + e^{-\lambda\bar{X}}\bar{X} = \frac{1}{\lambda}\left(1 - e^{-\lambda\bar{X}}\right). \tag{79}$$

To continue,

$$\begin{aligned}
\sigma(S) &= \int_0^{S-n\bar{X}} \lambda e^{-\lambda t}(t + \sigma(S-t))dt + \int_{S-n\bar{X}}^{\bar{X}} \lambda e^{-\lambda t}(t + \sigma(S-t))dt + e^{-\lambda\bar{X}}\bar{X} \\
&= \int_0^{S-n\bar{X}} \lambda e^{-\lambda t}\sigma(S-t)dt + \int_0^{S-n\bar{X}} \lambda e^{-\lambda t}tdt + \int_{S-n\bar{X}}^{\bar{X}} \lambda e^{-\lambda t}(t + \sigma(S-t))dt + e^{-\lambda\bar{X}}\bar{X} \\
&= \int_0^{S-n\bar{X}} \lambda e^{-\lambda t}\sigma(S-t)dt + \int_0^{\bar{X}} \lambda e^{-\lambda t}tdt + \int_{S-n\bar{X}}^{\bar{X}} \lambda e^{-\lambda t}\sigma(S-t)dt + e^{-\lambda\bar{X}}\bar{X} \\
&= \int_0^{S-n\bar{X}} \lambda e^{-\lambda t}\sigma(S-t)dt + \phi_n(S) + \xi.
\end{aligned} \tag{80}$$

Differentiating with respect to S and then integrating by parts yields

$$\sigma'(S) = \lambda\phi_n(S) + \phi_n'(S) + \lambda\xi. \tag{81}$$

So,

$$\sigma(S) = \sigma(n\bar{X}) + \int_{n\bar{X}}^{S} \sigma'(\nu)d\nu = \sigma(n\bar{X}) + \int_{n\bar{X}}^{S} (\lambda\phi_n(\nu) + \phi_n'(\nu) + \lambda\xi)d\nu, \tag{82}$$

where $\nu - t \leq n\bar{X}$ and $\sigma(\nu), \nu \leq \bar{X}$ is known from iteration in the previous round.

C.1 Part (i)

First, $\sigma(S)$ is increasing. For $S \leq \bar{X}$, we have $\sigma(S) = \frac{\sigma_1(\bar{X})}{\bar{X}}S$, and hence $\sigma'(S) > 0$. For $S > \bar{X}$,

$$\sigma(S) = \int_0^{\bar{X}} \lambda e^{-\lambda t}(t + \sigma(S - t))dt + e^{-\lambda\bar{X}}\bar{X}, \tag{83}$$

and taking the derivative yields

$$\sigma'(S) = \int_0^{\bar{X}} \lambda e^{-\lambda t}\sigma'(S - t)dt. \tag{84}$$

Hence $\sigma'(S) > 0$ because $\sigma'(S - t) < 0$ for $\forall 0 < t < S$.

Second, $\sigma(S)$ is concave for $S \in (n\bar{X}, (n+1)\bar{X}]$. When $n = 1$, we have $\sigma_1(S) = \frac{\sigma_1(\bar{X})}{\bar{X}}S$, which is weakly concave in S. Then, suppose concavity holds for $n - 1$, we show that it holds for n as well. Since $\sigma_n'(S) = \lambda\phi_n(S) + \phi_n'(S) + \lambda\xi$, differentiating yields $\sigma_n''(S) = \lambda\phi_n'(S) + \phi_n''(S)$. Differentiating ϕ yields

$$\phi_n'(S) = -\lambda e^{-\lambda(S-n\bar{X})}\sigma_{n-1}(n\bar{X}) + \int_{S-n\bar{X}}^{\bar{X}} \lambda e^{-\lambda t}\sigma_{n-1}'(S - t)dt \tag{85}$$

$$\phi_n''(S) = \lambda^2 e^{-\lambda(S-n\bar{X})}\sigma_{n-1}(n\bar{X}) - \lambda e^{-\lambda(S-n\bar{X})}\sigma_{n-1}'(n\bar{X}) \tag{86}$$

$$- \lambda e^{-\lambda(S-n\bar{X})}\sigma_{n-1}'(n\bar{X}) + \int_{S-n\bar{X}}^{\bar{X}} \lambda e^{-\lambda t}\sigma_{n-1}''(S - t)dt \tag{87}$$

