Setbacks, Shutdowns, and Overruns*

Felix Feng Curtis R. Taylor Mark Westerfield Feifan Zhang

February 2024

Abstract

We investigate optimal project management in a setting plagued by an indefinite number of setbacks that are discovered en route to project completion. The contractor can cover up delays in progress due to shirking either by making false claims of setbacks or by postponing the reports of real ones. The sponsor optimally induces work and honest reporting via a soft deadline and a reward for completion that specifies a bonus for early delivery. Late-stage setbacks trigger randomization between minimally feasible project extension and (inefficient) cancellation. Because extensions may be granted repeatedly, arbitrarily large overruns in schedule and budget are possible after which the project may still be canceled.

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) is from Duke University and Zhang (feifan.zhang@duke.edu) is from Duke Kunshan University. We are grateful for helpful comments and discussions 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, Adriano Rampini, Luis Rayo, Takuo Sugaya, and Mike Waldman. We likewise thank seminar participants at: Cornell's JGSM-SBE, Hong Kong UST, the 2021 SFS Cavalcade, the NBER Organizational Economics Working Group (spring 2021), the Econometric Society's 2021 NASM, the Finance Theory Group Fall 2021 meeting, the 2021 Webinar Series on Credence Goods and Expert Markets , The 2022 Canadian Economic Theory Conference, the SITE 2022 session on dynamic games and contracts, Indiana University's Kelley School of Business (BEPP), and the University of Colorado's Economics and Finance Departments. 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)

1 Introduction

Broad swaths of the modern economy are dedicated to the 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 prevalence of projects, it is important to understand the intrinsic characteristics of this mode of production 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. A prime example is South Carolina's V.C. Summer nuclear power plant construction project, canceled in 2017 after a series of setbacks generated substantial overruns in schedule and budget. Palmetto-State taxpayers were ultimately saddled with a bill of $9 billion and "nothing to show for it" (Lacy, 2019). Similarly, 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 for sharing files among agents. Irigoyen (2017) summarizes the VCF project 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."

Of course, South Carolina and the FBI are not alone in there project-management woes. According to Lineberger and Hussain (2016), "The combined cost overrun for Major Defense Acquisition programs in 2015 was $468 billion ... with an average schedule delay of 29.5 months." Similarly, according to a 2017 report from the Project Management Institute, "14 percent of IT projects fail... [and] Of the projects that

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

didn't fail outright, 31 percent didn't meet their goals, 43 percent exceeded their initial budgets, and 49 percent were late" (Greene, 2019).

In this paper we argue that schedule and cost overruns and project cancellations are not always the product of incompetence or inattention, but – at least to some degree – are unavoidable consequences of optimal project management 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: differing site conditions (construction), a design feature doesn't work as intended (manufacturing), or incompatibility of certain off-the-shelf subroutines (software engineering).

In our model, as in practice, the sponsor (the principal) must hire a contractor (the agent) to run the project. Both parties are risk-neutral, but the agent is protected by limited liability. Setbacks occur 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 occurrence of a setback. The first-best policy is simply to start and run the project until completed if and only if the value of the finished project exceeds the flow cost times expected project duration.

The agency 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 and must rely on 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. Indeed, 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 by submitting false reports of setbacks or postponing the reports of real ones, or any combination of these. 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 turns out to be the *No-Postponed-Setbacks* (NPS) condition. This 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 two important consequences. First, it implies that the agent also prefers not to cover up resource diversion with claims of false setbacks; that is, binding (NPS) is both necessary and sufficient for incentive compatibility. Second, it allows us to fully characterize the optimal contract, which closely resembles what the U.S. General Service Administration defines as a *cost-plus-award-fee contract* (see FAR 16.401). We discuss this more formally in Remark 2 in Subsection 4.4.

Under the optimal contract, the principal announces a *time budget* S_0, which is a type of soft deadline, and commits to pay the running cost until the project is completed or canceled. The time budget counts down deterministically with calendar time unless and until a reported setback makes project completion in the remaining time infeasible; i.e., at $S_{t-} < \bar{X}$. In this case, a binary randomization procedure is implemented under which the project is either canceled $S_t = 0$ without payment to the agent, or a minimally feasible schedule extension, $S_t = \bar{X}$ is granted. Subsequent reports of setbacks are treated similarly. Hence, the contract ultimately ends for one of two reasons, either because the project is delivered or because it is canceled. Upon delivery of the completed project, the agent is paid a linear reward consisting of a fixed fee plus an incentive payment proportional to any value remaining on the time budget. If one or more schedule extensions results in the project running longer than S_0, then an *overrun* (in time and operating cost) is said to occur. Given that the first-best policy is to run the project until it is completed, overruns are not in themselves problematic. However, a substantial overrun that ultimately ends in project cancellation obviously involves a large waste of resources.

We characterize the principal's value function by identifying two martingales and invoking the optional stopping theorem. This reveals that the principal's expected payoff equals the probability of 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 $S_0 = S^*$.

The probability of eventual success following a setback obeys a delayed differential equation (DDE) that results in a focal-point *kink* in the principal's value function at $S = \bar{X}$. This implies existence of a generic set of parameter values for which $S^* = \bar{X}$ is optimal. We call this a *short-leash* contract because the soft deadline equals the minimal time for project completion and every reported setback results in cancellation with positive probability. Although a short-leash contract has expected duration $\bar{X}$, the support of the stopping time is $(0, \infty)$ due to the possibilities of early cancellations and multiple extensions. Hence, even when the principal commits to keep the agent on a short leash, arbitrarily large cost and schedule overruns may occur. Importantly, every optimal contract possesses a short-leash phase that is triggered whenever the agent is granted an extension.

Our two main methodological contributions are the application of the optional stopping theorem to characterize the principal's value function and the bounding techniques we use to track the continuous state variable and prove optimality of the contract. We also present two noteworthy predictions.

The first prediction is that incentives change qualitatively over the life of the project – the contract has two phases owing to the minimum amount of time necessary to complete it. If a setback occurs in the first phase, at any $S_{t-} \geq \bar{X}$, then completion of the project in the remaining time is still possible and the contract does not respond to the setback directly, but continues to tick away the bonus the agent can earn from early delivery. If a setback occurs in the second phase, at any $S_{t-} < \bar{X}$, then it is no longer possible to complete the project in the remaining time and the contract calls for randomization between $S_t = 0$ (cancellation) to prevent shirking and $S_t = \bar{X}$ (extension) to restore feasibility. From this point on, the agent earns only a flat fee if the project is ultimately delivered. Thus, the contract exhibits a distinctive transition from monetary to deadline-based incentives at $S_t = \bar{X}$ corresponding to the focal kink in the principal's value function.

Second, as noted above, the optimal contract allows for the possibility of project cancellation after arbitrarily large overruns in time and budget. Of course, termination is a common feature of optimal dynamic contracts. What is more surprising is the possibility that the agent can be granted an indefinite number of extensions prior to cancellation, even in the face of his temptation to divert resources. To be sure, it is not difficult to imagine models in which projects may be canceled for exogenous reasons after operating for an arbitrarily long time (e.g., see Remark 3 and Corollary 3). Our focus here, however, is on projects that remain perpetually feasible. Thus, cancellation – when it occurs – is purely a manifestation of agency frictions. For example, at the time the projects were canceled it was presumably still possible to build a second and third nuclear power plant at the V.C. Summer site, where one had previously been built, or for SAIC to develop a functional case management software system for the FBI.[2].

The rest of the paper is organized as follows. In Section 2 we review related literature. Section 3 contains three subsections where we: introduce the model, derive the first-best when progress is publicly observed, and define an optimal contract when progress is only observed by the agent. In Section 4 we show that a time-budget contract maximizes the principal's value for any initial non-negative level of utility for the agent, and in Section 5 we derive the principal's value function and characterize the corresponding optimal initial utility to grant the agent. In Section 6 we sum up and suggest several avenues for future research. All proofs have been relegated to appendices.

2 Related Literature

The literature on the optimal provision of dynamic incentives is extensive and remains very active. Pioneering articles responsible for moving it forward 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), Williams (2011), and Bloedel, Krishna, and Strulovici (2021). Related to this paper, 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 ob-

[2]In fact, SAIC offered to complete the VCF project if granted one more year and an additional USD56 million, but the FBI's CIO rejected this offer and canceled the cost-plus-award-fee contract Marchewka (2010).

servable 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 with incentive provision. Biais, Mariotti, Rochet, and Villeneuve (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 a similar vein, Myerson (2015) explores a model in which a prince must provide a subordinate governor with incentives to exert hidden effort to reduce the Poisson arrival of observable crises. Interestingly, if the governor performs badly, then he is subjected to random termination, which mirrors a feature of our optimal contract, albeit for quite different reasons. In Myerson (2015) the threat of random termination induces the governor to exert high effort, whereas the promise of random extension in our model induces the agent to truthfully report the occurrence of late-stage setbacks. In contrast with all of these papers, the Poisson shocks in our model, the setbacks, are privately observed by the agent and arise as an unavoidable consequence of the production process – that is they are *discoveries*.[3] The potential for their occurrence gives the agent cover to commit moral hazard; i.e., to make plausible excuses for why project completion has been delayed.

Closest to ours are four papers that explore the optimal deadline for a project in the context of dynamic agency. The most salient of these is Green and Taylor (2016) (GT16 below) which, in a sense, explores the mirror image of the setting we investigate. In GT16, the project is complete once two Poisson breakthroughs occur, and the agent hired to run the project privately observes the occurrence of the first one. As in our setting, the agent can surreptitiously divert the principal's flow of investment in the project for private benefit, delaying progress.

We consider a complementary but substantially richer setting where progress advances continuously and in which a potentially infinite number of discrete setbacks may occur prior to completion. Methodologically this requires the contract in our model to track two continuous state variables, the agent's continuation utility and the current level of progress $X_t \in [0, \bar{X}]$. By contrast, the exogenous state in GT16 takes on only two possible values, 0 (no progress) or 1 (progress) and is irreversible.

[3] Indeed, if setbacks were publicly observed in our model, then they would provide evidence that the agent was working on the project.

Hence the agent in GT16 only makes a single report about when he observes the first breakthrough, whereas our agent may repeatedly report the occurrence of false setbacks or repeatedly postpone reporting the occurrence of real ones, or any combination of these. As a consequence, the agent in GT16 is eventually subjected to smooth stochastic termination at a constant rate for *not reporting* a breakthrough; whereas the agent in our model eventually faces discrete randomization between project extension and cancellation for *reporting* the occurrence of breakdowns. More, the probability of project termination in our model is not constant, as in GT16, but increases in the state, i.e., is higher the later in the schedule a setback is reported. Unlike the project in GT16, the discrete scope of our project $\bar{X} > 0$ generates a focal kink in the principal's value function that, among other things, results in the optimality of a short-leash contract for a generic set of parameter values.

