Online Appendix

C An Extra Result

It is possible to show the kink in the welfare function from our welfare bounding arguments, without calculating $\sigma(S)$. After Lemma A.3, we have

Lemma C.1 *If $\widehat{Y} \geq b\bar{X}$, then $h(\widehat{Y}) \leq \alpha_\infty b\bar{X}e^{-\lambda\bar{X}} + \alpha_\infty\widehat{Y}(1 - e^{-\lambda\bar{X}})$.*

Proof: Repeating (47)-(48), observing that $\widehat{Y} \geq b\bar{X}$ implies $\Phi(\bar{X}) \leq 1 \leq \frac{\widehat{Y}}{b\bar{X}}$, and substituting $\alpha_\infty =$ for α_n, we have

$$h(\widehat{Y}) \leq \int_0^{\bar{X}} \lambda e^{-\lambda x}\alpha_\infty\widehat{Y}dx + \left[e^{-\lambda\bar{X}}R - \int_0^{\bar{X}} \lambda e^{-\lambda x}\left(\alpha_\infty bx + \frac{c}{\lambda}\right)dx\right] \tag{69}$$

Using $\alpha_\infty = \frac{\lambda R}{b(e^{\lambda\bar{X}}-1)} - \frac{c}{b}$ and re-arranging yields the result.$\square$

Since Lemma A.4 shows $h(\widehat{Y} \leq b\bar{X}) = \alpha_\infty\widehat{Y}$, we have demonstrated the kink. The slope of the bound in Lemma C.1 matches slope of $\sigma(S)$ at $S = \bar{X}$ from (26).

D Relaxing Commitment

In this online appendix we explore an extension of the baseline model that relaxes the commitment of the principal to use explicit probabilities of project cancellation. First we note that a fully renegotiation-proof contract is not feasible in our setting. As discussed at the beginning of Subsection 4.2, termination is required for incentive compatibility. However, neither random nor deterministic termination are renegotiation-proof, as is usual in dynamic contracts that use termination to provide incentives. The project has positive expected value, so the principal would always prefer to grant an extension than to cancel the project and the agent likewise always prefers extension over cancellation.

It seems plausible (perhaps for reputational reasons) that the principal could commit to cancel the project upon the occurrence of a verifiable event; e.g., non-delivery at $S_t = 0$. What seems less plausible is that the principal's randomization procedure is verifiable. Hence, we suppose that randomization is possible but commitment to explicit randomization probabilities is not. In such a setting, either the contract must be canceled deterministically or it must be incentive compatible for *the principal* to randomize between extension and cancellation.

Thus, we seek contracts subject to a *commitment constraint*: when randomizing, the principal must receive the same payoff from extension as from cancellation, which is zero. To do so, we will look for contracts that respect the commitment constraint and are payoff-equivalent to implementing our time-budget contract under that constraint.[14]

It is easy to see that the optimal contract with deterministic cancellation involves paying the agent severance of $W_{t-} = bS_{t-}$ if he reports a setback with $S_t < \bar{X}$, which results in a terminal payoff of $-bS_{t-}$ for the principal. She can, however, do better with a contract under which it is incentive compatible to randomize. Formally, this involves a constraint that the principal's expected payoff when granting an extension is the same as from cancellation, namely 0.

D.1 Two Contracts

We demonstrate two such contracts. The first grants the agent a very large extension. Given the structure of the problem, there is always a positive initial grant of utility and an associated time budget $\bar{S}$ that gives the principal zero utility.[15] We modify the existing time-budget contract (Definition 4) so that an extension to $\bar{S}$ is granted with probability $S_t/\bar{S}$. This modification is incentive compatible for the agent and gives the principal zero utility for both extension and cancellation, so it respects the commitment constraint.