and

$$\begin{aligned}
\sigma_n''(S) &= \lambda\phi_n'(S) + \phi_n''(S) \\
&= \lambda \int_{S-n\bar{X}}^{\bar{X}} \lambda e^{-\lambda t}\sigma_{n-1}'(S-t)dt - 2\lambda e^{-\lambda(S-n\bar{X})}\sigma_{n-1}'(n\bar{X}) + \lambda e^{-\lambda t}\sigma_{n-1}''(S-t)dt \\
&\leq \lambda \int_{S-n\bar{X}}^{\bar{X}} \lambda e^{-\lambda t}\sigma_{n-1}'\left(S-(S-n\bar{X})\right)dt - 2\lambda e^{-\lambda(S-n\bar{X})}\sigma_{n-1}'(n\bar{X}) + 0 \\
&= -\lambda e^{-\lambda\bar{X}}\sigma_{n-1}'(n\bar{X}) - \lambda e^{-\lambda(S-n\bar{X})}\sigma_{n-1}'(n\bar{X}) < 0. \qquad (88)
\end{aligned}$$

The last step is from the fact that σ_{n-1} is a increasing function. So, by induction we have $\sigma_n(S)$ is concave on interval $S \in (n\bar{X}, (n+1)\bar{X}]$ for every n.

Third, we show that $\sigma_n(S)$ is a polynomial of order n. We have

$$\begin{aligned}
\lambda\phi_n(S) + \phi_n'(S) &= \lambda \int_{S-n\bar{X}}^{\bar{X}} \lambda e^{-\lambda t}\sigma_{n-1}(S-t)dt - \lambda e^{-\lambda(S-n\bar{X})}\sigma_{n-1}(n\bar{X}) \\
&\quad + \int_{S-n\bar{X}}^{\bar{X}} \lambda e^{-\lambda t}\sigma_{n-1}'(S-t)dt \\
&= \lambda \int_{S-n\bar{X}}^{\bar{X}} \lambda e^{-\lambda t}\sigma_{n-1}(S-t)dt - \lambda e^{-\lambda(S-n\bar{X})}\sigma_{n-1}(n\bar{X}) \\
&\quad + \lambda e^{-\lambda(S-n\bar{X})}\sigma_{n-1}(n\bar{X}) - \lambda e^{-\lambda\bar{X}}\sigma_{n-1}(S-\bar{X}) \\
&\quad - \lambda \int_{S-n\bar{X}}^{\bar{X}} \lambda e^{-\lambda t}\sigma_{n-1}(S-t)dt \\
&= -\lambda e^{-\lambda\bar{X}}\sigma_{n-1}(S-\bar{X}), \qquad (89)
\end{aligned}$$

and therefore (82) implies

$$\sigma_n(S) = \sigma(n\bar{X}) + \int_{n\bar{X}}^{S} \left(\lambda\xi - \lambda e^{-\lambda\bar{X}}\sigma_{n-1}(\nu-\bar{X})\right)d\nu. \qquad (90)$$

Since σ_1 is linear and hence polynomial of order 1, by (90) σ_2 is quadratic and σ_n is polynomial of order n.

C.2 Parts (ii) and (iii)

Because

$$\sigma(S) = \frac{1}{\lambda}\left[1 - \pi(S)e^{\lambda\bar{X}} - (1-\pi(S))\right] < \infty, \qquad (91)$$

$\sigma(S)$ is a bounded function. Since every monotone and bounded function in $\mathbb{R}$ converges, let the limit be

$$\sigma(\infty) = \lim_{S\to\infty} \sigma(S). \tag{92}$$

Since

$$\sigma(S) = \int_0^{\bar{X}} \lambda e^{-\lambda t}(t + \sigma(S-t))dt + e^{-\lambda \bar{X}}\bar{X}, \tag{93}$$

taking $S \to \infty$ yields

$$\sigma(\infty) = \int_0^{\bar{X}} \lambda e^{-\lambda t}(t + \sigma(\infty))dt + e^{-\lambda \bar{X}}\bar{X} = \frac{1}{\lambda}(1 - e^{-\lambda \bar{X}}) + \sigma(\infty)(1 - e^{-\lambda \bar{X}}), \tag{94}$$

and hence we have

$$\sigma(\infty) = \frac{e^{\lambda \bar{X}} - 1}{\lambda}. \tag{95}$$

Moreover, by the convergence of $\sigma(S)$, we have

$$\lim_{S\to\infty} \sigma'(S) = \lim_{S\to\infty} \lim_{\Delta S \to 0} \frac{\sigma(S+\Delta S) - \sigma(S)}{\Delta S} = 0. \tag{96}$$