Economically, the most significant difference between the two papers is in the type of enterprise they investigate. GT16 considers innovative projects for which the path to success must be discovered through a sequence of uncertain trials; e.g., searching for a new antibiotic by "sorting through thousands of promising compounds" Goodman (2023). Here, by contrast, we analyze projects for which the production process is planned and designed, but completion is subject to probabilistic mishaps; e.g., following blueprints to erect a building in the face of uncertainty about: differing cite conditions Amarasekara et al. (2018), weather delays Bae (2022), or equipment failure Baldwin et al. (1971). To be sure, many projects involve some weighted combination of breakthroughs and breakdowns depending on where they fall on the continuum between experimental and designed applications. The notable structural differences between the optimal incentive schemes identified in GT16 and here highlight the importance of understanding the predominant characteristics of the underlying mode of production and corresponding agency environment.

As for the other three related papers: first, Madsen (2022) 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. Second, Mayer (2022) 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 (2022), the principal must provide incentives for the agent to report that the project has gone bad and should be terminated. A third recent paper featuring a single privately observed transition is Curello and Sinander (2021). Similar in spirit but opposite in application to Madsen (2022), 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 differs from those studied in these three papers along a number of important dimensions. Rather than the arrival of a single irreversible transition, our agent may observe numerous setbacks, none of which render project completion infeasible. 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 delivery-contingent contract.

3 The Project, The First-Best, and the Contract

3.1 The Project

A risk-neutral principal (she/her) hires a risk-neutral agent (he/him) 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 $\bar{X} > 0$ (the *scope* of the project) before it is completed. As the agent works, *progress* X_t accumulates deterministically, but *setbacks* are discovered according to a Poisson process with arrival rate λ. We denote by N_t the process that counts the number of setbacks before or at time t. 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.

Three 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 are not due to the negligence or indolence of the agent.[4] Third, for analytic tractability, we assume that setbacks wipe out all progress (see Remark 3 in Subsection 4.4 for further discussion of this point).

The potential for moral hazard in this setting derives 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 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, whereas $a_t = 1$ represents *working*, which corresponds to using the supplied funds to develop the project. Shirking yields the agent a private flow benefit of b.

Assumption 1 *Shirking (or diversion) is socially inefficient: $b < c$.*

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.[5]

---

[4]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.

[5]Our main results including the form of the optimal contract continue to hold if the principal and agent share a subjective discount rate, $r > 0$. This analysis is available upon request.

3.2 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 complete 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

$$F^{\text{FB}} = R - c\Delta, \tag{3}$$

where Δ represents the expected duration of completing the project:

$$\Delta \equiv \frac{e^{\lambda \bar{X}} - 1}{\lambda}. \tag{4}$$

It is straightforward to verify that Δ increases in $\bar{X}$ and λ and that $\lim_{\lambda \to 0} \Delta = \bar{X}$.

It follows immediately that the first-best policy is to start the project and run it until completed if and only if $F^{FB} \geq 0$ However, in the second-best, incentivizing the agent requires paying him rents, so a stronger assumption on the gross value of the project to the principal is required:

Assumption 2 *$R - (c + b)\Delta > 0$.*

As we show in Proposition 6 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. Indeed, implementing such an outcome is impossible, as we discuss below.

3.3 Unobservable Progress

A fundamental incentive problem arises when the agent's expertise and decision to work on the project allow him to *privately* observe its state at each instant. Specifically 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, he could shirk for a spell before reporting it. Moreover, so long as he is not terminated, the agent may engage in these forms of misconduct ad infinitum.

However, the agent's reports must be consistent with *some* feasible path under the recommended actions; otherwise the principal can be certain the agent has lied. In particular, the agent cannot go longer than $\bar{X}$ time without reporting a setback or he will be fired for not delivering the completed project (see Lemma 1, below).

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 need 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. 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. Both can involve public randomization. We also allow the principal to provide the agent with rewards based on reported project status; however, this turns out not to be optimal in this setting as stated in Lemma 1 below. Specifically, 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, *the prize*, granted upon successful completion.

Definition 1 (Contract) *Denote the probability space as $(\Omega, \mathcal{F}, P)$, with the filtration as $\{\mathcal{F}_t\}_{t\geq 0}$ generated by the history of progress $\{X_t\}_{t\geq 0}$, reports $\{\hat{X}_t\}_{t\geq 0}$, and any public randomization process (the agent's information set). The filtration $\{\hat{\mathcal{F}}_t\}_{t\geq 0}$ is generated by the history of reports $\{\hat{X}_t\}_{t\geq 0}$ and any public randomization process (the principal's information set). Contingent on the filtration $\hat{\mathcal{F}}_t$, the contract specifies: a stopping time τ when the contract ends, from completion (τ^C) or termination (τ^T) with $\tau = \min\{\tau^C, \tau^T\}$, a terminal payment process $\{K_t\}_{t\geq 0}$ to the agent (a prize if $\tau = \tau^C$ and severance if $\tau = \tau^T$), and cumulative intermediate rewards $\{C_t\}_{t\geq 0}$. All quantities are assumed to be integrable and measurable under the usual conditions.*[6]

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

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

The agent's limited liability requires $K_\tau \geq 0$ and C_t to be non-decreasing, so $W_t \geq 0$.

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 reward to the agent:

$$F_t = \mathrm{E}\left[R_\tau - \int_t^\tau c ds - \int_t^\tau dC_s - K_\tau \middle| \hat{\mathcal{F}}_t\right], \tag{7}$$

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

[6] As is standard, we use $t-$ to mean the left-limit at t; e.g., if a setback occurs at time t, then $X_t = 0$ and X_{t-} is the value of progress immediately before the setback.

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

Lemma 1 (High Action and Intermediate Consumption) *(i) The principal always implements the high action ($a_t = 1$). (ii) The principal does not award the agent intermediate consumption ($dC_s = 0$). (iii) If the agent is verifiably detected misallocating resources or lying about the project's progress, then he is terminated at that point without severance.*

These claims are standard, so formal proofs are omitted.[7]

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

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

We allow the principal to randomize termination and payoffs as part of the contract. In fact, we will show that some randomization is required for the contract to be optimal. We describe the probability of random termination for incentive compatible contracts with

$$\Phi_t(s) = \mathrm{E}\left[1\{\tau^T > t+s\}\middle|\,\hat{\mathcal{F}}_t\right] \tag{8}$$

for $s \geq 0$. $\Phi_t(s)$ is the probability that the agent survives random termination through $t+s$ given the history of public randomization and reports at t. $\Phi_t(s)$ is weakly decreasing in s and non-negative. $d\Phi_t(s)$ is understood to mean the increment over s, holding t constant.

Randomization of termination or payoffs can be discrete or continuous. Continuous randomization of termination means that $\Phi_t(s)$ declines continuously in s. For

[7] Heuristically,, we note the following: (i) Shirking is never optimal because any contract that involves an interval $\delta > 0$ of shirking can be modified to eliminate the interval by giving the agent a lump-sum payment equal to what he would have obtained from shirking. Because X does not change while the agent shirks and neither party discounts future payments, this does not change incentives. However, it strictly increases the principal's expected payoff because $b\delta < c\delta$. (ii) Because both parties are equally patient, payments may be delayed without loss. (iii) The revelation principle ensures optimality of a truthful direct mechanism, and the harshest penalty for verifiable malfeasance should be imposed; with limited liability, this is termination without severance.

example, the principal may control an independent Poisson process with publicly observable path and intensity such that the agent is terminated when the Poisson process hits. Discrete randomization of termination means that $\Phi_t(s)$ may drop discretely. For example, there may be a stopping time at which the agent is terminated with positive probability, and we can have $\Phi_t(0) < 1$.

We define optimal contracts as:

Definition 3 (Optimal Contract) *A contract is optimal if it maximizes the principal's objective function (7) over the set of contracts that are IC and grant the agent non-negative initial utility:* $W_0 \geq 0$.

In Section 4, we show that a time-budget contract attains the principal's maximum value for any initial level of utility for the agent, W_0, while in Section 5, we characterize the principal's value and derive the optimal level of utility to grant the agent, W_0^*.

4 Optimal Contracts

This section contains five parts. After preliminaries, we introduce and discuss a crucial constraint in Subsection 4.1, the *No-Postponed-Setbacks* (NPS) condition, that is necessary for incentive compatibility. Then in Subsections 4.2 and 4.3 we characterize the optimal termination policy and prize for completion respectively, under the assumption that (NPS) is also sufficient for incentive compatibility. This analysis yields three exhaustive contractual properties. Finally, in Subsection 4.4 we present the concept of a time-budget contract and show that it satisfies the three properties and is IC, implying that it attains the maximal payoff for the principal for any given value of W_0.

We begin with a martingale representation of the agent's continuation utility:

Proposition 1 (Martingale Representation) *If a contract is incentive compatible, there is an* $\mathcal{F}_t$*-predictable, non-negative, integrable process* $\{J_t\}_{t\geq 0}$ *such that*

$$dW_t = J_t(\lambda dt - dN_t) + dM_t \tag{9}$$

where M_t is an integrable $\mathcal{F}_t$-martingale orthogonal to N_t that captures any public randomization. For future use, we define

$$Y_t = W_{t-} - J_t \tag{10}$$

to be the agent's expected continuation utility after a setback but before any randomization is applied. Then, W_t is the agent's continuation utility after randomization is applied at t.

This proposition is proved in Appendix A.2.

To proceed, we use recursive methods for incentive compatible contracts that focus on one attempt at completion. This means that we start at some time t for which $X_t = 0$ and continue until the attempt ends at termination, a setback, or completion. Because the contract is incentive compatible, $dX_t = dt$ until the next setback, so we will track $x \in [0, \bar{X}]$ instead of time and suppress the t dependence to conserve on notation. We will use $Y(x)$ to mean Y_{t+x} for $x > 0$ and $Y(0) = \lim_{x \downarrow 0} Y(x)$; $\Phi_x(s)$ to mean $\Phi_{t+x}(s)$ or just $\Phi(s)$ for $x = 0$; $K(x) = K_{t+x}$ to mean the agent's severance pay if he is randomly terminated at $x < \bar{X}$ and $K(\bar{X})$ to mean the (possibly random) prize on completion.

4.1 Incentive Compatibility

We begin the analysis with 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 (NPS) is also sufficient to prevent any other deviation.

Proposition 2 (Incentive Compatibility) *A necessary condition for incentive compatibility is that for all $\delta < \bar{X} - x$*

$$Y(x) \geq b \int_0^\delta \Phi_x(s) ds + (1 - \Phi_x(0))K(0) - \int_0^\delta K(x+s) d\Phi_x(s) + \Phi_x(\delta) Y(x+\delta) \tag{NPS}$$

If the constraint (NPS) is also sufficient for incentive compatibility, then the principal does not award severance pay, $K(x < \bar{X}) = 0$, and the constraint holds with equality:

$$Y(x) = b \int_0^\delta \Phi_x(s) ds + \Phi_x(\delta) Y(x + \delta) \tag{11}$$

This is proved in Appendix A.3. To understand (NPS), suppose the state of progress is $X_{t-} = x$ when a setback hits. Consider two possible paths the agent can take:

- [Work] The agent reports the setback immediately, and then works as desired.
- [Shirk] The agent delays reporting the setback and attempts to shirk for time $\delta \leq \bar{X} - x$. Then, if he is not terminated while shirking, he reports a (bigger) setback and works as desired.

