Expert Advice and Optimal Project Termination

Erik Madsen*

February 10, 2019

Abstract

I analyze how a firm should elicit advice from an expert on terminating a project with a stochastic lifespan. The firm cannot directly observe the project's lifespan, but imperfectly monitors its state by observing incremental output. The expert directly observes the state of the project, but collects on-the-job benefits and so prefers to prolong operation as much as possible. He possesses no capital and enjoys limited liability, preventing efficient trade even in case the expert has no initial private information. The optimal long-term contract involves a stochastic project deadline and a completion bonus to the expert which declines as the deadline approaches. The deadline is responsive to good and bad runs of output, and exhibits variable output sensitivity over the lifetime of the project, in particular becoming more sensitive the closer the project is to termination. Elicitation of expert advice increases the ex post operational efficiency of the project, but asymmetrically - late terminations are completely resolved, while early terminations are mitigated but not entirely eliminated. These features are robust to extensions in which expert has limited initial capital, can be replaced, or can have his on-the-job benefits dissipated by busywork.

JEL Classification: D82, D86, J33, J41, L14, M52

Keywords: Dynamic contracting, mechanism design, optimal stopping, limited liability, expert advice

*Department of Economics, New York University. Email: emadsen@nyu.edu. An earlier version of this paper circulated under the title "Optimal project termination with an informed agent". I am deeply indebted to my advisers Andrzej Skrzypacz, Robert Wilson, and Sebastian Di Tella for the essential guidance they have contributed to the development of this paper. I am especially grateful to Andrzej Skrzypacz for the time and tireless energy he has devoted at many critical junctures. I also thank Mohammad Akbarpour, Doug Bernheim, Jeremy Bulow, Gabriel Carroll, Jeff Ely, Ben Golub, Brett Green, Johannes Hörner, Nicolas Lambert, Michael Ostrovsky, Takuo Sugaya, Juuso Toikka, and Jeffrey Zwiebel for valuable discussions and feedback.

1 Introduction

Many firms hire external consultants or agencies (henceforth “experts”) to perform open-ended projects for the firm. For instance, management consultants improve operational efficiency and implement industry best practices; ad agencies orchestrate campaigns to advertise a new product or boost an existing line; and investment bankers search for promising acquisition opportunities for a firm looking to grow. The scope of such projects may vary significantly from case to case, and is often initially uncertain to both parties. As the project proceeds, the expert can use past experience and their domain expertise to determine when the project has reached the point of diminishing returns. In the examples just listed, management consultants will recognize when all high-value process improvements have been fully implemented; ad agencies can conduct internal polling to determine when their campaign has saturated the market and how it is being received; and bankers may sense when they’ve exhausted the attractive potential targets in the market.

The firm may naturally wish to rely on the expertise of their expert to size the project correctly, by identifying when all productive work has been exhausted and the project should be wrapped up. Unfortunately, experts often have natural incentives to prolong projects as much as possible. This is particularly true when, as is common in many applications, the expert bills for time spent or is otherwise paid a regular retainer for the duration of the project. Alternatively, the expert may benefit from learning-by-doing or gather valuable data on the firm through continued contact, or may simply be motivated by empire-building concerns. Whatever the source, this incentive is distinct from the more-studied moral hazard concern that an expert may prolong a project by shirking and failing to complete it in a timely manner. Indeed, in many contractor settings the firm may be able to monitor reasonably accurately how much time and manpower the expert is devoting to their project, for instance by keeping their consulting staff on-site and demanding regular in-depth updates on work completed. In such cases the more difficult task is to discern whether resources devoted to the project are producing enough results to justify their cost.

To discipline the incentive to prolong projects, firms often have some ability to monitor “deliverables”, an output signal generated by the project which is less precise than the expert’s insider information, but is correlated with it. For instance, management consultants may be benchmarked on the basis of ongoing cost reductions; ad agencies may be checked by the sales of the products they advertise; and investment banks may be judged on the quality and number of potential deals they bring to the firm. These metrics naturally sag as the project becomes unviable, and by conditioning payments and termination on them the firm can provide a counterforce to the expert's incentive to prolong their project. In this paper I study how firms should contract on deliverables to elicit the expert's private information on optimal project duration and maximize project returns.

I build a model in which a firm employs an expert to run a project which is viable for only a limited time, initially unknown to both firm and expert. While the expert monitors the project state directly over time, the firm observes only a noisy output signal whose increments are correlated with the current state. I allow for a very general specification of project output, nesting settings in which output evolves as a Brownian motion or via Poisson jump processes. The Brownian case corresponds to environments in which profit fluctuates incrementally with each unit of output produced, while Poisson models reflect environments with periodic "breakthroughs" or "accidents". I also permit arbitrary project lifespan distributions. In particular, my model encompasses the case in which the project is surely viable initially but decays at a constant rate, as well as the polar opposite case in which the project is either viable forever or else immediately unviable. The firm and the expert are both risk-neutral with equal discount rates, and the firm can make transfers to the agent as well as commit to shuttering the project. However, the expert enjoys limited liability and cannot be sold the project. Expert incentives to prolong the project are modeled by a stream of unpledgeable flow benefits which accrue so long as the project is active, regardless of its state.