Then, at $S = n\bar{X}$, the left slope is

$$\begin{aligned}
\sigma'_{n-1}(n\bar{X}) &= \lambda \phi_{n-1}(n\bar{X}) + \phi'_{n-1}(n\bar{X}) + \lambda\xi \\
&= \lambda \int_{\bar{X}}^{\bar{X}} \lambda e^{-\lambda t}\sigma_{n-1}(n\bar{X} - t)dt - \lambda e^{-\lambda \bar{X}}\sigma_{n-1}(n\bar{X}) \\
&\quad + \int_{\bar{X}}^{\bar{X}} \lambda e^{-\lambda t}\sigma'_{n-1}(n\bar{X} - t)dt + \lambda\xi \\
&= -\lambda e^{-\lambda \bar{X}}\sigma_{n-1}(n\bar{X}) + \lambda\xi.
\end{aligned} \tag{97}$$

The right slope is

$$\begin{aligned}
\sigma'_n(n\bar{X}) &= \lambda \phi_n(n\bar{X}) + \phi'_n(n\bar{X}) + \lambda\xi \\
&= \lambda \int_0^{\bar{X}} \lambda e^{-\lambda t}\sigma_{n-1}(n\bar{X} - t)dt - \lambda\sigma_{n-1}(n\bar{X}) + \lambda\xi.
\end{aligned} \tag{98}$$

Hence

$$\begin{aligned}\sigma'_{n-1}(n\bar{X}) - \sigma'_n(n\bar{X}) &= -\lambda e^{-\lambda\bar{X}}\sigma_{n-1}(n\bar{X}) - \lambda\int_0^{\bar{X}} \lambda e^{-\lambda t}\sigma_{n-1}(n\bar{X} - t)dt + \lambda\sigma_{n-1}(n\bar{X}) \\ &> -\lambda e^{-\lambda\bar{X}}\sigma_{n-1}(n\bar{X}) - \lambda\int_0^{\bar{X}} \lambda e^{-\lambda t}\sigma_{n-1}(n\bar{X})dt + \lambda\sigma_{n-1}(n\bar{X}) \\ &= -\lambda e^{-\lambda\bar{X}}\sigma_{n-1}(n\bar{X}) - \lambda\sigma_{n-1}(n\bar{X}) + \lambda\sigma_{n-1}(n\bar{X}) - \lambda\sigma_{n-1}(n\bar{X}) \\ &= 0,\end{aligned}$$

where the inequality comes from the fact that σ_{n-1} is a (strictly) increasing function. $\square$

D Proof of Proposition 4 and Proposition 5

D.1 Value Function Without Commitment to Randomization

The same steps used to prove Proposition 1 apply in this setting, except that upon receiving an extension the agent must randomize between working and shirking such that the principal's expected payoff from granting the extension is the same as from canceling the project, namely 0.

First observe that the agent's expected utility from working immediately following an extension is $b\bar{X}$, which is the same as his utility from shirking out the clock – so he is willing to randomize at this point. Now, suppose the agent receives an extension and works. Let $s \in [0, \bar{X}]$ be the amount of time since the extension was granted. Any setback during the extension will reset the principal's expected utility to 0. Therefore her expected utility when the agent works as a function of s is

$$\big(\mathcal{R} - c(\bar{X} - s)\big)\, e^{-\lambda(\bar{X}-s)} - \int_0^{\bar{X}-s} ct\lambda e^{-\lambda t}dt \tag{99}$$

$$= \mathcal{R}e^{-\lambda(\bar{X}-s)} - \frac{c}{\lambda}\left(1 - e^{-\lambda(\bar{X}-s)}\right) \equiv \hat{F}(\bar{X} - s, s), \tag{100}$$

where $\mathcal{R} \equiv R - \frac{b}{\lambda}\left(e^{\lambda\bar{X}} - 1\right)$ is the principal's surplus at the completion of the project.

Let q be the probability the agent shirks out the clock following an extension. Then randomization between cancelation and granting an extension is incentive compatible for the principal iff

$$(1 - q)\hat{F}(\bar{X}, 0) - qc\bar{X} = 0. \tag{101}$$

Solving for q yields the formula given in the claim.