Now, we compare the two paths, with working first. Since working is optimal, truthful reporting gives the agent $Y(x)$, his promised expected utility after a setback.

Next, we consider shirking. 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 $X_{t+\delta}$ 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. In this case, the agent's continuation utility at x is

$$\underbrace{b \int_0^\delta \Phi_x(s) ds}_{\text{benefits from shirking}} + \underbrace{(1 - \Phi_x(0)) K(0) - \int_0^\delta K(x + s) d\Phi_x(s)}_{\text{severance pay}} + \underbrace{\Phi_x(\delta) Y(x + \delta)}_{\text{late setback report}}$$

The agent has three sources of utility: the direct benefit from shirking (first term), any severance pay from termination (middle term), and any utility granted after reporting the delayed setback (last term). We do not have to keep track of intermediate changes in utility because there are only three possible ways for an attempt at progress to end, and the third (completion) cannot happen if the agent is shirking.

The agent's utility from shirking is complicated by the fact that the principal can perform random termination; the probability that the agent survives for an additional time δ after x is $\Phi_x(\delta)$. So, if the agent attempts to shirk for δ time, the average length of time he will shirk before being randomly terminated is $\int_0^\delta \Phi_x(s)ds$. The agent survives to report the setback at $t+\delta$ with probability $\Phi_x(x+\delta)$, and dissipates his private information to receive $Y(x+\delta)$. Similarly, the severance pay term captures any payments the agent receives after random termination.[8] The incentive constraint (NPS) just stipulates 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.

Next, we observe that if (NPS) is sufficient for incentive compatibility, then the principal will never award severance pay. Severance pay weakens incentives: $\Phi_t(s)$ decreasing in s implies that the severance pay terms in (NPS) are positive, making the incentive constraint harder to meet. The principal has fewer paths of $Y(x)$ that meet the constraint and are incentive compatible. As we show in the proof, this is sufficient to exclude severance pay from the optimal contract.

Proposition 2 also implies that the agent loses expected utility between truthfully reported setbacks. Substituting $x = 0$ into (NPS) or (11), $Y(0)$ is the agent's expected continuation utility after truthfully reporting a setback, and $Y(\delta)$ is his expected continuation utility after truthfully reporting a setback δ time later. If the first setback occurs at t, W_t is the agent's utility after the setback *and* any immediate randomization. Thus, we have

$$W_t - \Phi_t(\delta)Y_{t+\delta} \geq b\int_0^\delta \Phi_t(s)ds \tag{12}$$

and the loss of continuation utility between truthfully reported setbacks is at least equal to the expected time the agent could have spent shirking between them. We call this a "round trip" because the agent goes from $X = 0$ through some path and back to $X = 0$. If setbacks were observable, the agent would not be punished for

[8] The term $(1 - \Phi_x(0))K(0)$ captures incentives like "if the agent does not report a setback at x, terminate him with positive probability." Similarly, $\int_0^\delta K(x+s)d\Phi_x(s)$ captures the possibility of random termination with severance over $(x, x+\delta)$.

them because they are evidence that he is working. However, because the principal cannot detect them, the agent can manipulate reporting setbacks so as to provide cover for shirking. Therefore, the agent's loss in utility between reported setbacks must be at least the expected benefit from shirking during that spell.

The round trip penalty demonstrates how the incentive compatibility constraint makes time into a scarce resource: as progress and setbacks accumulate, the agent's utility moves over round trips toward zero, where the limited liability constraint can bind. In the next section, we show how this process determines the optimal randomization procedure.

4.2 Termination and Randomization

We now consider the optimal termination policy 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 with positive probability. However, the agent could simply mimic that path with his reports while shirking, and thus obtain infinite utility. Put differently, the NPS round trip penalty implies that between any two setbacks, the agent loses at least as much expected continuation value as he could have obtained by shirking. However, absent random termination, termination must still 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 either be randomly terminated or eventually run out of time.

Second, termination is random. Imagine that the principal does not randomize termination and a setback occurs at t resulting in $W_t = Y_t \in (0, b\bar{X})$. 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 (12) with $\Phi_t(s) = 1$ would result in $Y_{t+\delta} < 0$, which is not feasible. The agent would prefer to shirk rather than to report the second setback. One option to prevent this is simply to terminate the contract at t and give the agent a severance payment of W_t, but this is not optimal (Proposition 2). Instead, there is a better alternative: use randomization to either fire the agent without severance 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.

In general, randomization can be discrete or continuous. One possibility is to wait until some randomization is required, and then randomize as much as necessary all at once: when a setback results in $Y_t < b\bar{X}$, randomize immediately to $W_t = b\bar{X}$ or 0. Another possibility is continuous randomization: let $\Phi_t(s)$ decline continuously so as to keep W_{t+s} high enough after surviving random termination that the agent can report a setback without violating limited liability. In an extreme form, the principal could set $\Phi_t(s)$ to keep continuation utility from surviving termination as constant. These methods could be combined: randomize after a setback results in $Y_t < b\bar{X}$, but just enough to keep the agent's utility non-negative after a second setback. In fact, it is optimal only to randomize discretely and only when necessary:

Proposition 3 (Optimal Randomization) *If (NPS) is sufficient for incentive compatibility, then the principal randomizes only when $X_t = 0$ and only if $Y_t < b\bar{X}$. The agent survives random termination with probability $\frac{Y_t}{b\bar{X}}$ and is terminated with the complementary probability. Upon survival, the agent's continuation utility is set to $W_t = b\bar{X}$. The project is not otherwise terminated.*

This is proved in Appendix A.4. We utilize an inductive approach, beginning by showing the optimal randomization policy for a project where no setbacks are possible, then applying a sequence of bounding arguments to show that the result continues to hold for projects that may have an arbitrarily large number of setbacks.

A key friction is that the agent can always gain utility from shirking, so there is a lower bound on how much utility the contract can promise at the start of an attempt ($X_t = 0$), and that bound is higher than zero due to limited liability: $W_t \geq b\int_0^{\bar{X}} \Phi_t(s)ds$. This constraint is still valid even in the modified no-setback economy because the agent can still shirk, mis-report progress for $\bar{X}$ time, and fail to deliver. In this modified economy, we can fix the survival probability at completion ($\Phi(\bar{X})$) and ask what is the best way to get there. Randomizing as early as possible means setting $\Phi(s) = \Phi(\bar{X})$ for all $0 \leq s < \bar{X}$, and this policy reduces expected operating cost $c\int_0^{\bar{X}} \Phi(s)ds$, and relaxes the incentive constraint without impacting the expected project return $\Phi(\bar{X})R$. Hence, the principal's problem in this no-setback economy reduces to looking for a corner solution to the following problem:

$$\max_{\Phi(\bar{X})} \Phi(\bar{X})(R - c\bar{X}) - W_t, \quad s.t. \quad W_t \geq b\Phi(\bar{X})\bar{X}, \tag{13}$$

for which the randomization policy in Proposition 3 is optimal. We apply a sequence of bounding arguments to show that the corner solution continues to obtain in an economy in which an arbitrary number of setbacks are possible. A setback may trigger random termination, but, conditional on surviving, it is optimal never to terminate the agent before the next setback.

Propositions 2 and 3 together yield the following useful result:

Corollary 1 (Round-trip Utility Loss) *If (NPS) is sufficient for incentive compatibility, then the agent's expected utility loss between two setbacks occurring at t and $t+\delta$ is $b\delta$ under the optimal contract, i.e., $W_t - Y_{t+\delta} = b\delta$.*

This is proved in Appendix A.4.

This reveals a second key feature of the optimal contract, a linear countdown. Incentive compatibility (Section 4.1) showed that time is a scarce resource for the agent. Optimality now shows that the relationship between time and the agent's utility is linear. Next we consider the prize for completion.

4.3 The Prize for Completion

Subsections 4.1 and 4.2 reveal two of the three key features of an optimal contract. First, there is a soft deadline, meaning the agent's utility is randomly re-set whenever he reports a setback with $Y_t < b\bar{X}$. Second, there is a linear countdown, meaning that between two reports of setbacks, the agent's utility must decline linearly at rate b. The third and final feature to determine is the structure of the prize for delivery of the completed project.

In fact, randomization only at $X_t = 0$ and a linear countdown are enough to fully characterize the agent's utility process (9) and thus the prize upon completion. For an attempt that starts at t with $X_t = 0$ and lasts through $t + \delta$, there are no setbacks and no randomization on $[t, t+\delta]$. Because there are no setbacks and no randomization, we can treat $W_{t+\delta} - W_t$ as a deterministic function of δ. Then the martingale representation (9), definition of Y_t, and Corollary 1 show that

$$\frac{\partial[W_{t+\delta} - W_t]}{\partial\delta} = \lambda J_{t+\delta} = \lambda\left(W_{t+\delta} - Y_{t+\delta}\right) = \lambda\left([W_{t'+\delta} - W_{t'}] + b\delta\right) \tag{14}$$

$$W_{t+\delta} - W_t = -b\delta + \frac{b}{\lambda}\left(e^{\lambda\delta} - 1\right) \tag{15}$$

where the second line follows from integrating the first with respect to δ.

The prize is thus

$$K_{t+\bar{X}} = W_{t+\bar{X}} = W_t - b\bar{X} + \frac{b}{\lambda}\left(e^{\lambda\bar{X}} - 1\right) = W_t - b\bar{X} + b\Delta \tag{16}$$

where Δ is the expected duration given in (4).

Starting at time 0 and re-applying the results of Corollary 1 over repeated setbacks shows the third required feature of the contract: conversion from time to money. The agent begins with a prize of $W_0 + b\Delta$. Before randomization, time spent reduces the prize at rate b until it equals $b\Delta$. Once randomization begins, the agent is re-set to $W_t = b\bar{X}$ after each extension and (16) yields a constant prize $b\Delta$. We formalize this below:

Corollary 2 (The Prize) *The prize begins at $K_0 = W_0 + b\Delta$, drifts down at rate $b\,dt$ until it equals $b\Delta$, after which it remains constant.*

Given a starting value for the agent's payoff W_0, having specified the law of motion for his continuation utility after and between setbacks (equation (15) and Corollary 1), randomization only at the time of setbacks (Proposition 3), and the prize upon completion (16), there are no remaining contractual degrees of freedom. In particular, if (NPS) is sufficient for incentive compatibility, then all elements of an optimal contract specified in Definitions 1 and 3 (i.e., stopping time τ, terminal payment process $\{K_t\}$, and cumulative intermediate rewards $\{C_t\}$) have been determined. Thus, we need only find an implementation that generates these elements and show that it is IC, which we do in the next subsection.

4.4 Optimality of Time-Budget Contracts

We are now ready to present our main result: the optimal mechanism can be implemented with a *time-budget contract*. We do so in three steps: first, we formally define a time-budget contract. Second, we show that a time-budget is IC. Third, we show that the time-budget satisfies all three properties of an optimal contract obtained in the previous subsections. Since those properties were derived under the assumption that (NPS) is sufficient to be IC, the fact that the time-budget contract is itself IC implies that it must be optimal.