In the absence of an informative signal about the project's state, the firm's optimal contract would be very simple - the firm sets a deterministic public deadline at which the project is surely terminated, and then pays the expert a "completion bonus" to compensate for lost benefits in case the expert advises the project should be shuttered early. The length of the deadline is then chosen to balance the efficiency gains from longer project operation against the incentive payments needed to make truthtelling optimal for the expert. When the firm can additionally learn from the project's output history, the problem becomes more complex. On the one hand, conditioning the deadline on the quality of past output reduces incentive payments by decreasing the project's expected lifespan when it is unviable. On the other hand, this variance increases the probability of early project termination, which due to discounting reduces the firm's expected profits.

An optimal public deadline balances these forces and conditions on the history of project output in a very elegant way. I show that an optimal deadline is a potentially time-dependent threshold rule in the firm's "naive beliefs" about the project state, i.e. the beliefs the firm would have held from monitoring output if it ignored all reports by the expert. If the expert advises that the project be terminated prior to this deadline, the firm does so, and compensates the expert with a completion bonus equal to the rents he would have collected by staying silent forever. If the project reaches the deadline, no incentive payments are made.

The optimal deadline can be implemented via a sequence of forecasted deadlines, which approach at rate 1 on average but fluctuate in response to good or bad runs of output and the arrival of discrete breakthroughs or accidents which are informative about the project's state. Good news extends the forecasted deadline, while bad news hastens its arrival. The forecast exhibits time-varying sensitivity to runs of output, with increasing sensitivity the closer the project is to the deadline. Further, under a regularity condition it exhibits the same sensitivity pattern in response to discrete output events so long as the project is sufficiently far from termination. All of these results are robust to extensions in which the expert has initial capital to contribute to the project, the firm can replace the expert rather than shuttering the project, and the expert can be assigned busywork to reduce his incentive to prolong project operation.

Relative to a benchmark in which the firm is unable to solicit the expert's advice (for instance, if incentive payments were not feasible or could not be committed to), expert advice improves the efficiency of project operation in every state of the world, but asymmetrically. Whenever the project would have been operated past viability without expert advice, it is now operated ex post efficiently. But if the project would have been shuttered too soon, its operational lifespan is prolonged, but not necessarily to the efficient termination point. In terms of comparative statics, the efficiency of the optimal contract is increasing in the informativeness of project output about the state, decreasing in the rate of state transitions, and decreasing in the severity of the expert's incentive misalignment. In particuar, contractability of output increases project efficiency.

My model suggests several important qualitative lessons for overcoming expert incentives to prolong open-ended projects. First, it has been understood at least since Lazear (1983) that variable severance pay, whose NPV decreases over time, can function as a tool to induce efficient separation by workers who privately observe changes to match quality. I show that this tool works best for the firm when paired with deadlines, which economize on the termination bonuses needed to avoid inefficient project prolongment (at the cost of introducing inefficient early separations). Second, when observable signals of project performance are available, the contract can be made simultaneously more profitable and more efficient by conditioning on these signals. Third, the sensitivity of contractual outcomes to current project performance should be modulated by aggregate project performance over its history. When aggregate performance has been good, incentives should be low-powered, i.e. relatively unresponsive to current output movements. By contrast, when aggregate performance has been poor, incentives should power up.

1.1 Related literature

This paper solves a dynamic mechanism design problem featuring limited liability, lack of single crossing in the agent's payoff function, and imperfect public monitoring of the agent's type. These features collectively depart significantly from the assumptions of most existing models.

A large class of models assume fully flexible transfers and agent marginal valuations for allocations which are increasing in type, commonly known as the single crossing condition. Papers in this tradition can be thought of as extending the canonical static model of Myerson (1981) to multi-period settings, though they often also allow for more general agent preferences and types which enter the principal's objective function. Important entries include Baron and Besanko (1984), Courty and Li (2000), Besanko (1985), Battaglini (2005), Pavan, Segal, and Toikka (2014), and Williams (2011). In these papers the agent's private information is elicited by substituting away from monetary transfers and toward current and future allocations as reported type increases.[1] By contrast, in my model the expert's type is payoff-relevant only to the firm, and so tradeoffs between payments and allocations cannot separate different types. Instead, the firm observes public signals correlated with type and uses them to tie the expert's payoff to his type. In common between my setting and the papers above, optimal contracting boils down to a tradeoff between allocational efficiency and the payment of information rents. However, as emphasized in Eso and Szentes (2017), under fully flexible transfers the agent typically receives no information rents for any (orthogonalized) private information received after time zero. (This is true regardless of the allocation implemented.) Therefore the nature of the information rents is quite different in the two settings, as my problem features information rents even in case the expert possesses no time-zero private information.[2]

Garrett and Pavan (2012) replace agent-payoff-relevant types with imperfect public monitoring of the type. In their model observable output is the sum of type and random noise as

[1] A related paper, Kruse and Strack (2015), departs from the typical revelation contract framework by restricting attention to contracts which delegate a decision to halt to the agent and receive no other communication. While this restricts the set of implementable allocations, a single-crossing condition leads to the usual tradeoff between transfers and allocations (i.e. project lifespans).

[2] My model allows for the possibility that the project is initially nonviable with positive probability, in which case the expert possesses private information at time zero. However, the basic structure of an optimal contract is the same whether or not the expert possesses such information.

well as unobserved effort, building in a career concerns dynamic. The principal's allocation decision is worker-task matching, as he can replace the worker. As in my model, worker-firm match value is ephemeral and the analysis focuses on characterizing the firm's optimal termination policy. Also in common with my model, type is payoff-relevant only to the principal, but payments can be linked to output to separate types.[3] Unlike my model, the agent is not protected by limited liability and so extract rents solely via time-zero private information. This distinction, along with the presence of moral hazard, leads to starkly different optimal termination dynamics. In particular, the decision to terminate is completely independent of the history of output and can eventually become unresponsive to reported bad news from the agent.