There is one further item that must be checked: the principal must be willing to continue the contract if the agent does not report a setback. This must be checked because the principal does not know if the agent is working or shirking out the clock and shirking does not generate setbacks. The probability that the principal believes the agent is shirking conditional on having reported no setbacks by s is

$$q(s) = \frac{\hat{F}(\bar{X}, 0)}{\hat{F}(\bar{X}, 0) + c\bar{X}e^{-\lambda s}}.$$

Then, the principal is willing to let the clock run during the extension iff for all $s \in [0, \bar{X}]$,

$$\begin{aligned}
&(1 - q(s))\hat{F}(\bar{X} - s, s) - q(s)c(\bar{X} - s) \geq 0 \\
\iff &\bar{X}e^{-\lambda s}\hat{F}(\bar{X} - s, s) \geq (\bar{X} - s)\hat{F}(\bar{X}, 0) \\
\iff &\bar{X}e^{-\lambda\bar{X}}\left(\mathcal{R}e^{-\lambda(\bar{X}-s)} - \frac{c}{\lambda}\left(1 - e^{-\lambda(\bar{X}-s)}\right)\right) \geq (\bar{X} - s)\left(\mathcal{R}e^{-\lambda\bar{X}} - \frac{c}{\lambda}\left(1 - e^{-\lambda\bar{X}}\right)\right) \\
\iff &\frac{c}{\lambda}\left(\bar{X} - s + se^{-\lambda\bar{X}} - \bar{X}e^{-\lambda s}\right) \geq -s\mathcal{R}e^{-\lambda\bar{X}}.
\end{aligned}$$

The right side is non-positive, so the claim will hold if the left side is non-negative for $s \in [0, \bar{X}]$:

$$\bar{X} - s + se^{-\lambda\bar{X}} - \bar{X}e^{-\lambda s} \geq 0.$$

At $s \in \{0, \bar{X}\}$ the inequality above evidently binds. For $s \in (0, \bar{X})$, differentiating the left side w.r.t. s yields $-1 + e^{-\lambda\bar{X}} + \lambda\bar{X}e^{-\lambda s}$. This is positive at $s = 0$, negative at $s = \bar{X}$, and 0 at a single point between 0 and $\bar{X}$.$\square$

D.2 Optimal Initial Time Budget

Define $\hat{\pi}(S_0)$ to be the probability that the project succeeds without commitment given an initial time budget S_0 and initial progress $X_0 = 0$. Then, the principal's value function can be written

$$\hat{F}(S_0, X = 0) = \hat{\pi}(S_0)E\left[R - bS_{\hat{\tau}^{proj}} - \frac{b}{\lambda}(e^{\lambda\bar{X}} - 1)\middle| S_0, X_{\hat{\tau}^{proj}} = \bar{X}\right] - cE[\hat{\tau}^{proj}|S_0]$$

The first term is the expected reward minus payment, the second term is expected running cost.

Define $\hat{\tau}^{work}$ as the random time the agent stops working (either by completion of project, or by the agent shirking) and $\hat{\tau}^{proj}$ as the random time the project stops (either by completion of project, or by termination of project). Both $\hat{\tau}^{work}$ and $\hat{\tau}^{proj}$ are $\mathcal{F}_t^N-$ stopping times, and $\hat{\tau}^{work} \leq \hat{\tau}^{proj}$.

Next, we continue with the martingale analysis, following the main text. For $t \leq \hat{\tau}^{work}$, we have that $dX_t = dt - X_t dN_t$ implies $d(e^{\lambda X_t} - \lambda t) = (1 - \lambda e^{\lambda X_t})(dN_t - \lambda dt)$, which is a martingale. Applying the optional stopping theorem regarding $e^{\lambda X_t} - \lambda t$, we have

$$\begin{aligned}
1 &= \mathrm{E}[e^{\lambda X_t} - \lambda t | S_0, t = \hat{\tau}^{work}] \\
&= \hat{\pi}(S_0)\mathrm{E}[e^{\lambda X_{\hat{\tau}^{work}}} | S_0, X_{\hat{\tau}^{work}} = \bar{X}] + (1 - \hat{\pi}(S_0))\mathrm{E}[e^{\lambda X_{\hat{\tau}^{work}}} | S_0, X_{\hat{\tau}^{work}} = 0] - \lambda \mathrm{E}[\hat{\tau}^{work} | S_0] \\
&= \hat{\pi}(S_0)e^{\lambda \bar{X}} + (1 - \hat{\pi}(S_0)) - \lambda E[\hat{\tau}^{work} | S_0] \qquad (102)
\end{aligned}$$

Notice that we have used the fact that success or failure of the project is fully determined at $t = \hat{\tau}^{work}$; one does not have to wait until $t = \hat{\tau}^{proj}$. Solving, we have

$$\hat{\pi}(S_0) = \frac{\lambda}{e^{\lambda \bar{X}} - 1}\mathrm{E}[\hat{\tau}^{work} | S_0] \qquad (103)$$