Time-budget contracts are mechanisms that give the agent a *quantity of time* for which the principal will fund the project and a reward for completion. The agent spends time working on the project and can request extensions that are stochastically granted. If the agent runs out of time, project funding ends and the agent is terminated without severance. If the agent completes the project, he receives a reward that includes both a fixed payment and an amount proportional to the remaining time. Thus, the initial endowment of time – the time budget – is converted into either funding for the project or a reward for early completion.

Definition 4 (Time-Budget Contracts) *A time-budget contract is an initial grant of time S_0 and a payment on success at time τ of $K_\tau = b(\Delta + S_\tau)$. S_t counts down deterministically ($dS_t = -dt$) unless the agent requests an extension, which he is free to do at any time. Requests for extensions are ignored if $S_t \geq \bar{X}$. If $S_t < \bar{X}$, an extension is granted with probability $\frac{S_{t-}}{\bar{X}}$, resetting $S_t = \bar{X}$. If the extension is not granted, the agent is terminated without severance (S_t is set to zero).*

Under a time-budget contract, the agent's objective is to maximize (6), subject to the evolution of the project (1) and the time budget: $dS_t = -dt + L_t d\hat{N}_t$, with $\hat{N}_t$ counting extension requests. L_t is the realized length of the extension; if $S_t < \bar{X}$, L_t takes a value of $\bar{X} - S_{t-}$ with probability $\frac{S_{t-}}{\bar{X}}$ and a value of $-S_{t-}$ with probability $1 - \frac{S_{t-}}{\bar{X}}$.

A time-budget contract is IC; it is optimal for the agent never to shirk and to truthfully and immediately report a setback (request an extension):

Proposition 4 (Time-Budget Incentive Compatibility) *A time-budget contract is IC. The agent obtains continuation utility $W(X, S) = bS + \frac{b}{\lambda}\left(e^{\lambda X} - 1\right)$.*

This is proved in Appendix A.5 using dynamic programming techniques to derive the agent's value function if given a time-budget contract. If $S_t \geq \bar{X}$, the agent is always indifferent between working and shirking because shirking requires reporting a setback, and the round trip penalty between setbacks exactly cancels out the private benefits from shirking. If $S_t < \bar{X}$, then the agent strictly prefers not to shirk if $X_t > 0$ and is indifferent for $X_t = 0$. If $X_t > 0$, then shirking requires him to report a setback at some point, and he risks losing his positive progress if he is terminated. If $S_0 < \bar{X}$, the agent will request an extension immediately .

A time-budget contract features the three key properties of an optimal contract identified in the previous subsection: a soft deadline, a linear countdown, and a conversion from time to money on completion. Because these three properties of an optimal contract are exhaustive and were derived under the assumption that (NPS) is IC, the fact that a time-budget contract *is* IC proves that it must be optimal:

Theorem 1 (Optimal Contracts) *A time-budget contract with $bS_0 = W_0$ maximizes the principal's objective function among incentive compatible contracts that deliver the agent initial value W_0.*

This is completed in Appendix A.6.

Remark 1 (Relaxing Reporting) Optimal incentives can be implemented with less stringent reporting requirements than specified. Rather than requiring a report of each setback, at project inception the principal can announce a *target date* $T = S_0 - \bar{X}$ and then commit to fund the project until this point *no-questions-asked*. If the agent delivers the completed project at $\tau \leq T$, he receives K_τ as specified in Definition 4. Once the target date has passed, the principal requires setbacks to be reported, and she follows the random termination procedure given in the definition of a time-budget contract from that point on.

Remark 2 (Procurement Contracts) As noted in the introduction, the implementation of the optimal incentive mechanism characterized in Definition 4 is a cost-plus-award-fee contract. The principal commits to: cover the operating cost of the project $c\tau$, pay a fixed fee $b\Delta$ upon project delivery, and pay an incentive award bS_τ for early completion.

Remark 3 (Generalizations) Although a full analysis of heterogenous setbacks is beyond the scope of this investigation, the structure outlined in Definition 4 remains qualitatively intact under various alternative specifications such as the one concerning fatal setbacks, below. Other environments (e.g., partial setbacks) in which the agent's persistent private information is not cleanly dissipated are more technically challenging. Analysis of such settings requires development of new methods and/or stronger assumptions (e.g., pre-specifying state variables and using Markovian contracts).

To see one way the model can be generalized, suppose that when a setback occurs it is either an un-fixable dead end with probability $\gamma \in (0, 1)$ or progress is wiped out but the project remains feasible with probability $1 - \gamma$. The agent privately observes the occurrence and the type of each setback.

Corollary 3 (Fatal Setbacks) *The value-maximizing contract in this setting is a modified time-budget contract with $W_0 = bS_0$. If the agent reports a fatal setback at τ, then the contract is terminated with a severance payment of bS_τ. Reports of non-fatal setbacks are treated as in Proposition 3.*

We sketch a proof of this in Appendix A.7.

In this version of the model the project can be canceled for one of two reasons: the exogenous arrival of a fatal setback or the endogenous decision not to grant an extension. Although fatal setbacks are plausible in some settings, as noted in the Introduction, our focus is mainly on planned or designed projects that remain perpetually feasible.

5 The Value of the Project

5.1 The Initial Value

We can use the fact that a time-budget contract is an optimal implementation to derive the principal's value function $F(S, X)$. What is most important is her valuation of a given time budget when starting from scratch: $F(S, 0)$.[9] In fact, $F(S, 0)$ is

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

characterized by a Delay-Differential Equation (DDE) and is discontinuous in its first- and second- derivatives at a focal point, $S_t = \bar{X}$.[10] So, instead of the usual techniques involving PDEs and dynamic programming, we will derive $F(S, 0)$ using two martingales, the optional stopping theorem, and a relatively simple DDE. This allows for a more fundamental understanding of the contract and its economic characteristics.

In order to simplify the principal's problem, we will define two useful auxiliary functions. 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 the contract ends (from either completion or termination). We have

$$\pi(S) = \mathrm{E}_t \left[1_{X_\tau = \bar{X}} | S_t = S, X_t = 0 \right] \tag{17}$$

$$\sigma(S) = \mathrm{E}_t \left[\tau - t | S_t = S, X_t = 0 \right], \tag{18}$$

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 schedule (the time budget) which reduces both the probability of success and the time allowed.

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

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

$$= \pi(S)R - bS - c\,\sigma(S) \tag{20}$$

where $W(X = 0, S) = bS$ from Proposition 4. 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 to the agent by the time budget, minus the expected direct running cost.

[10] It is possible to derive the existence of the kink using the welfare bounding arguments of Appendix A, as we demonstrate in the On-line Appendix, Lemma C.1. This kink, and its simple characterization using a DDE, motivate the methods of this section.

We continue by applying the optional stopping theorem[11] to two martingales in order to relate the probability of success to the expected time remaining. First, the agent's continuation utility $W(X,S) = bS + \frac{b}{\lambda}\left(e^{\lambda X} - 1\right)$ is a martingale that can only stop at two boundaries: project success or termination. Second, the randomization probabilities (Definition 4) imply that $S_t + t$ is a martingale that also only stops upon completion or termination. , so the optional stopping theorem implies $S_0 + 0 = \mathrm{E}\left[S_\tau + \tau\right]$. Thus we have

$$bS_0 = W(0,S_0) = \mathrm{E}\left[W(X_\tau, S_\tau)\right] = bS_0 + \mathrm{E}\left[-b\tau + \frac{b}{\lambda}\left(e^{\lambda X_\tau} - 1\right)\right] \tag{21}$$

$$0 = -b\sigma(S_0) + \mathrm{E}\left[\frac{b}{\lambda}\left(e^{\lambda X_\tau} - 1\right)\right] \tag{22}$$

$$= -b\sigma(S_0) + b\Delta\pi(S_0). \tag{23}$$

The first line follows from plugging in the form of W and applying the optional stopping theorem; the second line follows from the definition of $\sigma(S)$; the third line follows from the definitions of $\pi(S)$ and Δ and the optional stopping theorem applied to success or termination.

Re-arranging and generalizing to any time with $X_t = 0$, we have

$$\pi(S) = \frac{\sigma(S)}{\Delta}. \tag{24}$$

which illustrates the direct link between reducing the projects expected time and reducing the projects expected success. In other words, the probability of successful completion is the ratio of the expected duration of the project under the time-budget $\sigma(S)$ to the expected duration in the first-best setting. The intuition for this result comes from the martingale property of the agent's continuation utility. The agent's utility $W(X,S)$ 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 the probability that the agent receives the

[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).

constant part of his reward, and the average time to completion must be proportional to the probability of success.

Taking (24) to the principal's valuation (20), we have four equivalent expressions:

$$\begin{aligned} F(S,0) &= \left(\frac{R}{\Delta} - c\right)\sigma(S) - bS \\ &= \pi(S)\left(R - c\Delta\right) - bS \\ &= \pi(S)F^{\mathrm{FB}} - bS. \end{aligned} \tag{25}$$

This highlights the principal's trade-off in choosing an optimal schedule length. Giving the agent more time (S or $\sigma(S)$) grants an increase in the probability of success.[12] However, more time also grants the agent more opportunity for malfeasance and therefore increases the initial utility that the contract must promise the agent, bS. In fact, the trade-off is so direct that the principal's expected payoff is the probability that the project is completed in the allotted time times the first-best payoff, minus the agent's information rents. 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. Hence, there is an optimal amount of time $S_0 = S^*$ and corresponding optimal level of promised utility $W_0^* = bS_0^*$ for the principal to initially grant the agent:

Proposition 5 *The principal's initial valuation is given by any of* (25)*. These functions exist, are unique, are weakly concave, attain a single local and global maximum in* S*, and are differentiable except at* $S = \bar{X}$*. We have* $\sigma(S) = S$ *when* $S \leq \bar{X}$*, and* $\sigma(S) < S$ *when* $S > \bar{X}$*. Furthermore,*

$$\lim_{S\uparrow\bar{X}} \sigma'(S) = 1 > 1 - e^{-\lambda\bar{X}} = \lim_{S\downarrow\bar{X}} \sigma'(S) \tag{26}$$

[12]In fact, a long time budget implies a small probability of cancellation, and in the limit as $S \to \infty$, the project is never canceled so that the expected stopping time corresponds to the expected first-best duration of the project. i.e., $\lim_{S\to\infty} \sigma(S) = \Delta$. Similarly, (and matching (24)), 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 converges to 1.

The maximizing value of S is $S^* \geq \bar{X}$. The optimal contract grants the agent initial utility $W_0^* = bS^*$ and initial time $S_0 = S^*$.

This is proved in Appendix B. $F(S, 0)$ and $\sigma(S)$ are plotted in Figure 1.