Grenadier, A. Malenko, and N. Malenko (2016) study an assisted optimal stopping problem, in which a principal solicits advice from an expert with knowledge of the first-best exercise time of a real option. The agent is biased toward either early or late exercise, but by a fixed amount inducing a linkage between the agent's preferred allocation and their type analogous to single crossing. Unlike my paper, no transfers are allowed, and the real option problem abstracts from any dynamics which affect incentives - the principal receives no signal of the agent's type and there is no learning about a commonly unknown payoff-relevant state. As a result the problem has a very simple, essentially static structure under commitment, and the bulk of their analysis focuses on the case where the principal cannot commit to a stopping rule.

Guo (2016) studies a bandit experimentation model with an agent who has private information about the quality of the risky arm and a bias toward experimentation. The principal can commit to an experimentation policy following an initial report by the agent, and as the outcome of experimentation is correlated with the agent's private information the principal can generate public signals about the agent's type. However, transfers are absent, and so the main force shaping an optimal policy is single crossing between the agent's type and the length of her preferred experimentation period. My paper abstracts from this force and focuses instead on the optimal linkage of transfers and public signals.

Finally, Varas (2017) and Green and Taylor (2016) study long-term contracting problems with transfers, limited liability, and changing types, all features in common with my model. In Varas (2017) the state is payoff-relevant only to the principal and can be imperfectly monitored via a public output process, as in my model. Unlike my model, the output signal is observed *after* the project is completed, rather than before, and the focus is on

[3] In particular, the linkage of output to payments induces a single crossing condition in type and disutility of effort, as demonstrated in Proposition 4 of that paper.

the limits to deferred compensation and resolution of moral hazard when the agent has a higher discount rate than the principal. Meanwhile, in Green and Taylor (2016) the privately observed project state is not payoff relevant to either party, and its reporting is instead used to partially alleviate an underlying moral hazard problem.

2 The model

2.1 The technology

A firm operates a project with a limited but uncertain scope. The project delivers average output $r_G > 0$ per unit time in continuous time up to a catastrophic failure time $\Lambda \in \mathbb{R}_+ \cup \{\infty\}$, after which it produces average output $r_B < 0$ per unit time. I will refer to Λ as the project's *lifespan*, with θ the associated project *state process*, where $\theta_t = G$ if $t < \Lambda$ and $\theta_t = B$ otherwise.

The firm must decide when to terminate the project, with the goal of maximizing expected discounted project output. However, it is uncertain ex ante about the value of Λ and cannot directly observe the project's state due to random output variability. Specifically, the firm believes ex ante that $\Lambda \sim H(\cdot)$, where H is an arbitrary distribution function over $\mathbb{R}_+ \cup \{\infty\}$. It learns about Λ by observing the process Y, where Y_t is the project's cumulative output up to time $t \in \mathbb{R}_+$. I assume that

$$Y_t = Y^G_{t\wedge\Lambda} + (Y^B_t - Y^B_{t\wedge\Lambda}),$$

where Y^G and Y^B are mutually independent stochastic processes with stationary, independent increments. In other words, Y evolves according to Y^G as long as the project is good (i.e. in the Good state), and evolves according to Y^B when it is bad (in the Bad state). The increments of Y^θ have mean r_θ per unit time, with random variability independent of the project lifespan and past noise realizations.

I assume that each Y^θ is decomposable as the sum of a Brownian motion with drift and a finite set of scaled, compensated Poisson counting processes:[4]

$$Y^\theta_t = r_\theta t + \sigma Z^\theta_t + \sum_{i=1}^n d_i(N^\theta_i(t) - \lambda^\theta_i t),$$

[4] This assumption amounts to a mild restriction on the set of all possible stationary, independent increment processes. (The set of all such processes is characterized by the the Lévy-Itô decomposition.) Essentially, I allow for only a finite arrival rate of jumps and a finite set of jump sizes. These assumptions substantially simplify the analysis of inferring the state from output, without losing much economic realism.

with each Z^θ a standard Brownian motion, each N_i^θ a Poisson processes with rate $\lambda_i^\theta \geq 0$, and $Z^\theta, N_1^\theta, ..., N_n^\theta$ mutually independent. I let $\mathcal{D} = \{d_1, ..., d_n\}$ represent the set of possible jump sizes exhibited by the output process, with each $d_i \in \mathbb{R} \setminus \{0\}$. The Brownian term in each state exhibits a common volatility $\sigma \geq 0$, which is without loss given that the state would otherwise be immediately detectable.

This flexible specification nests common signal processes used in economic problems with dynamic learning, in particular the Brownian and Poisson good and bad news frameworks. Brownian learning corresponds to $Y_t^\theta = r_\theta t + \sigma Z_t^\theta$, with no jump terms. The Poisson good news setting models an output process which consumes a stream of input costs and occasionally generates breakthroughs yielding a fixed windfall. The corresponding output processes are $Y_t^\theta = RN_t^\theta - ct$ with $\lambda^G > \lambda^B$. Here $R > 0$ is the return to a breakthrough and $c > 0$ is the flow cost of operation; they are chosen so that $r_\theta = R\lambda^\theta - c$. Meanwhile the Poisson bad news setting captures a project which generates a steady flow of profits, occasionally punctuated by costly accidents incurring a set cost. The output processes in this case are $Y^\theta = rt - DN_t^\theta$ with $\lambda^B > \lambda^G$. Here $r > 0$ is the profit flow from operation, while $D > 0$ is the cost of an accident; they are chosen so that $r_\theta = r - D\lambda^\theta$.