The second contract uses mixed strategies. To construct such a contract we need to relax the concept of incentive compatibility given in Definition 2. Specifically, suppose that a contract is IC if it is optimal for the agent to follow the principal's *recommended action* at each point in time and to honestly report setbacks. Now consider an altered time-budget contract under which the principal recommends $a_t = 1$ unless and until a setback with $S_t < \bar{X}$ is reported. Then, if the agent is granted an extension (i.e., the clock is reset to $S_t = \bar{X}$) the recommended action is for him to randomize between quitting while requesting a severance payment of $b\bar{X}$ and continuing to work. Because the agent is indifferent between these two alternatives, randomization with any probabilities is IC for him. Moreover, the recommended probabilities can be calibrated such that the expected payoff to the principal from granting an extension is 0. This would leave her indifferent between cancellation and extension and hence willing to randomize with the probabilities given in Definition 4:

Proposition D.1 (Non-verifiable Randomization) *The principal can implement a contract that respects the commitment constraint with an altered time budget and the same prize structure K_τ. If $S_t < \bar{X}$ and a setback is reported, then the principal extends the schedule to $\bar{X}$ with probability $\frac{S_{t-}}{\bar{X}}$ and cancels the project with probability*

[14] Notice that our time-budget contract has randomization as soon as the principal knows that randomization will be required, but not earlier. Any delay would result in the principal receiving a negative payoff from the flow cost c before randomizing and receiving zero.

[15] The principal's payoff must be less than $F^{FB} - W_0$ and continuous in W_0. If there is some level of W_0 for which the principal obtains positive value, then there is some higher level such that she obtains zero.

$1-\frac{S_{t-}}{\bar{X}}$. *Upon receiving an extension, the agent randomizes between quitting and requesting severance of* $b\bar{X}$ *with probability* $q=\frac{\hat{F}(\bar{X},0)}{\hat{F}(\bar{X},0)+b\bar{X}}$, *and continuing to work with probability* $1-q$, *where*

$$\hat{F}(S=\bar{X},0)\equiv Re^{-\lambda\bar{X}}-\frac{c+b}{\lambda}\left(1-e^{-\lambda\bar{X}}\right). \tag{70}$$

Proof: This result holds because $\hat{F}(S=\bar{X},0)$ is the principal's payoff when the agent works, so she is indifferent between cancellation and extension iff

$$0=(1-q)\hat{F}(\bar{X},0)-qb\bar{X}. \tag{71}$$

The two contracts grant the same payoff to the principal because they are both incentive compatible for the agent and both grant the principal zero utility after the first setback with $S_t<\bar{X}$. □

D.2 Payoffs and the Initial Time Budget

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

$$\hat{F}(S=\bar{X},0)=F(S=\bar{X},0)\left(\frac{1-e^{-\lambda\bar{X}}}{\lambda\bar{X}}\right). \tag{72}$$

The fraction on the right side of this equation is less than 1 since $\lambda\bar{X}>0$. A setback during the course of a short-leash contract under full commitment still leaves the principal with a positive expected payoff as shown in Proposition 6, whereas a setback in the randomization region absent commitment results in an expected payoff of zero. Perhaps surprisingly, the lack of commitment to randomize leads the principal to grant the agent *more* initial time:

Proposition D.2 (Optimal Initial Time Budget) *Define* S^* *to be the principal's optimal initial time budget when randomization is verifiable, and* $\hat{S}^*$ *to be the principal's optimal initial time budget when randomization is not verifiable (randomization implies the principal's value must be zero). Then,* $S^*\leq\hat{S}^*$.

Proof: Define $\tau_R=\inf\{t>0:X_t=0,S_t<\bar{X}\}$ as the first time a setback occurs that requires randomization and τ to be contract completion, both in economies with verifiable randomization. Define $\hat{\tau}\equiv\min\{\tau,\tau_R\}$, and define $\hat{F}(S,0)$ and $\hat{\sigma}(S)$ analogously as (7) and (18) with $\hat{\tau}$ replacing τ. Then, by the optional stopping theorem, $\mathrm{E}_0\left[W_{\hat{\tau}}\right]=W_0=bS_0$. Because the principal's payoff at $\hat{\tau}$ is either $R-W_{\hat{\tau}}$ or 0, which is identical to her payoff at τ in the baseline model, the same arguments in Section 5.1 applies, which yields:

$$\hat{F}(S,0) = \left(\frac{\lambda R}{e^{\lambda \bar{X}} - 1} - c\right)\hat{\sigma}(S) - bS. \tag{73}$$