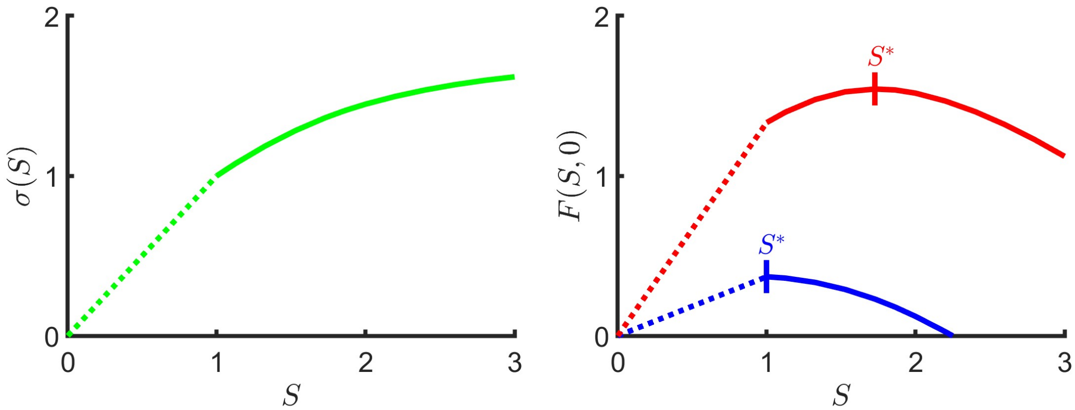

Figure 1: **The Expected Remaining Time and The Principal's Value Function**
The left panel plots $\sigma(S)$ (the expected remaining time until the contract ends, green line) and the right panel plots $F(S, 0)$ (the principal's value function at $X = 0$) for two values of R, 2.75 (blue line) and 4 (red line), with vertical markers indicating the maximizing value of S, labelled S^*. All plots use a dashed line for the linear segment $S \leq \bar{X}$, exhibit a non-differentiable point at $S = \bar{X} = 1$, and use a solid line for $S > \bar{X}$. Other parameters are $\lambda = 1$, $\bar{X} = 1$, $c = 1$, $b = 3/4$.

By definition, the expected time remaining, $\sigma(S)$, is the sum of the time to the next setback or completion and any residual expected time. For $S \leq \bar{X}$, the optional stopping theorem implies $S_0 + 0 = \mathrm{E}\left[S_\tau + \tau\right]$, and (18) gives us $\sigma(S) = S$. For $S > \bar{X}$, the recursive formula for $\sigma(S)$ is

$$\sigma(S) = \int_0^{\bar{X}} \lambda e^{-\lambda t}(t + \sigma(S - t))dt + e^{-\lambda \bar{X}} \bar{X}. \tag{27}$$

As we show in Appendix B.1, this corresponds to the DDE

$$\sigma'(S) = \lambda e^{-\lambda \bar{X}} \left(\Delta - \sigma(S - \bar{X})\right) \tag{28}$$

which allows us to characterize the principal's value function in Proposition 5.

Observe that the time-budget contract exhibits a strong policy switch by the principal coinciding with the kink at $S = \bar{X}$. For $S < \bar{X}$ the principal extends or cancels the project based on the agent's reports, and randomization after setbacks implies that the project is never completed with extra time remaining. In contrast, for $S > \bar{X}$ the principal can simply wait, ignoring reports and paying on delivery if and when it occurs. This policy switch at $S = \bar{X}$ marks an important focal point in the contract.

First, the focal point marks a transition in the expected duration of the project. Since $\sigma(S) = S$ for all $S \leq \bar{X}$, if a setback occurs at $S \leq \bar{X}$, then (prior to randomization) the expected duration of the project coincides with the remaining value of the time budget. In contrast, for $S > \bar{X}$ setbacks are responded to passively as time simply continues to tick down. In fact, since the project still can be completed early, before running into the $S \leq \bar{X}$ region, we have $\sigma(S) < S$ (as shown in Proposition 5). Thus although the optimal contract allows for extensions, it is not the case that the expected duration of the project is longer than the initial time budget. Instead, overruns in which the realized time taken τ is greater than S^* are a probabilistic consequence of extensions granted following late-stage setbacks. Indeed, if $S^* > \bar{X}$, then the contract initially builds *slack* or *slippage time* into the schedule in the sense that the expected allotted time S^* is longer than the expected duration of the project $\sigma(S^*)$.

Second, the focal point kink at $S = \bar{X}$ also delineates a transition in the optimal incentive scheme from monetary-based to deadline-based incentives owing to the minimum time required to perform the project. If $S_{t-} \geq \bar{X}$ when a setback hits, then completion of the project in the remaining time is feasible and the optimal contract does not respond directly to the setback, but continues to tick away the bonus for early delivery. On the other hand, if $S_{t-} < \bar{X}$ when a setback occurs, then completion in the remaining time is not possible. As we have shown, it is optimal at this point to randomize between $S_t = 0$ and $S_t = \bar{X}$. Specifically, if the outcome of randomization were any $S_t \in (0, \bar{X})$, then the agent would find it optimal to shirk away the allotted time. Hence, $S_t = 0$ (cancellation) prevents shirking and $S_t = \bar{X}$ (extension) restores project feasibility and induces continued working.

The kink arises from a discontinuous change in the marginal value of time, due to the all-or-nothing nature of the optimal randomization. As we showed in Section 4.2, the principal optimizes discretely at $X = 0$, and the reason is that policy allows the principal to grant the agent a low level of average utility in the cheapest way. However, key to this optimization is that the principal knows how much time the project would take to complete without a setback: the project has a known *scope*. The uncertainty is over whether the agent will be able to execute, but a successful attempt always costs $c\bar{X}$ and grants the opportunity to shirk $b\bar{X}$. Those values are re-set whenever an extension is granted, and so the agent's utility can be re-set fully. The smooth-pasting condition that one might expect instead of a kink originates in a continuous optimization problem over the location of a boundary, but our boundary is discrete. Any randomization above $\bar{X}$ implies an excess risk of cancellation and thus the marginal value of time drops discretely, not continuously, at $\bar{X}$.

The foregoing discussion implies that the minimum optimal value for the initial time budget is $S_0 = S^* = \bar{X}$, and, because of the kink at this point, there is a generic set of parameter values that all result in the same optimal contract. We explore this in the next subsection.

Remark 4 (Commitment to Randomize) In the online appendix we present an extension of the baseline model that relaxes the commitment of the principal to use explicit probabilities for project cancellation. As is usual in dynamic mechanisms that use termination to provide incentives, a fully renegotiation-proof contract is not feasible in our setting because the principal prefers extension to cancellation. Nevertheless, the principal's commitment to explicit probabilities can be relaxed by considering either a very long extension granted with very low probability, or a mixed-strategy time-budget implementation. In the latter, the principal randomizes between cancellation and extension just as in the baseline model, and upon receiving an extension, the agent randomizes between continuing to work and shirking out the clock. Although both relaxations reduce the initial value of the project to the principal relative to the commitment case, we show that the principal responds to her inability to commit to randomize by *increasing* the time and resources allocated to the project. By granting the agent a longer initial time budget, S_0, the principal raises the probability that the project will be completed before the randomization region, $S_t < \bar{X}$, is reached.

5.2 A Short-Leash Contract

A particularly salient situation is when the principal finds it optimal to set the expected schedule to the minimal time necessary to complete the project, $S^* = \bar{X}$ and responds to any reported setback by either canceling the project or resetting the clock. Such a *short-leash contract* possesses no slack in the sense that the expected stopping time equals the initial time allotted; i.e., early completion is not possible:

Proposition 6 (Short-Leash Contract) *If $S \leq \bar{X}$, then $F(S,0) = S\left(\frac{R}{\Delta} - c - b\right)$, which is increasing. The optimal contract involves the minimal initial time budget $S_0 = \bar{X}$ and a fixed prize $b\Delta$ iff*

$$(c+b)\Delta < R < \left(c + \frac{b}{1-e^{-\lambda \bar{X}}}\right)\Delta. \tag{29}$$

This is proved in Appendix B.

The first inequality in (29) is a restatement of Assumption 2, and it ensures that $F(S,0)$ is positive and increasing for $S < \bar{X}$. Therefore, it is never optimal to start the contract with $S_0 < \bar{X}$. The second inequality then ensures that $F(S,0)$ is decreasing for $S > \bar{X}$. Hence, when (29) 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 $S^* = \bar{X}$. The blue line on the right panel of Figure 1 demonstrates this case. In other words, the principal should keep the agent on a short leash: granting him only the minimal 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 (29) 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 (29) fails and moral hazard precludes the project from ever getting off the ground. On the other hand, if b is too small, then the right inequality fails. In this case, hidden action rents are less of a concern, 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 (29) 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 expected rent of $b\bar{X}$ necessary to induce the agent to work on the project.

At inception of a short-leash contract, the expected duration of the project is $\sigma(S = \bar{X}) = \bar{X}$. However, if the agent reports a setback with S_{t-} left on the schedule, then he is granted an extension of $\bar{X} - S_{t-}$ 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. We describe the probability of on-time completion, early cancellation, overruns, and overruns with cancellation analytically and graphically in On-line Appendix E.

6 Conclusion

At a general level, projects are often viewed as possessing three defining features: scope, schedule, and budget, the so-called *iron triangle*. The scope of a project is the quality of the deliverable; The schedule is the time allotted to production; and the budget is the monetary or other physical resources dedicated to it. In this paper we hold scope fixed, and present 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., the FBI's VCF system), its progress is likely to be hampered by agency frictions. To study this, we embed a model of production involving random setbacks into a dynamic agency environment and solve for the optimal contract. This yields a number of novel insights and conclusions. Among the most robust are: 1) an optimal contract can be implemented with a soft deadline and a terminal payment resembling a cost-plus-award-fee contract; 2) penalties for reports of setbacks or delays are generally more severe the later they occur in development; and 3) mishaps that are reported near the end of the allotted schedule either result in project cancellation or minimally feasible project extension.

Many avenues remain open for future research. For example, expanding our current treatment to incorporate common strategies for dealing with the time pressure created by unanticipated setbacks is intriguing. 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. Also, there is the question of scope itself. Throughout we assume that the project is 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. We leave these considerations and others for future work, having judged this particular project to be completed (albeit after a setback or two).

Appendix – Proofs and Derivations

A The Optimal Contract

This section is a series of lemmas and propositions that lead up to our core result, Theorem 1.

A.1 Re-writing the Principal's Objective

We re-write the principal's optimization problem for incentive compatible contracts. We use recursive methods and the notation introduced in Section 4.1 that focuses on one attempt at completion. This means that we start at some time t for which $X_t = 0$ and continue until the attempt ends at termination, a setback, or completion.

In addition to what is mentioned in the text, we use $\widehat{Y}$ to mean the agent's continuation utility at the beginning of the attempt but before any discrete randomization is applied. Notice that $\widehat{Y} \neq \lim_{x \downarrow 0} Y(x)$: $\widehat{Y}$ is the value at the beginning of the attempt and before randomization at $X_t = 0$, while $Y(x)$ is the value after the next setback at $t + x$. Similarly, $g(Y)$ means the principal's continuation utility from a given attempt that begins with Y, before any discrete randomization is applied. If discrete randomization is applied when a setback is reported, we have $\Phi(0) < 1$. Social value is denoted with $h(Y) = g(Y) + Y$.