To streamline statements of results, I do not formally allow for discrete-time production processes. Such a process would correspond to $Y_t^\theta = \sum_{n=0}^{\lfloor t \rfloor} X_n^\theta$ for a sequence of iid random variables $\{X_n^\theta\}_{n=0}^\infty$, which does not have stationary increments in continuous time. The major qualitative results of this paper would go through in discrete time; however, the characterization of the optimal contract's response to output surprises would be significantly more cumbersome.

Formally, I model the exogenous uncertainty of this setting by a complete probability space $(\Omega, \mathcal{F}, \mathbb{P})$ supporting Λ, Y, and a public randomization device independent of both.[5] For each $T \in \mathbb{R}_+ \cup \{\infty\}$, let $\mathbb{P}^T$ be the probability measure under which $\mathbb{P}^T\{\Lambda = T\} = 1$ and Y is identical in law to $Y_{t\wedge T}^G + (Y_t^B - Y_{t\wedge T}^B)$. The measure $\mathbb{P}$ is defined in terms of the $\mathbb{P}^T$ as

$$\mathbb{P} = H(0)\mathbb{P}^0 + \int_0^\infty \mathbb{P}^T \, dH(T) + (1 - H(\infty))\mathbb{P}^\infty.$$

I also define a pair of auxiliary measures which will prove very useful for constructing optimal contracts. Let $\mathbb{P}^G$ denote the measure under which Λ has distribution H and Y is identical in law to Y^G and independent of Λ; and similarly let $\mathbb{P}^B$ denote the measure under which Λ has distribution H and Y is identical in law to Y^B and independent of Λ. These measures

[5]Formally, the public randomization device can be modeled by an adding states to Ω and enlarging the natural filtration $\mathbb{F}^0$ generated by Y and $\mathbf{1}\{\Lambda \leq t\}$, such that the enlarged filtration $\mathbb{F}$ is a standard extension of $\mathbb{F}^0$ under $\mathbb{P}$. See Kallenberg (1997), pg. 298, for details.

induce the same marginal distribution over Λ as $\mathbb{P}$, but assign probabilities to output paths as if the state were "always Good" or "always Bad". (All measures leave the distribution of the randomization device and its independence of Y and Λ unchanged.)

I let $\mathbb{F}^Y = \{\mathcal{F}_t^Y\}_{t\geq 0}$ denote the $\mathbb{P}$-augmented filtration of $\mathcal{F}$ generated by Y and the randomization device; this filtration captures the information available to the firm from its observation of past output. I write $\mathbb{E}_t^Y$ for the conditional expectation under $\mathbb{P}$ with respect to $\mathcal{F}_t^Y$. I also let $\mathbb{F} = \{\mathcal{F}_t\}_{t\geq 0}$ denote the $\mathbb{P}$-augmented filtration of $\mathcal{F}$ generated by both Y and θ as well as the randomization device, with $\mathbb{E}_t$ the conditional expectation under $\mathbb{P}$ with respect to $\mathcal{F}_t$.

The firm is a risk-neutral expected-profit maximizer with discount rate ρ. Supposing the firm operates the project until some ($\mathbb{F}^Y$-stopping) time τ^Y, it receives expected profits

$$\Pi = \mathbb{E}\left[\int_0^{\tau^Y} e^{-\rho t}\,dY\right] = \mathbb{E}\left[\int_0^{\tau^Y} e^{-\rho t}(\pi_t r_G + (1-\pi_t) r_B)\,dt\right], \tag{1}$$

where $\pi_t = \mathbb{P}_t^Y\{\Lambda > t\} = \mathbb{E}_t^Y[\mathbf{1}\{\Lambda > t\}]$ are the firm's posterior beliefs at time t that the project's lifetime has not yet been exceeded, based on its observation of past output.

2.2 The agency problem

The firm employs an expert to oversee the project and monitor its state. The expert costlessly and privately observes the state process θ in addition to the public output process Y. The filtration $\mathbb{F}$ therefore captures the flow of information to the expert. Note that the expert does not directly observe Λ; in particular, he possesses no ex ante private information about the project's lifespan except whether the project is initially viable. Rather, the expert observes everything the firm does, plus precisely enough additional information to terminate the project efficiently.

The expert may make reports on the state of the project to the firm to aid in decision-making. However, he faces an agency problem discouraging honest communication: he enjoys intrinsic benefits from employment, in the form of flow benefits $b > 0$ per unit time accrued while the project operates, regardless of its state.[6] I assume that $r_B + b < 0$, so that it is jointly unprofitable for the project to be operated in the Bad state.[7] I also assume that

[6] None of the results of this paper would impacted if the expert's flow benefits were made state-contingent so long as $b_G \geq b_B > 0$. This is because, whenever $b_G \geq b_B$, all incentive constraints for truthtelling when the project is Good are slack in an optimal contract. Thus the optimal contract is independent of b_G.