Next, apply the optional stopping theorem regarding $S_t + t$

$$\begin{aligned}
S_0 &= \mathrm{E}[S_t + t | S_0, t = \hat{\tau}^{proj}] \\
&= \hat{\pi}(S_0)\mathrm{E}[S_{\hat{\tau}^{proj}} | S_0, X_{\hat{\tau}^{proj}} = \bar{X}] + (1 - \hat{\pi}(S_0))\mathrm{E}[S_{\hat{\tau}^{proj}} | S_0, X_{\hat{\tau}^{proj}} = 0] + \mathrm{E}[\hat{\tau}^{proj} | S_0] \\
&= \hat{\pi}(S_0)\mathrm{E}[S_{\hat{\tau}^{proj}} | S_0, X_{\hat{\tau}^{proj}} = \bar{X}] + \mathrm{E}[\hat{\tau}^{proj} | S_0] \qquad (104)
\end{aligned}$$

The third line follows from that $\mathrm{E}[S_{\hat{\tau}^{proj}} | S_0, X_{\hat{\tau}^{proj}} = 0] = 0$. If the project has ended at $X_t = 0$, there must not be any time remaining.

Examining the terms in the principal's value function one at a time, we have

$$\hat{\pi}(S_0)R = \frac{R\lambda}{e^{\lambda \bar{X}} - 1}\mathrm{E}[\hat{\tau}^{work} | S_0] \qquad (105)$$

$$-\hat{\pi}(S_0)E[bS_{\hat{\tau}^{proj}} | S, X_{\hat{\tau}^{proj}} = \bar{X}] = b\mathrm{E}[\hat{\tau}^{proj} | S_0] - bS_0 \qquad (106)$$

$$-\hat{\pi}(S_0)\frac{b}{\lambda}(e^{\lambda \bar{X}} - 1) = -b\mathrm{E}[\hat{\tau}^{work} | S_0] \qquad (107)$$

$$-cE[\hat{\tau}^{proj} | S_0] = -cE[\hat{\tau}^{proj} | S_0] \qquad (108)$$

Adding these up and re-arranging, we obtain

$$\hat{F}(S_0, X = 0) = R\frac{\lambda \mathrm{E}[\hat{\tau}^{work} | S_0]}{e^{\lambda \bar{X}} - 1} - cE[\hat{\tau}^{proj} | S_0] - bS_0 + b\left(\mathrm{E}[\hat{\tau}^{proj} | S_0] - \mathrm{E}[\hat{\tau}^{work} | S_0]\right)$$

which implies

$$\hat{F}(S_0, X=0) = \left(\frac{R\lambda}{e^{\lambda\bar{X}}-1} - c\right) \mathrm{E}[\hat{\tau}^{proj}|S_0] - bS_0 + \left(\frac{R\lambda}{e^{\lambda\bar{X}}-1} + b\right) \left(\mathrm{E}[\hat{\tau}^{work}|S_0] - \mathrm{E}[\hat{\tau}^{proj}|S_0]\right) \tag{109}$$

Next, we observe that upon entering the short-leash region, we have $E[\hat{\tau}^{proj}|S = \bar{X}, X = 0] = E[\tau|S = \bar{X}, X = 0] = \bar{X}$. This means that commitment does not change the average project duration. In addition, outside the short-leash region, the agent is not shirking, and the evolution of X and S are the same with and without commitment. Thus, we have $E[\hat{\tau}^{proj}|S_0] = E[\tau|S_0]$ (where τ is the with-commitment stopping time), and we can write the principal's value function without commitment as

$$\hat{F}(S_0, X=0) = F(S_0, X=0) - \left(\frac{R\lambda}{e^{\lambda\bar{X}}-1} + b\right) \left(\mathrm{E}[\hat{\tau}^{proj}|S_0] - \mathrm{E}[\hat{\tau}^{work}|S_0]\right) \tag{110}$$

As a reminder, F is the principal's value function with commitment. We want to show that $\left(\mathrm{E}[\hat{\tau}^{proj}|S_0] - \mathrm{E}[\hat{\tau}^{work}|S_0]\right)$ is decreasing in S_0.