First, the principal's objective function (7) with $X = 0$ and initial agent's utility $\widehat{Y}$ can be written as

$$g(\widehat{Y}) = \int_0^{\bar{X}} \lambda e^{-\lambda x} \Phi(x) \left[g(Y(x)) - cx \right] dx + e^{-\lambda \bar{X}} \Phi(\bar{X}) \left[R - c\bar{X} - K(\bar{X}) \right] \quad (30)$$
$$- (1 - \Phi(0)) K(0) + \int_0^{\bar{X}} e^{-\lambda x} (K(x) + cx) d\Phi(x)$$

This is the sum of value after a setback (first term), and value after completion (second term), and value after paid severance (third term), which are the only three ways a run from $X = 0$ can end in an incentive compatible contract.

Second, because the agent's continuation utility is a martingale that ends in a setback, termination and severance pay, or project completion, we have

$$\widehat{Y} = \int_0^{\bar{X}} \lambda e^{-\lambda x} \Phi(x) Y(x) dx + e^{-\lambda \bar{X}} \Phi(\bar{X}) K(\bar{X}) \quad (31)$$
$$+ (1 - \Phi(0)) K(0) - \int_0^{\bar{X}} e^{-\lambda x} K(x) d\Phi(x)$$

Third, through integration by parts, we have

$$\int_0^{\bar{X}} cxe^{-\lambda x} d\Phi(x) = e^{-\lambda\bar{X}}\Phi(\bar{X})c\bar{X} - \int_0^{\bar{X}} \left(ce^{-\lambda x} - \lambda cxe^{-\lambda x}\right)\Phi(x)dx \tag{32}$$

Adding (30), (31), and (32) shows that the principal's objective function (7) at any time with $X = 0$ can be written as

$$g(\widehat{Y}) = \int_0^{\bar{X}} \lambda e^{-\lambda x}\Phi(x)\left[h(Y(x)) - \frac{c}{\lambda}\right] dx + e^{-\lambda\bar{X}}\Phi(\bar{X})R - \widehat{Y} \tag{33}$$

where $h(Y(x)) = g(Y(x)) + Y(x)$ is the sum of the principal's and agent's continuation values and the social welfare function. $h(Y)$ is weakly increasing because intermediate consumption is never optimal (Lemma 1).

This re-writing of the objective function shows that we can consider the principal's problem to be optimization of social welfare over the path of Y and Φ, subject to the IC constraint and the limited liability constraint (Y weakly positive).

A.2 Martingale Representation (Proposition 1)

We have $dC_t = 0$ for all $t < \tau$ (Lemma 1), and incentive compatibility implies $a_t = 1$ (Definition 2). Therefore (6) under incentive compatible contracts becomes $W_t = E\left[K_\tau \,|\mathcal{F}_t\right]$. Since W_t is an $\mathcal{F}_t$-martingale, and since N_t and any public randomization are orthogonal, a martingale representation theorem (e.g., Protter (2005), Theorem 44) implies there exists a $\mathcal{F}_t$-predictable, integrable process $\{J_t\}_{t\geq 0}$ such that

$$dW_t = J_t(\lambda dt - dN_t) + dM_t. \tag{34}$$

where M_t is an integrable $\mathcal{F}_t$-martingale orthogonal to N_t that captures any public randomization. Truthful reporting requires J_t to be non-negative.

A.3 Incentive Compatibility (Proposition 2)

This section proves Proposition 2, then Lemma A.1, which is a corollary lemma following from Section A.1.

Following the arguments in the text (Section 4.1), we require for all $\delta < \bar{X} - x$

$$Y(x) \geq b\int_0^{\delta} \Phi_x(s)ds + (1 - \Phi_x(0))K(0) - \int_0^{\delta} K(x+s)d\Phi_x(s) + \Phi_x(\delta)Y(x+\delta) \tag{35}$$

Next, we observe that the impact of severance pay is to increase the right hand side of (35) by $(1-\Phi_x(0))K(0)-\int_0^\delta K(x+s)d\Phi_x(s) > 0$. Thus, if (35) is sufficient for incentive compatibility, then positive severance pay after termination strictly reduces the set of possible paths of Y that the principal can use, while having no impact on the objective function (33) for incentive compatible contracts; the principal can weakly improve any incentive compatible contract with severance pay with one without severance pay. Thus, we can consider only contracts without termination-severance pay. From (35) and setting $K(x < \bar{X}) = 0$, we obtain

$$Y(x) \geq b\int_0^\delta \Phi_x(s)ds + \Phi_x(\delta)Y(x+\delta) \tag{36}$$

Finally, recall that the principal's objective function for incentive compatible contracts (33) is pointwise (weakly) increasing in Y. Then, for any given randomization policy, the path of Y is maximized point-wise if (36) binds everywhere. Thus, if (35) is sufficient for incentive compatibility, then we can consider only contracts that induce (36) constraint to hold with equality. □

Lemma A.1 (The Principal's Incentive Compatible Problem) *Assume* (35) *is sufficient for incentive compatibility; no other constraints are required. Then the principal acts to maximize the constrained social welfare function h characterized by*

$$h(\widehat{Y}) = \max_{\Phi} \int_0^{\bar{X}} \lambda e^{-\lambda x}\Phi(x)\left[h(Y(x)) - \frac{c}{\lambda}\right]dx + e^{-\lambda\bar{X}}\Phi(\bar{X})R \tag{37}$$

$$s.t. \quad Y(x) = \frac{\widehat{Y} - b\int_0^x \Phi(s)ds}{\Phi(x)} \tag{38}$$

$$\widehat{Y} \geq b\int_0^{\bar{X}} \Phi(x)dx \tag{39}$$

Further, $\Phi(\bar{X}) \leq \frac{\widehat{Y}}{b\bar{X}}$.

Proof: The principal's objective (37) is a rewriting of (33) using $g(\widehat{Y}) + \widehat{Y} = h(\widehat{Y})$. Since $\widehat{Y}$ is the agent's utility at the beginning of the run, (11) implies (38). Limited liability requires $Y(x) \geq 0$ and implies (39); adding that $\Phi(x)$ is weakly decreasing implies $\Phi(\bar{X}) \leq \frac{\widehat{Y}}{b\bar{X}}$. □

A.4 Optimal Randomization (Proposition 3)

This section shows that the principal optimally randomizes only at $X = 0$ and only the minimally feasible amount: $\Phi(x) = \min\left\{\frac{\widehat{Y}}{b\bar{X}}, 1\right\}$, a constant. To do so, we will go through several preliminary lemmas that follow Lemma A.1 in assuming that (NPS) is sufficient for incentive compatibility; no other constraints are required. We then finish with a proof of Corollary 1.

A.4.1 An Economy Without Setbacks

Lemma A.2 $h(\widehat{Y}) \leq \min\left\{\frac{\widehat{Y}}{b\bar{X}}, 1\right\}\left(R - c\bar{X}\right)$.

Proof: The principal's welfare must be lower than in an economy without setbacks (e.g. with $\lambda = 0$). In such an economy, the agent starts with utility $\widehat{Y}$ and setbacks do not occur. The agent can always receive $b\int_0^{\bar{X}} \Phi(x)dx$ from endeavoring to shirk for time $\bar{X}$, so we must have $\widehat{Y} \geq b\int_0^{\bar{X}} \Phi(x)dx$. The principal's objective is $\Phi(\bar{X})(R - K(\bar{X})) - c\int_0^{\bar{X}} \Phi(x)dx$, where $K(\bar{X})$ is the prize upon successful completion. In an incentive compatible contract, we have payment only on completion, so $\Phi(\bar{X})K(\bar{X}) = \widehat{Y}$. Thus, the principal's problem is

$$\max_{\Phi}\left[\Phi(\bar{X})R - \widehat{Y} - c\int_0^{\bar{X}} \Phi(x)dx\right] \tag{40}$$

$$s.t. \quad \widehat{Y} \geq b\int_0^{\bar{X}} \Phi(x)dx \tag{41}$$

We observe that for any given $\Phi(\bar{X})$, the principal gains by minimizing $\Phi(x)$, which both reduces the expected running cost and slackens the constraint. Thus $\Phi(x) = \Phi(\bar{X})$ and the principal's problem becomes

$$\max_{\Phi(\bar{X})}\left[\Phi(\bar{X})R - \widehat{Y} - c\Phi(\bar{X})\bar{X}\right] \tag{42}$$

$$s.t. \quad \widehat{Y} \geq b\bar{X}\Phi(\bar{X}) \tag{43}$$

If $\widehat{Y} \geq b\bar{X}$, then $\Phi(\bar{X}) = 1$ and $K(\bar{X}) = \widehat{Y}$, and the contract is incentive compatible; the principal's value is $R - c\bar{X} - \widehat{Y}$ and the social surplus is $R - c\bar{X}$. If $\widehat{Y} \leq b\bar{X}$, then $\Phi(\bar{X}) = \frac{\widehat{Y}}{b\bar{X}}$ and $K(\bar{X}) = b\bar{X}$, and the contract is incentive compatible; the principal's value is $\frac{\widehat{Y}}{b\bar{X}}\left(R - c\bar{X}\right) - \widehat{Y}$ and social surplus is $\frac{\widehat{Y}}{b\bar{X}}\left(R - c\bar{X}\right)$. □

A.4.2 Iterating the Bound on Welfare

Lemma A.3 *Define* α^* *such that*

$$e^{-\lambda\bar{X}}R - \int_0^{\bar{X}} \lambda e^{-\lambda x}\left(\frac{c}{\lambda} + \alpha^* bx\right)dx = 0 \tag{44}$$

and consider the sequence

$$\alpha_{n+1} = \alpha_n \int_0^{\bar{X}} \lambda e^{-\lambda x} dx + \frac{1}{b\bar{X}} \max \left[0, e^{-\lambda \bar{X}} R - \int_0^{\bar{X}} \lambda e^{-\lambda x} \left(\alpha_n b x + \frac{c}{\lambda} \right) dx \right] \tag{45}$$

with $\alpha_0 = \frac{1}{b\bar{X}} \left(R - c\bar{X} \right)$*. Then* α_n *is decreasing and* $\alpha_\infty = \frac{\lambda R}{b(e^{\lambda \bar{X}} - 1)} - \frac{c}{b} < \alpha^*$*.*

We have that $h(\widehat{Y}) \leq \alpha_\infty \widehat{Y}$*.*

Proof: We begin with the properties of the sequence. If $\alpha_n \geq \alpha^*$, then the maximand in (45) is zero; if $\alpha_n < \alpha^*$, then the maximand in (45) is $e^{-\lambda \bar{X}} R - \int_0^{\bar{X}} \lambda e^{-\lambda x} \left(\alpha_n b x + \frac{c}{\lambda} \right) dx$. Either way, inspection shows that the coefficient on α_n is between zero and one, the sequence is decreasing, and the limit of the sequence is

$$\lim_{n \to \infty} \alpha_n = \alpha_\infty = \frac{\frac{e^{-\lambda \bar{X}} R}{b\bar{X}} - \int_0^{\bar{X}} \lambda e^{-\lambda x} \frac{c}{\lambda b \bar{X}} dx}{1 - \int_0^{\bar{X}} \lambda e^{-\lambda x} \left(1 - \frac{x}{\bar{X}} \right) dx} = \frac{\lambda R}{b(e^{\lambda \bar{X}} - 1)} - \frac{c}{b} \tag{46}$$

Now we show that (45) characterizes a sequence such that $h(\widehat{Y}) \leq \alpha_n \widehat{Y}$. The claim holds for $n = 0$ from Lemma A.2, and we show that if it holds for n, it must hold for $n + 1$.