[7] When $r_B + b \geq 0$, the optimal contract under limited liability is uninteresting: the firm makes no incentive payments and does not condition termination decisions on the expert's reports. The logic behind

the flow benefits are unpledgeable, either because they are non-pecuniary or because they are collected only after the expert separates from the firm (and cannot be borrowed against during the project).

Finally, I assume that the firm has no technology for imposing non-pecuniary costs on the expert. For instance, the firm cannot assign unpleasant busywork to offset the expert's flow benefits. This restriction best approximates environments in which the expert performs additional non-monitoring tasks requiring attention and mental acuity that are significantly degraded by busywork. In particular, if imposition of a dollar's worth of flow costs on the expert degrades the quality of other work performed by the expert by more than a dollar's worth of output, nonpecuniary costs will never be imposed by the firm. In Section 6.4 I relax this assumption and show how the optimal contract changes when the firm has access to a convex-cost technology that reduces the expert's flow benefits.

The expert is endowed with no initial wealth, has no access to capital markets to borrow, and is protected by limited liability, so cannot be sold the project. (In Section 6.2 I show how the optimal contract changes when the expert can contribute capital to the project.) He is risk-neutral and possesses the same discount rate ρ as the firm.

2.3 Contracts

The firm commits to a long-term contract eliciting reports from the expert over time and specifying a termination date τ and a path of cumulative payments Φ, which may condition on the public history of output and the expert's reports as well as the public randomization device. I thus allow for randomized contracts in my framework. In accordance with the revelation principle, I restrict attention to contracts eliciting a sequence of reports $\theta'_t \in \{G, B\}$ of the current project state at each time t. Given that the state switches only once irreversibly, a revelation contract equivalently elicits a single report at time Λ.

Formally, I assume that the expert makes a report Λ' from the set of $\mathbb{F}$-stopping times.

Definition 1. *A* reporting policy Λ' *is an $\mathbb{F}$-stopping time. The associated* reported state process θ' *is the process defined by $\theta'_t = G$ if $t < \Lambda'$ and $\theta'_t = B$ otherwise.*

Under a revelation contract the firm observes both Y and θ', and commits to a payment process Φ and a termination stopping time τ adapted to the natural filtration $\mathbb{F}'$ generated by Y and θ'.[8]

this result is simple: when $r_B + b \geq 0$ there is scope for gains from trade by operating the project in the Bad state, but because the expert has no ability to pay the firm none of these gains can be realized.

[8]This construction is somewhat informal, as Φ and τ are not well-defined processes on the exogenous

Definition 2. *A* revelation contract $\mathcal{C} = (\Phi, \tau)$ *is a stochastic process* $\Phi \geq 0$ *and a stopping time* τ*, both adapted to* $\mathbb{F}'$*, such that* Φ *is right-continuous, increasing, and satisfies* $\Phi_t = \Phi_\tau$ *for all* $t > \tau$*.*

Limited liability corresponds to the requirement that Φ be positive and increasing.[9] To simplify formulae, I assume that a revelation contract makes no transfers subsequent to termination. This is without loss of generality, since both parties have linear consumption utility with equal discount rates and no information arrives after termination.

The expected payoff to an expert under a revelation contract (Φ, τ) and reporting policy Λ' is

$$\mathbb{E}^{\Lambda'}\left[\int_0^\tau e^{-\rho t}\,(b\,dt + d\Phi_t)\right],$$

where $\mathbb{E}^{\Lambda'}$ averages over output and reported state paths conditional on the reporting policy Λ'. Incentive-compatibility is then defined in the natural way:

Definition 3. *A revelation contract* (Φ, τ) *is* incentive-compatible*, or an* IC contract*, if*

$$\mathbb{E}^{\Lambda}\left[\int_0^\tau e^{-\rho t}\,(b\,dt + d\Phi_t)\right] \geq \mathbb{E}^{\Lambda'}\left[\int_0^\tau e^{-\rho t}\,(b\,dt + d\Phi_t)\right]$$

for all reporting policies Λ'*.*

The firm's profits under any IC contract $\mathcal{C} = (\Phi, \tau)$ are

$$\Pi[\mathcal{C}] = \mathbb{E}^{\Lambda}\left[\int_0^\tau e^{-\rho t}(dY_t - d\Phi_t)\right].$$

The firm's problem is to maximize $\Pi[\cdot]$ over all IC contracts. I refer to any contract achieving this maximum as an *optimal contract*. Implicit in this formulation of the problem is the assumption that the firm requires the expert to operate the project in addition to monitoring it. Therefore the firm cannot terminate the expert without also ceasing operation of the project.[10]

---

probability space, but are properly families of processes indexed by the choice of Λ'. As this technicality does not impact the developments in the body of the paper, I leave a formal discussion of the details to Appendix A.

[9] In fact, the optimal contract would be unchanged if I allowed any $\Phi \geq 0$, since the firm optimally defers all compensation until termination.

[10] In Section 6.3 I analyze alternative settings in which the firm can replace the expert at a cost or continue operating the project on its own. The major qualitative features of an optimal contract remain unchanged, though unsurprisingly the firm chooses a more aggressive termination policy given its improved outside option.

3 Deriving an optimal contract

3.1 Relaxing the incentive constraints

Incentive compatibility amounts to the requirement that the expert not benefit from falsely reporting a state change either before or after Λ, regardless of when the state switches or what run of output occurs. It will be very helpful to characterize incentive compatibility as the concatenation of two sets of constraints which separately rule out deviations to early and late reporting.