Since $\hat{\tau} \leq \tau$, we have $\hat{\sigma}(S) \leq \sigma(S)$. (63) implies that

$$\sigma'(S) - \hat{\sigma}'(S) = \lambda e^{\lambda \bar{X}} \left(\hat{\sigma}(S - \bar{X}) - \sigma(S - \bar{X})\right) \leq 0 \tag{74}$$

Thus, $\hat{F}(S,0) \leq F(S,0)$ and $\hat{F}_S(S,0) \geq F_S(S,0)$. Standard increasing differences logic thus implies that $\hat{S}^* \geq S^*$. □

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

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

E Explicit Probabilities

A nice feature of a short-leash contract is that it allows straightforward calculation of the probabilities of extensions, cancellations, and overruns. Figure 2 plots the probabilities of these events. Define $\mu \equiv \lambda \bar{X}$ to be the expected number of setbacks experienced while the project is in operation. Then,

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

2. $P_{\text{EC}}(\mu)$ is the probability that the project is canceled early, before its initial ex-

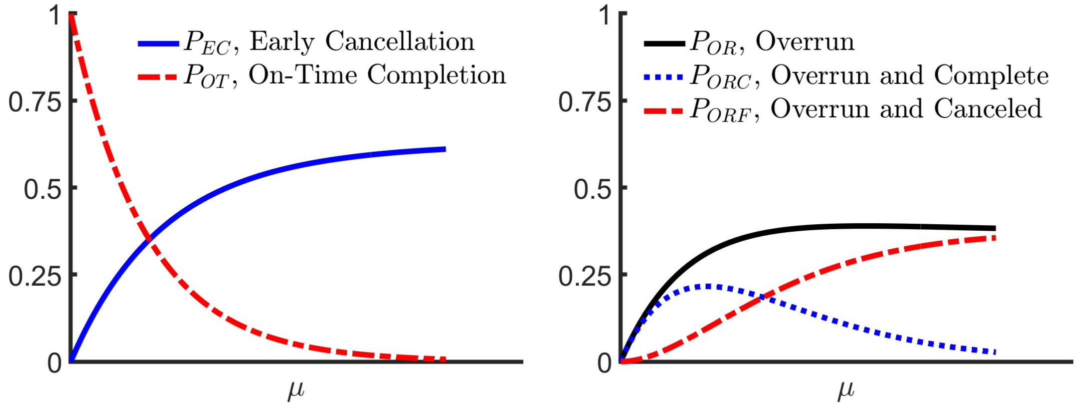

Figure 2: **Probability of an Overrun Under a Short-leash Contract**
The left panel of this figure plots the probability of early cancellation (P_{EC}, blue solid curve) and the probability of on-time completion (P_{OT}, red dash-dot curve). The right panel of this figure plots the probability of an overrun (P_{OR}) $\mu \equiv \lambda \bar{X}$ is the expected number of setbacks experienced while the project is in operation.

pected duration of $\bar{X}$ (left panel, blue solid curve).[16] This function increases from 0 because as the expected number of setbacks rises, it becomes ever more likely that the project will not survive the requisite randomizations before time $\bar{X}$ has elapsed. Indeed, $\lim_{\mu \to \infty} P_{\text{EC}}(\mu) = 1$ because a constant stream of setbacks must result in early project cancellation for any $\bar{X} > 0$.

3. $P_{\text{OR}}(\mu) = 1 - P_{\text{EC}}(\mu) - P_{\text{OT}}(\mu)$ is the probability of an overrun, $\Pr\{\tau > \bar{X}\}$: the probability that the project ends, either from completion or cancellation, after the initial expected duration $\bar{X}$ (right panel, black solid curve). It is low for small values of μ because the project will most likely be completed on time. It rises until achieving a maximum of approximately 0.39 when $\mu = 3.34$ and then decreases as the probability of early cancellation becomes ever more likely. We break this into two sub-possibilities: efficient overruns for which the project is eventually completed P_{ORC}, and inefficient overruns followed by eventual cancellation P_{ORF}.[17]

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

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