Starting with (37), we have

$$h(\widehat{Y}) = \max_{\Phi} e^{-\lambda \bar{X}} \Phi(\bar{X}) R + \int_0^{\bar{X}} \lambda e^{-\lambda x} \Phi(x) \left(h(Y(x)) - \frac{c}{\lambda} \right) dx \tag{47}$$

$$\leq \max_{\Phi} e^{-\lambda \bar{X}} \Phi(\bar{X}) R + \int_0^{\bar{X}} \lambda e^{-\lambda x} \Phi(x) \left(\alpha_n Y(x) - \frac{c}{\lambda} \right) dx$$

$$= \max_{\Phi} e^{-\lambda \bar{X}} \Phi(\bar{X}) R + \int_0^{\bar{X}} \lambda e^{-\lambda x} \left(\alpha_n \widehat{Y} - \alpha_n b \int_0^{x} \Phi(s) ds - \frac{\Phi(x) c}{\lambda} \right) dx$$

$$\leq \max_{\Phi} e^{-\lambda \bar{X}} \Phi(\bar{X}) R + \int_0^{\bar{X}} \lambda e^{-\lambda x} \left(\alpha_n \widehat{Y} - \alpha_n b x \Phi(\bar{X}) - \Phi(\bar{X}) \frac{c}{\lambda} \right) dx$$

$$= \int_0^{\bar{X}} \lambda e^{-\lambda x} \alpha_n \widehat{Y} dx + \max_{\Phi} \Phi(\bar{X}) \left[e^{-\lambda \bar{X}} R - \int_0^{\bar{X}} \lambda e^{-\lambda x} \left(\alpha_n b x + \frac{c}{\lambda} \right) dx \right] \tag{48}$$

Since $0 \leq \Phi(\bar{X}) \leq \frac{\widehat{Y}}{b\bar{X}}$ (Lemma A.1), (48) is maximized at either $\Phi(\bar{X}) = 0$ or $\Phi(\bar{X}) = \frac{\widehat{Y}}{b\bar{X}}$, and we have that $h(\widehat{Y}) \leq \alpha_{n+1} \widehat{Y}$. Since $\alpha_n \downarrow \alpha_\infty$, we have $h(\widehat{Y}) \leq \alpha_\infty \widehat{Y}$.
□

Lemma A.4 *If* $\widehat{Y} \leq b\bar{X}$*, then* $h(\widehat{Y}) = \alpha_\infty \widehat{Y}$ *with* $\Phi(x) = \frac{\widehat{Y}}{b\bar{X}}$*, constant.*

Proof: The contract of Section 5.1 is incentive compatible and attains the given bound with $\Phi(x) = \frac{\widehat{Y}}{b\bar{X}}$ for $\widehat{Y} \leq b\bar{X}$. $\Box$

Lemma A.5 *For all Y and δ, we have $\frac{h(Y)}{Y} \geq \frac{h(Y+\delta)}{Y+\delta}$.*

Proof: Let $\Phi(x)$ be optimal starting from $\widehat{Y} = Y + \delta$, and let $\hat{\Phi}(x) = \Phi(x)\frac{Y}{Y+\delta}$. Examining (38), we see that using $\hat{\Phi}$ starting from $\widehat{Y} = Y$ and using Φ starting from $\widehat{Y} = Y + \delta$ generate the same path of $Y(x)$. Because $\Phi(x)$ is optimal for $\widehat{Y} = Y + \delta$ but $\hat{\Phi}(x)$ could be sub-optimal for $\widehat{Y} = Y$, inspection of (37) shows that $h(Y) \geq h(Y+\delta)\left(\frac{Y}{Y+\delta}\right)$. $\Box$

A.4.3 Randomization

Lemma A.6 *If $\widehat{Y} \geq b\bar{X}$, then $\Phi(x) = 1$ is strictly optimal.*

Proof: $\widehat{Y} \geq b\bar{X}$ implies the feasibility constraint (39) is slack and the principal's maximization problem (37) is

$$h(\widehat{Y}) = \max_{\Phi} \int_0^{\bar{X}} \lambda e^{-\lambda x}\Phi(x)\left[h(Y(x)) - \frac{c}{\lambda}\right] dx + e^{-\lambda\bar{X}}\Phi(\bar{X})R \tag{49}$$

subject to (38). Consider any $\Phi(x)$ with $\Phi(\bar{X}) < 1$. We will show that the gain from increasing $\Phi(x)$ to 1 is strictly positive.

Let $Y(x)$ be defined as in (38) and $Y_\phi(x) = \widehat{Y} - bx$ be the analogous function after replacing Φ with 1. Then $\Phi(x) \leq 1$ implies $Y(x) \geq Y_\phi(x)$. The change in (49) from increasing $\Phi(x)$ to 1 is

$$\begin{aligned} G =& e^{-\lambda\bar{X}}\left(1 - \Phi(\bar{X})\right)R - \int_0^{\bar{X}} \lambda e^{-\lambda x}\left(1 - \Phi(x)\right)\frac{c}{\lambda}dx \\ &+ \int_0^{\bar{X}} \lambda e^{-\lambda x}\left[h(Y_\phi(x)) - \Phi(x)h(Y(x))\right] dx \end{aligned} \tag{50}$$

Looking at the first line in (50), $\Phi(x) \geq \Phi(\bar{X})$ implies

$$-\int_0^{\bar{X}} \lambda e^{-\lambda x}\left(1 - \Phi(x)\right)\frac{c}{\lambda}dx \geq -\left(1 - \Phi(\bar{X})\right)\int_0^{\bar{X}} \lambda e^{-\lambda x}\frac{c}{\lambda}dx \tag{51}$$

Continuing with the second line in (50): Lemma A.5 implies $\frac{h(Y_\phi)}{Y_\phi} \geq \frac{h(Y)}{Y}$. Multiplying by Y_ϕ and adding $-\Phi(x)h(Y)$ to both sides yields $h(Y_\phi) - \Phi(x)h(Y) \geq$

$-\frac{h(Y)}{Y}\left[\Phi(x)Y - Y_\phi\right]$. Thus,

$$\int_0^{\bar{X}} \lambda e^{-\lambda x}\left[h(Y_\phi) - \Phi(x)h(Y)\right] dx \geq -\int_0^{\bar{X}} \lambda e^{-\lambda x}\frac{h(Y)}{Y}\left[\Phi(x)Y - Y_\phi\right] dx \tag{52}$$

Next, we observe that because $\Phi(x) \geq \Phi(\bar{X})$, we have

$$0 < \Phi(x)Y - Y_\phi = bx - b\int_0^x \Phi(t)dt \leq b\left(1 - \Phi(\bar{X})\right)x \tag{53}$$

Substituting (53) into (52) yields

$$\begin{aligned}\int_0^{\bar{X}} \lambda e^{-\lambda x}\left[h(Y_\phi) - \Phi(x)h(Y)\right] dx &\geq -\int_0^{\bar{X}} \lambda e^{-\lambda x}\frac{h(Y)}{Y}\left[b\left(1 - \Phi(\bar{X})\right)x\right] dx \\ &\geq -\left(1 - \Phi(\bar{X})\right)\int_0^{\bar{X}} \lambda e^{-\lambda x}\alpha_\infty bx\, dx\end{aligned} \tag{54}$$

with α_∞ from Lemma A.3. Putting together (50), (51), and (54), we obtain

$$G \geq (1 - \Phi(\bar{X}))\left[e^{-\lambda\bar{X}}R - \int_0^{\bar{X}} \lambda e^{-\lambda x}\left(\frac{c}{\lambda} + \alpha_\infty bx\right) dx\right] \tag{55}$$

Since $\Phi(\bar{X}) < 1$ and the term in square brackets is strictly positive (Lemma A.3), we have $G > 0$. Thus, there is a single optimal policy in (49) which is $\Phi(x) = 1$. □

Lemmas A.4 and A.6 together give the randomization policy in Proposition 3.

Corollary 1 follows from substituting the policy from Proposition 3 into (11) with $x = 0$. Then, (38) implies $Y(0) = \lim_{x\downarrow 0} Y(x) = \frac{\hat{Y}}{\Phi(0)} = W_t$.

A.5 The Time-Budget Contract (Proposition 4)

We show that a time-budget contract is sufficient to induce the agent to take the high action and to request an extension if and only if he has a setback or the initial utility W_0 is less than $b\bar{X}$. In doing so, the agent achieves continuation utility $W(X, S) = bS + \frac{b}{\lambda}\left(e^{\lambda X} - 1\right)$.

We will first solve an altered problem in which the agent's payoffs are higher, creating an upper bound on his continuation utility in the original problem. In the altered problem, in addition to the standard payment on success, the agent receives a payment of $\frac{b}{\lambda}\left(e^{\lambda X} - 1\right)$ for partially completed projects $X \leq \bar{X}$ that are delivered when $S_t = 0$ or when the agent is terminated.[13] In the altered problem, the candidate time-budget contract with initial time $S_0 > 0$ is sufficient to induce the agent to report truthfully and take the high action. In doing so, the agent achieves continuation utility $\bar{W}(X,S) = bS + \frac{b}{\lambda}\left(e^{\lambda X} - 1\right)$.

The agent's HJB equation is

$$\begin{aligned} 0 = \max \Big\{ & \mathrm{E}\bar{W}(X, S+L) - \bar{W}(X,S), \\ & \bar{W}_X a + (1-a)b - \bar{W}_S \\ & + a\lambda \max\left\{\mathrm{E}\bar{W}(0, S+L) - \bar{W}(X,S), \bar{W}(0,S) - \bar{W}(X,S)\right\} \Big\} \end{aligned} \tag{56}$$

where L is the length of an extension (Section 4.4), and with the following boundary conditions:

$$\bar{W}(\bar{X}, S) = bS + \frac{b}{\lambda}\left(e^{\lambda \bar{X}} - 1\right) \tag{57}$$

$$\bar{W}(X, 0) = \frac{b}{\lambda}\left(e^{\lambda X} - 1\right). \tag{58}$$

Equation (56) represents a branched choice. The first line of (56) represents the change of utility from requesting an extension when there is no setback. The second line represents the flow utility from working or shirking, and the third line represents the change of utility from requesting an extension or not when there is a setback. The first max is over whether to announce a false setback or take the change over dt. If there is no false setback, the agent chooses to work or shirk, and the second max is over whether to postpone the report of a setback when one occurs.

The boundary condition (58) is the key difference between the original problem and the altered problem. The proposed continuation utility function $\bar{W}(X,S) = bS + \frac{b}{\lambda}\left(e^{\lambda X} - 1\right)$ is consistent with the boundary conditions.