Definition 4. *A revelation contract* (Φ, τ) *satisfies* IC-G *(respectively,* IC-B*) if*

$$\mathbb{E}^{\Lambda}\left[\int_0^{\tau} e^{-\rho t}\,(b\,dt + d\Phi_t)\right] \geq \mathbb{E}^{\Lambda'}\left[\int_0^{\tau} e^{-\rho t}\,(b\,dt + d\Phi_t)\right]$$

for all reporting policies $\Lambda' \leq \Lambda$ *(respectively,* $\Lambda' \geq \Lambda$*).*

A contract satisfying IC-G (IC-B) ensures that no reporting policy which always reports sooner (later) than Λ is preferable to truthful reporting. The IC-G and IC-B constraints collectively represent only a subset of the constraints required for incentive-compatibility, since an IC contract must also deter mixed misreporting policies that sometimes report early and sometimes late. The following lemma verifies that IC-G and IC-B together nonetheless rule out all such deviations and ensure incentive-compatibility.

Lemma 1. *A revelation contract is incentive-compatible iff it satisfies both IC-G and IC-B.*

This partition of the set of IC constraints turns out to separate the constraints which bind at the optimum from those which lie slack. In particular, the set of IC-G constraints will be slack under an optimal contract and can be dropped from the optimization problem. Intuitively, given the expert's preference to delay project termination absent contractual incentives, at least some of the IC-B constraints must bind. But since provisioning incentives to report on-time is costly, the firm should impose them as lightly as possible. It would therefore be surprising if the firm provided such strong incentives for reporting a switch on-time that IC-G were violated and the expert profited by reporting a switch prematurely. To validate this conjecture, I solve the relaxed problem of characterizing an optimal IC-B contract. I then verify afterward that the resulting contract is indeed incentive-compatible and therefore an optimal contract with respect to the full suite of incentive constraints.

3.2 Optimal usage of the expert's report

Designing an optimal contract entails determining how payments and allocations should respond both to runs of output and the timing of the expert's report. In Section 4.1 I showed that when the contract does not condition on output, the cost-minimizing IC contract implementing any deadline terminates as soon as the expert reports a state switch and defers all payments until the time of a report. I show now that this result generalizes to arbitrary IC-B contracts which may condition on the history of output. Thus the design of the contract's output dependence can be cleanly separated from its response to the expert's report.

Lemma 2 (No late termination)**.** *Suppose $\mathcal{C} = (\Phi, \tau)$ is an IC-B contract. Then there exists an IC-B contract $\mathcal{C}' = (\Phi', \tau')$ such that:*

- $\tau' = \tau \wedge \Lambda'$;
- $\Pi[\mathcal{C}'] \geq \Pi[\mathcal{C}]$;
- $\Pi[\mathcal{C}'] > \Pi[\mathcal{C}]$ *if* $\mathbb{P}^{\Lambda}\{\tau > \Lambda\} > 0$.

This lemma establishes that optimal IC-B contracts never terminate inefficiently late - that is, after the expert reports the state has switched. In principle late termination could be desirable as a way to compensate the expert with flow benefits for truthful reporting. However, the assumption that $|r_B| > b$ means that the firm can always compensate the expert more cheaply with a monetary transfer at the time of the state switch. The proof of Lemma 2 exploits this observation, modifying a given contract by halting at the time of a report and adding an additional transfer equal to the expert's expected flow benefits plus future transfers under the original contract. This change preserves IC-B while improving profitability if the original contract continued operations in the Bad state with positive probability.

Remark. *There exist IC contracts for which the modified contract constructed in the proof of Lemma 2 is not incentive-compatible.*

To see this, consider a contract which provides large termination payments early in the contract but attenuates these payments quickly as the project proceeds. Such setups create an incentive for the expert to report a state switch early in order to maximize his termination payment, endangering IC-G. Incentive-compatibility can be enforced by maintaining operations following a report in order to monitor output, punishing the expert for a misreport by reducing the termination payment following good runs of output. If such a verification phase were removed by truncating project operation, the transformed contract would violate IC-G even if the initial contract is IC. It is therefore critical in Lemma 2 that the transformed contract is allowed to be merely IC-B and not fully incentive-compatible. This caveat illustrates the tractability brought by passing to the relaxed problem, as well as the importance of verifying that an optimal IC-B contract does not violate IC-G.

Lemma 3 (Backloading)**.** *Suppose* $\mathcal{C} = (\Phi, \tau)$ *is an IC-B contract satisfying* $\tau \leq \Lambda'$ *and* $\mathbb{P}^{\Lambda}\{\tau < \infty\} = 1$. *Then there exists an* $\mathbb{F}^Y$*-adapted process* $F \geq 0$*, inducing payment process* $\Phi'_t = F_\tau \mathbf{1}\{t \geq \tau\}$*, and an* $\mathbb{F}^Y$*-stopping time* τ^Y *such that* $\mathcal{C}' = (\Phi', \tau^Y \wedge \Lambda')$ *is an IC-B contract satisfying* $\Pi[\mathcal{C}'] \geq \Pi[\mathcal{C}]$.

This lemma affords several simplifications to the structure of contracts featuring no late termination. First, any termination rule halting no later than the time of a report must base termination only on public information whenever the project is terminated early. Hence an optimal termination policy can always be formulated as a rule of the form "Terminate the project as soon as the expert reports a state switch or τ^Y has been reached, whichever comes first," for some *public deadline* τ^Y which conditions only on the history of output.