We will use the fact that outside of the short-leash region, the agent never shirks. Thus, if the project succeeds before entering the short-leash region, we have $\hat{\tau}^{work} = \hat{\tau}^{proj}$. Then, we can write the difference between $E[\hat{\tau}^{proj}|S_0]$ and $E[\hat{\tau}^{work}|S_0]$ entirely as a function of the probability of entering the short-leash region and the expected time spent there:

$$E[\hat{\tau}^{proj}|S_0] - E[\hat{\tau}^{work}|S_0] = (1 - \Pr(X_{\hat{\tau}^{proj}} = \bar{X}, S_{\hat{\tau}^{proj}} > \bar{X})) \times \left(\mathrm{E}[\hat{\tau}^{proj}|S_{\hat{\tau}^{proj}} \leq \bar{X}] - \mathrm{E}[\hat{\tau}^{work}|S_{\hat{\tau}^{work}} \leq \bar{X}]\right) \tag{111}$$

where $(1-\Pr(X_{\hat{\tau}^{proj}} = \bar{X}, S_{\hat{\tau}^{proj}} > \bar{X})) = (1-\Pr(X_{\hat{\tau}^{work}} = \bar{X}, S_{\hat{\tau}^{work}} > \bar{X}))$ is the probability of entering the short leash region. Then we notice: $\mathrm{E}[\hat{\tau}^{proj}|S_{\hat{\tau}^{proj}} \leq \bar{X}] - \mathrm{E}[\hat{\tau}^{work}|S_{\hat{\tau}^{work}} \leq \bar{X}]$, the difference in the expected time taken by the end of work and the end of the project in the short leash region, does not depend on the starting value of S_0.

Since the probability that the agent does not succeed before the short-leash region is decreasing in S_0, we must have $E[\hat{\tau}^{proj}|S_0] - \mathrm{E}[\hat{\tau}^{work}|S_0]$ is decreasing in S_0. Thus, from (110), the principal's value function without commitment is equal to the value function with commitment minus a term that is decreasing in S_0; without-commitment and with-commitment have increasing differences in S_0. Both value functions are bounded from above and are $-\infty$ for $S_0 \to \infty$. By the standard logic of increasing differences, the optimal S_0 without commitment must be weakly larger than the optimal S_0 with commitment. $\square$

E Optimal Contract with Discounting

This section presents the key results of the baseline model when both contracting parties have the same discount rate $r > 0$. Details of the derivations are mostly omitted in the interest of space, but are available upon request.

First, suppose the project is operated until completed. Compared to (3), F^{FB}, the first-best value to the principal at inception, is now given by

$$F^{FB} = R\left(\frac{e^{-(r+\lambda)\bar{X}}}{\lambda e^{-(r+\lambda)\bar{X}} + r}\right)(\lambda + r) - c\left(\frac{1 - e^{-(r+\lambda)\bar{X}}}{\lambda e^{-(r+\lambda)\bar{X}} + r}\right) \tag{112}$$

In the baseline model, F^{FB} is given by the final reward R minus the expected cost of operating the project. The same interpretation applies here, except that both costs and rewards are discounted. Taking the $r \to 0$ limit, wee see that (112) goes to (3).

We first consider the case that the agent's initial continuation utility (W_0) is less than b/r. In this case, the contract must terminate, otherwise the agent would be better off obtaining the perpetuity value of shirking b/r by adopting the report strategy that yields no termination. Conditional on the project being feasible, Lemma 1 still applies: high action is always optimal and no payment except for project completion. The No-Postponed-Setback (NPS) constraint can be written as

$$W_t - J_t \geq \int_0^\delta e^{-rt} b dt + e^{-r\delta}\left[W_{t+\delta} - J_{t+\delta}\right] \tag{113}$$

The law of motion for the agent's continuation utility W_t is

$$dW_t = rW_t dt + J_t(\lambda dt - dN_t) \tag{114}$$

Like in the baseline model without discounting, the agent's value drifts up over time in the absence of a setback. However, the drift is composed of two parts: the interest over time (captured by rW_t) and the reward for having no setback (captured by λJ_t).

We can show that binding NPS is still optimal. This, combined with (113) and (114), implies the utility penalty for reporting a setback is

$$J(X) = \frac{b}{\lambda + r}\left[e^{(\lambda+r)X} - 1\right] \tag{115}$$

Clearly, (115) becomes (12) when $r \to 0$.[16]

[16]Without discounting, $J(X)$ also equals b times the project's expected duration. However, this is a

The optimal contract has the same structure as characterized in Proposition 1 in the baseline model: a time budget S with a soft deadline that involves randomization. Under this contract, the agent's utility when $X = 0$ (either at the outset of the contract or immediately after experiencing a setback) is

$$W(S, X = 0) = \frac{b}{r}\left(1 - e^{-rS}\right) \tag{116}$$

Randomization occurs when $S < \bar{X}$ and a setback is reported. Upon randomization, either $S = 0$, in which case the contract is terminated with no payment, or S is extended to $\bar{X}$. The probability of the extension is,

$$p(S) = \frac{1 - e^{-rS}}{1 - e^{-r\bar{X}}} \tag{117}$$

Finally, the agent is rewarded with a prize K for completion, where

$$K(S) = \frac{b}{r}\left(\frac{\lambda}{r+\lambda} - e^{-rS}\right) + \frac{b}{r+\lambda}e^{(r+\lambda)\bar{X}} \tag{118}$$