Outside the randomization region (for $S > \bar{X}$), L equals 0. Thus (56) becomes

$$0 = \max\left\{0, \bar{W}_X a + (1-a)b - \bar{W}_S + a\lambda\left\{\bar{W}(0,S) - \bar{W}(X,S)\right\}\right\} \tag{59}$$

Substituting in $\bar{W}(X,S) = bS + \frac{b}{\lambda}\left(e^{\lambda X} - 1\right)$ yields $0 = \max\{0, 0\}$, so the agent is indifferent between working and shirking and between asking for an extension or not.

Inside the randomization region, L_t equals $\bar{X} - S_{t-}$ with probability $S_{t-}/\bar{X}$ and

[13] The original problem has a fundamental discontinuity: the agent receives the reward for success if $S = 0$ and $X = \bar{X}$ but zero if $S = 0$ and $X < \bar{X}$. The altered problem does not have this discontinuity.

equals $-S_{t-}$ with probability $1 - S_{t-}/\bar{X}$. Thus (56) becomes

$$0 = \max\left\{\left(1 - \frac{S}{\bar{X}}\right)\bar{W}(X,0) + \frac{S}{\bar{X}}\bar{W}(X,\bar{X}) - \bar{W}(X,S), \bar{W}_X a + (1-a)b - \bar{W}_S \right.$$
$$\left. + a\lambda \max\left\{\left(1 - \frac{S}{\bar{X}}\right)\bar{W}(0,0) + \frac{S}{\bar{X}}\bar{W}(0,\bar{X}) - \bar{W}(X,S), \bar{W}(0,S) - \bar{W}(X,S)\right\}\right\}.$$

Substituting in $\bar{W}(X,S) = bS + \frac{b}{\lambda}\left(e^{\lambda X} - 1\right)$ yields

$$0 = \max\left\{0, b\left(e^{\lambda X} - 1\right)(a-1)\right\}, \tag{60}$$

Then (60) implies that the agent prefers working to shirking if $X > 0$ and is indifferent if $X = 0$. The agent is indifferent between announcing a false setback or not and is indifferent between postponing the report of a setback or not when one occurs. Thus, working and truthfully reporting are (weakly) optimal policies. A standard verification argument shows that our solution to the HJB equation is also sufficient.

Next, we observe that, in the original problem, if the agent reports truthfully and takes the high action, then direct calculation shows that the agent never enters the region $S < \bar{X} - X$, never finishes the contract with a partially completed project, and obtains a continuation utility $W(X,S) = bS + \frac{b}{\lambda}\left(e^{\lambda X} - 1\right)$. Since the altered problem has weakly higher payoffs than the original problem, but the continuation utility from the altered problem is attained in the original problem with high action and truthful reporting, then it must be the case that those policies are also optimal in the original problem. □

A.6 The Optimal Contract Resolved (Theorem 1)

We can now complete the proof of Theorem 1: Proposition 2 shows that (NPS) is necessary for a contract to be IC. If (NPS) is also sufficient (no other constraints are required), then optimality implies a law of motion for the agent's utility (Proposition 2) and the optimal randomization scheme (Proposition 3). The law of motion for utility implies a value for K_τ (16). There are no additional degrees of freedom. The time budget contract (Definition 4) satisfies these properties and is incentive compatible (Proposition 4) without any additional constraints. Thus, (NPS) is sufficient, and the time-budget contract attains the principal's maximum value. Since $X_0 = 0$, Proposition 4 shows that $W_0 = bS_0$. □

A.7 Fatal Setbacks (Corollary 3)

Proof Sketch: If the project has suffered a fatal setback, the principal would like the agent to report the project's death so that she can stop paying the running cost. To make the agent willing to report a setback's type, the principal grants the agent the same level of utility for both fatal setback and non-fatal setback reports. Implementing this incentive compatibility constraint in a time-budget contract implies that since a non-fatal setback is met with continuation utility $W(X = 0, S) = bS$ as the project continues, a fatal setback is met with severance pay of bS as the project is canceled. A higher level of severance would induce the agent to report all setbacks as fatal. A lower level would induce the agent to report a non-fatal setback and shirk out the contract, earning bS. $\square$

B The Value of the Project (Propositions 5 and 6)

B.1 A Formula for $\sigma(S)$

First, (24) implies $\sigma(S)$ is bounded because $\pi(S)$ is bounded between 0 and 1. Clearly, $\lim_{S\to\infty} \pi(S) = 1$. Thus, $\lim_{S\to\infty} \sigma(S) = \Delta = \frac{1}{\lambda}\left(e^{\lambda \bar{X}} - 1\right)$.

From our in-text martingale analysis, the optional stopping theorem implies $S_0 + 0 = \mathrm{E}\left[S_\tau + \tau\right]$. Since $S \leq \bar{X}$ implies that the contract can only end with $S_\tau = 0$, (18) gives us $\sigma(S) = S$.

Next, define the constant $\xi = \int_0^{\bar{X}} \lambda e^{-\lambda t} t dt + e^{-\lambda \bar{X}} \bar{X} = \frac{1}{\lambda}\left(1 - e^{-\lambda \bar{X}}\right)$. Then, for any $S \geq \bar{X}$, we have the following recursive formula:

$$\sigma(S) = \int_0^{\bar{X}} \lambda e^{-\lambda t}(t + \sigma(S - t))dt + e^{-\lambda \bar{X}} \bar{X} = \int_0^{\bar{X}} \lambda e^{-\lambda t} \sigma(S - t)dt + \xi \tag{61}$$

where a setback at time t yields an expected time remaining of $\sigma(S-t)$, while success without a setback implies that $\bar{X}$ time was spent. Then, we use a change of variables $\nu = S - t$ and some algebra to find that (61) can be re-written as

$$e^{\lambda S}\sigma(S) = \int_{S-\bar{X}}^{S} \lambda e^{\lambda \nu} \sigma(\nu) d\nu + \xi e^{\lambda S} \tag{62}$$

This implies $\sigma(S)$ is continuous for all S and differentiable in S for all $S > \bar{X}$ from the arguments in Smith (2011). Differentiating and simplifying, we obtain

$$\sigma'(S) = \lambda\xi - \lambda e^{-\lambda \bar{X}} \sigma(S - \bar{X}) \tag{63}$$

This is a delay differential equation (DDE), with initial condition $\sigma(S) = S$ for all

$S \leq \bar{X}$. By the arguments in Smith (2011), the solution exists and is unique.

Furthermore, because $\sigma(S) < \Delta$, we have $\sigma'(S) > \lambda\xi - \lambda e^{-\lambda\bar{X}}\Delta = 0$, and $\sigma(S)$ is increasing. The fact that $\sigma(S)$ is continuous implies $\sigma'(S)$ is also continuous for all $S > \bar{X}$. Since $\sigma(S) = S$ for $S \leq \bar{X}$, we have $\lim_{S\uparrow\bar{X}} \frac{\partial}{\partial S}\sigma(S) = 1$. However,

$$\lim_{S\downarrow\bar{X}} \sigma'(S) = \lambda\xi - \lambda e^{-\lambda\bar{X}}\sigma(0) = \lambda\xi = 1 - e^{-\lambda\bar{X}} < 1 \tag{64}$$

That is, $\sigma(S)$ has one kink at $\bar{X}$. Similarly, for all $S > \bar{X}$ except at $S = 2\bar{X}$,

$$\sigma''(S) = -\lambda e^{-\lambda\bar{X}}\sigma'(S - \bar{X}) < 0, \tag{65}$$

while at $S = 2\bar{X}$,

$$\lim_{S\uparrow 2\bar{X}} \sigma''(S) = -\lambda e^{-\lambda\bar{X}} \lim_{S\uparrow\bar{X}} \sigma'(S) = -\lambda e^{-\lambda\bar{X}} < 0 \tag{66}$$

$$\lim_{S\downarrow 2\bar{X}} \sigma''(S) = -\lambda e^{-\lambda\bar{X}} \lim_{S\downarrow\bar{X}} \sigma'(S) = -\lambda e^{-\lambda\bar{X}} \left(1 - e^{-\lambda\bar{X}}\right) < 0. \tag{67}$$

implying that $\sigma(S)$ is concave and attains a single local/global maximum in S.

B.2 The Value Function

Equation (25) follows directly from the arguments in Section 5.1. $\sigma(S)$ increasing and concave (Section B.1) implies $F(S, 0)$ is concave in S and attains a single local/global maximum. Similarly, $S \leq \bar{X}$ implies $\sigma(S) = S$ and $F(S, 0) = \left(\frac{\lambda R}{e^{\lambda\bar{X}}-1} - c - b\right) S$, which is increasing from Assumption 2.

Following Section B.1, there is at most one kink point at $S = \bar{X}$, where

$$\begin{aligned}
&\lim_{S\downarrow\bar{X}} \frac{\partial}{\partial S} F(S, 0) = \left(\frac{\lambda R}{e^{\lambda\bar{X}} - 1} - c\right)\left(1 - e^{-\lambda\bar{X}}\right) - b \\
&< \lim_{S\uparrow\bar{X}} \frac{\partial}{\partial S} F(S, 0) = \left(\frac{R}{\Delta} - c\right) - b.
\end{aligned} \tag{68}$$

and (29) holds if and only if $F(S, 0)$ is maximized at $S = \bar{X}$.

The value of the prize follows from (16) and the fact that if there is a setback with $W_t \leq b\bar{X}$, it is no longer possible for $S_\tau > 0$ when the project is completed. □

References

Amarasekara, W., B. Perera, and M. Rodrigo (2018). Impact of differing site conditions on construction projects. *Journal of Legal Affairs and Dispute Resolution in Engineering and Construction 10*(3), 04518006.

Bae, S. (2022). New tech helps construction battle extreme weather and labor disruptions. 2022-03-17.

Baldwin, J. R., J. M. Manthei, H. Rothbart, and R. B. Harris (1971). Causes of delay in the construction industry. *Journal of the Construction Division 97*(2), 177–187.

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.

Bloedel, A., R. V. Krishna, and B. Strulovici (2021). Persistent private information revisited. Working paper.

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.

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.

Goodman, B. (2023). A new antibiotic, discovered with artificial intelligence, may defeat a dangerous superbug. *CNN*. 2023-05-26.

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*.

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.

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. (2022). Designing deadlines. *American Economic Review 112*(3), 963–97.

Marchewka, J. T. (2010). The fbi virtual case file: A case study. *Communications of the IIMA 10*(2), 1.

Mayer, S. (2022). Financing breakthroughs under failure risk. *Journal of Financial Economics 144*(3), 807–48.

Myerson, R. B. (2015). Moral hazard in high office and the dynamics of aristocracy. *Econometrica 83*(6), 2083–2126.

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

Protter, P. E. (2005). *Stochastic Integration and Differential Equations*. Number 21 in Stochastic Modelling and Applied Probability. Springer.

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

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

Smith, H. (2011). *An Introduction to Delay Differential Equations with Applications to the Life Sciences*. Springer.

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.

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