Second, all payments can be backloaded to a single *completion bonus*, denoted F, paid when the project is shuttered. This result is straightforward - both parties are risk-neutral and share the same discount rate, so profits and incentive-compatibility are undisturbed by deferring all promised payments until termination and accruing interest on them at rate ρ.

Finally, the size of the bonus need not be conditioned on the date of a past report, hence F is $\mathbb{F}^Y$-adapted. Establishing this result requires care, since $\mathcal{C}$ may pay the expert differently depending on whether he reports a state switch "just in time" when τ^Y arrives. The proof of the lemma shows that in this case $\mathcal{C}$ may always be modified to pay the lower of the bonuses promised depending on whether or not the expert reports a switch at τ^Y, without disturbing incentive-compatibility or firm profits.

In light of Lemmas 2 and 3, I restrict attention to IC-B contracts (Φ, τ) which may be written $\tau = \tau^Y \wedge \Lambda'$ and $\Phi_t = F_\tau \mathbf{1}\{t \geq \tau\}$ for some $\mathbb{F}^Y$-adapted process $F \geq 0$ and $\mathbb{F}^Y$-stopping time τ^Y.[11] Such contracts are summarizable by the pair (F, τ^Y). Crucially, both

[11]If $\Lambda < \infty$ a.s., this restriction is without loss of generality. Otherwise, it excludes contracts which operate the project forever with positive probability and disburse payments prior to termination in such histories. This is because under such contracts, there is no terminal date at which to backload payments in some histories.

Still, the profit of any such contract can be approximated arbitrarily closely by a sequence of IC-B contracts with bounded termination dates. Optimality of a contract which never terminates and gives interim payments would then manifest via the supremum of contractual profits being unattainable by backloaded contracts. As I shall show, the supremum is achievable and so the restriction is innocuous.

F and τ^Y condition only on the public output history and not the reports of the expert. I therefore pass from the general problem of designing the contract's dependence on both reports and the history of output, to the simpler problem of designing just its dependence on output.

3.3 Optimal implementation of public deadlines

I next solve the problem of how to optimally implement an arbitrary public deadline τ^Y via an IC-B contract. This step reduces the contractual design problem from the simultaneous choice of both F and τ^Y to the choice of τ^Y only.

Definition 5. *An $\mathbb{F}^Y$-stopping time τ^Y is* implementable *if there exists a completion bonus process F such that (F, τ^Y) is an IC-B contract. In this case, F* implements *τ^Y.*

Remark. *Every τ^Y is implementable via the bonus schedule $F_t = b/\rho$.*

This remark follows from the fact that if the firm fixes $F_t = b/\rho$ for all time, then the expert's total profits are the same under any reporting strategy, ensuring he has no incentives to delay reporting. (In fact, this argument shows that there exists a fully IC contract implementing any τ^Y.) Given this positive result, it is meaningful to search for a profit-maximizing implementation of an arbitrary τ^Y.

I derive the optimal implementation by solving a (further) relaxed problem which isolates the binding subset of IC-B constraints. Suppose that the expert's reporting strategy is constrained to satisfy $\Lambda'(\omega) \in \{\Lambda(\omega), \infty\}$ for all $\omega \in \Omega$. In other words, the expert can either report a state switch immediately or not at all, but cannot deviate in any other way. This is equivalent to a model in which the expert's reports are verifiable, as the firm could then costlessly deter any false reports by mandating immediate termination with no payments.

Definition 6. *A contract $\mathcal{C} = (F, \tau^Y)$ satisfies* IC-∞ *if*

$$\mathbb{E}\left[\int_0^{\tau^Y \wedge \Lambda} e^{-\rho t} b\, dt + e^{-\rho(\tau^Y \wedge \Lambda)} F_{\tau^Y \wedge \Lambda}\right] \geq \mathbb{E}\left[\int_0^{\tau^Y \wedge \Lambda'} e^{-\rho t} b\, dt + e^{-\rho(\tau^Y \wedge \Lambda')} F_{\tau^Y \wedge \Lambda'}\right]$$

for all reporting policies Λ' such that $\Lambda'(\omega) \in \{\Lambda(\omega), \infty\}$ for all $\omega \in \Omega$.

Clearly all IC-B contracts satisfy IC-∞, but not vice versa. I solve the relaxed problem of maximizing profits while implementing τ^Y subject to IC-∞, and then show that the resulting contract satisfies IC-B as well.

To state the result, I define a new conditional expectation operator $\mathbb{E}_t^B$ which averages over uncertainty in future output "assuming the state has already switched."

Definition 7. *For any random variable X, let $\mathbb{E}^B[X] = \int X \, d\mathbb{P}^B$. For each $t \in \mathbb{R}_+$, let $\mathbb{E}^B_t[X]$ be the expectation of X under $\mathbb{P}^B$ conditional on $\mathcal{F}_t$.*

Informally, one can think of $\mathbb{E}^B_t$ as satisfying $\mathbb{E}^B_t[X] = \mathbb{E}[X \mid (Y_s)_{s \le t}, \Lambda \le t]$. The latter expression, however, is not a well-defined random variable, and isn't meaningful if $\mathbb{P}\{\Lambda \le t\} = 0$. Definition 7 resolves these issues through a more careful construction.

The following remarks highlight several simple properties satisfied by this conditional expectation operator.