Unlike the baseline model, the probability of extension and the prize for completion with discounting are no longer linear functions of the remaining balance of the time budget. However, they are both still increasing functions in S, meaning that later reports of the setback result in a lower probability of extension, and later completion of the project results in a smaller prize. Moreover, straightforward algebra shows that they both converge to their baseline model counterparts; i.e., (117) becomes (13), and (118) becomes (15), when $r \to 0$.

As in the baseline model, the principal's initial value function at the outset of the contract $F(S)$, has kinks at the integer multiples of $\bar{X}$. We begin by defining the auxiliary functions,

$$\pi(S) = \mathrm{E}_t\left[e^{-r(\tau-t)}1_{X_\tau=\bar{X}}|X_t = 0\right] \tag{119}$$

$$\sigma(S) = \mathrm{E}_t\left[\int_t^\tau e^{-r(s-t)}ds|X_t = 0\right] \tag{120}$$

which are the discounted probability of success and expected duration of the project, analogous to (16) and (17). This allows us to write the principal's value function as

$$F(S) = \pi(S)R - \frac{b}{r}\left(1 - e^{-rS}\right) - c\sigma(S) \tag{121}$$

which is analogous to (19).

coincidence and does not hold with $r > 0$.

We proceed using the same martingale methods as in the baseline model. We observe that $e^{-rt}W_t$ and $1 - e^{-r(S_t+t)}$ are martingales, and thus we can repeat the steps beginning with (20) to obtain an analogy to (24):

$$\pi(S) = \frac{(r+\lambda)\sigma(S)}{e^{(r+\lambda)\bar{X}} - 1}. \tag{122}$$

The intuition is the same as in the baseline model: the agent's continuation utility has two terms, one that counts down as the time budget is used and one that counts up as the project accumulates progress. These two changes must cancel out on average to maintain incentives. Thus the passage of time is matched by an increase in the probability of eventual success. Then, combining (121) and (122), we obtain an analogy to (25):

$$F(S, X=0) = \left(\frac{(r+\lambda)}{e^{(r+\lambda)\bar{X}} - 1} R - c\right) \sigma(S) - \frac{b}{r}\left(1 - e^{-rS}\right). \tag{123}$$

In the "short-leash" region ($S \leq \bar{X}$), every setback triggers randomization, and the probability of retention is given by (117) above. Then, direct calculation shows that $\sigma(S) = \frac{1}{r}\left(1 - e^{-rS}\right)$, so (123) becomes

$$F(S \leq \bar{X}, X=0) = \frac{1 - e^{-rS}}{r}\left(\frac{(r+\lambda)}{e^{(r+\lambda)\bar{X}} - 1} R - (b+c)\right) \tag{124}$$

$F(S)$ in any region $[n\bar{X}, (n+1)\bar{X}]$ (where $n \in \mathbb{N}$) can be obtained from the same iterative procedure used to derive $F(S)$ in the baseline model in Propositions 2 and 3, albeit with more cumbersome algebra. One can show through direct calculation that the limit as $r \to 0$ of $\sigma(S)$ in (120) converges to the value with $r = 0$ in (77). Then, using (123), one can show that the principal's value $F(S)$ for $r > 0$ converges to that in the baseline model (i.e. equation 26) when $r \to 0$ for all $S > 0$.

The above analysis requires $W_0 < \frac{b}{r}$. If $W_0 \geq \frac{b}{r}$ and $F^{\text{FB}} - W_0 > 0$, then the optimal contract is to offer a fixed prize of $K = W_0\left(\frac{\lambda + re^{(\lambda+r)\bar{X}}}{\lambda+r}\right)$ upon completion, request no intermediate reports, and never fire the agent. This contract implements the first best policy, but is not feasible for small r, e.g., $r < b/R$.

References

Bergemann, D. and U. Hege (1998). Venture capital financing, moral hazard, and learning. *Journal of Banking and Finance 22*(6-8), 703–735.

Biais, B., T. Mariotti, G. Plantin, and J.-C. Rochet (2007). Dynamic Security Design: Convergence to Continuous Time and Asset Pricing Implications. *Review of Economic Studies 74*, 345–390.

Biais, B., T. Mariotti, J.-C. Rochet, and S. Villeneuve (2010). Large risks, limited liability, and dynamic moral hazard. *Econometrica 78*(1), 73–118.

Billante, P. (2017). 3 Ways to Avoid Common Construction Schedule Setbacks. *BIM 360*. August 8.

Brandt, W. (2020). Berlin Airport Opens 10 Years late and Three Times over Budget. *Euronews*. February 11.

Clementi, G. L. and H. a. Hopenhayn (2006). A Theory of Financing Constraints and Firm Dynamics. *Quarterly Journal of Economics 121*, 229–265.

Cohen, S. and R. Elliott (2015). *Stochastic Calculus and Applications*. Birkhäuser, New York, NY.

Curello, G. and L. Sinander (2021). Screening for breakthroughs. Working paper.

DeMarzo, P. M. and Y. Sannikov (2006). Optimal security design and dynamic capital structure in a continuous-time agency model. *Journal of Finance 61*, 2681–2724.

Ely, J. C. and M. Szydlowski (2020). Moving the Goalposts. *Journal of Political Economy 128*(2), 468–506.

Goldstein, H. (2005). Who Killed the Virtual Case File? How the FBI blew more than $100 million on case-management software it will never use. *IEEE Spectrum*. September 1.

Green, B. and C. R. Taylor (2016). Breakthroughs, deadlines, and self-reported progress: Contracting for multistage projects. *American Economic Review 106*(12), 3660–99.

Greene, J. (2019). The top 9 reasons for IT project failure: Is your project at risk? *atSpoke*.

Haynes, W. (2007). Boston's big dig project: A cautionary tale. *Bridgewater Review 27*(1), 3–7.

Hoffmann, F., R. Inderst, and M. M. Opp (2020). Only Time Will Tell: A Theory of Deferred Compensation. *Review of Economic Studies*. Forthcoming.

Irigoyen, C. (2017). The FBI Virtual Case File System. *CPI Case Study*. June 20.

Lacy, A. (2019). South Carolina Spent $9 Billion To Dig A Hole In The Ground And Then Fill It Back In. *The Intercept*. February 6.

Lewis, G. and P. Bajari (2011). Procurement contracting with time incentives: Theory and evidence. *Quarterly Journal of Economics 126*(3), 1173–1211.

Lewis, T. R. (2012). A theory of delegated search for the best alternative. *RAND Journal of Economics 43*(3), 391–416.

Lineberger, R. S. and A. Hussain (2016). Program management in Aerospace & Defense: Sector programs are still over budget and late. Deloitte Development LLC, Hermitage, TN, USA.

Madsen, E. (2021). Designing deadlines. *American Economic Review*. Forthcoming.

Mason, R. and J. Välimäki (2015). Getting it done: dynamic incentives to complete a project. *Journal of the European Economic Association 13*(1), 62–97.

Mayer, S. (2021). Financing breakthroughs under failure risk. *Journal of Financial Economics*. Forthcoming.

Phelan, C. and R. M. Townsend (1991). Computing multi-period, information-constrained optima. *Review of Economic Studies 58*(5), 853–881.

Quadrini, V. (2004). Investment and liquidation in renegotiation-proof contracts with moral hazard. *Journal of Monetary Economics 51*, 713–751.

Rahmani, M., G. Roels, and U. S. Karmarkar (2017). Collaborative work dynamics in projects with co-production. *Production and Operations Management 26*(4), 686–703.

Sannikov, Y. (2008). A continuous-time version of the principal-agent problem. *Review of Economic Studies 75*, 957–984.

Shenhar, A., V. Holzmann, B. Melamed, and Y. Zhao (2016). The challenge of innovation in highly complex projects: what can we learn from boeing's dreamliner experience? *Project Management Journal 47*(2), 62–78.

Spear, S. E. and S. Srivastava (1987). On repeated moral hazard with discounting. *Review of Economic Studies 54*(4), 599–617.

Toxvaerd, F. (2006). Time of the essence. *Journal of Economic Theory 129*(1), 252–272.

Vasama, S. (2017). On moral hazard and persistent private information. Working paper.

Wild, C. (2015). Building the Sydney Opera House: An architectural icon, 10 years late and 1,457% over budget. *Mashable*. November 7.

Williams, N. (2011). Persistent Private Information. *Econometrica 79*(4), 1233–1275.

Wyngaard, C. J. V., J. Pretorius, and L. Pretorius (2012). Theory of the Triple Constraint – a Conceptual Review. *IEEM Procceedings*. December 10.