Remark. *If X is $\mathcal{F}^Y_\infty$-measurable, then $\mathbb{E}^B_t[X]$ is $\mathcal{F}^Y_t$-measurable for each t.*

In general, because $\mathbb{E}^B_t[X]$ may condition on the history of the indicator variable $\mathbf{1}\{\Lambda \le t\}$, it is not measurable with respect to just the history of output. However, when X is a function only of the path of output and the randomization device, then under $\mathbb{P}^B$ its distribution is independent of the value of Λ. Thus its expectation conditional on $\mathcal{F}_t$ is the same as conditional on $\mathcal{F}^Y_t$.

Remark. *Suppose X is a stochastic process and Λ' is an $\mathbb{F}$-stopping time satisfying $\Lambda' \ge \Lambda$. Then $\mathbb{E}^B_{\Lambda'}[X_{\Lambda'}] = \mathbb{E}_{\Lambda'}[X_{\Lambda'}]$ a.s.*

This remark simply reflects the fact that, after Λ, the conditional distribution over future output under $\mathbb{P}$ is the same as under $\mathbb{P}^B$.

With the operator $\mathbb{E}^B_t$ in hand, I can characterize the optimal IC-∞ contract implementing an arbitrary τ^Y:

Lemma 4. *For any $\mathbb{F}^Y$-stopping time τ^Y, define a bonus process F^* via*

$$F^*_t = \mathbb{E}^B_t \left[\int_{t \wedge \tau^Y}^{\tau^Y} e^{-\rho(s-t)} b \, ds \right].$$

Then F^* is an $\mathbb{F}^Y$-adapted process, (F^*, τ^Y) satisfies IC-∞, and $F^*_{\Lambda \wedge \tau^Y} \le F_{\Lambda \wedge \tau^Y}$ a.s. for every F such that (F, τ^Y) satisfies IC-∞. In particular, (F^*, τ^Y) maximizes expected profits among all IC-∞ contracts (F, τ^Y).

This result is proven in the following way. Suppose that τ^Y is implemented via IC-∞ bonus process F. If the state switches at time $t < \tau^Y$, the expert has two options - report immediately or withhold his report forever. The expert receives F_t from reporting immediately, versus $\mathbb{E}^B_t \left[\int_t^{\tau^Y} e^{-\rho(s-t)} b \, ds + e^{-\rho(\tau^Y - t)} F_{\tau^Y} \right]$ from withholding his report forever.

IC-∞ therefore amounts to the requirement that

$$F_\Lambda \geq \mathbb{E}^B_\Lambda \left[\int_\Lambda^{\tau^Y} e^{-\rho(s-\Lambda)} b \, ds + e^{-\rho(\tau^Y - \Lambda)} F_{\tau^Y} \right] = F^*_\Lambda + \mathbb{E}^B_\Lambda \left[e^{-\rho(\tau^Y - \Lambda)} F_{\tau^Y} \right]$$

whenever $\Lambda < \tau^Y$. In particular, given $F \geq 0$, the weaker inequality $F_\Lambda \geq F^*_\Lambda$ must also hold whenever $\Lambda < \tau^Y$. And since $F^*_{\tau^Y} = 0$ by construction, it must be that $F_{\tau^Y \wedge \Lambda} \geq F^*_{\tau^Y \wedge \Lambda}$ for any IC-∞ bonus process implementing τ^Y. Hence if F^* itself satisfies IC-∞, it must be profit-maximizing among all bonus processes implementing τ^Y. And indeed F^* is IC-∞ given $F^*_{\tau^Y} = 0$, proving the lemma. Note that under F^*, the expert obtains exactly the same expected payoffs under the strategies $\Lambda' = \Lambda$ and $\Lambda' = \infty$; the payment he receives at the time of his report is always just enough to make him indifferent between reporting immediately and withholding his report forever.

The following lemma shows that the solution to the relaxed problem established in Lemma 4 satisfies IC-B and so solves the original unrelaxed problem.

Lemma 5. *For any $\mathbb{F}^Y$-stopping time τ^Y, (F^*, τ^Y) satisfies IC-B when F^* is as defined in Lemma 4.*

To understand this result, consider the expert's payoff from delayed reporting strategy under F^*. Delaying a report leads to the collection of flow rents for some time, followed by payment of F^*, which by construction is exactly equal to the expected flow rents he would have collected by continuing to withhold his report forever. Hence all delayed reporting policies yield precisely the same expected payoff as the policy $\Lambda' = \infty$, which in turn provides the same expected payoff as truthful reporting.

The simplicity of this characterization flows from the stationarity of the setting once the Bad state has been reached. Were it the case that the state continued to evolve after the Bad state were reached, Lemma 4 would still hold under a suitable generalization of the conditional expectation used to define F^*. However, that bonus process would not be guaranteed to satisfy IC-B. In section 6.1 I show how the technique just developed can be generalized to analyze projects with richer state spaces.

3.4 The firm's virtual profit function

Lemmas 4 and 5 establish that any $\mathbb{F}^Y$-stopping time τ^Y is implementable by an IC-B contract, with associated profit-maximizing bonus process

$$F_t^* = \mathbb{E}_t^B \left[\int_t^{\tau^Y} e^{-\rho(s-t)} b \, ds \right].$$

The following proposition leverages this fact to prove the first major result of the paper. The following proposition establishes that when F^* is eliminated from the firm's profit function, the resulting optimization problem for τ^Y can be stated elegantly in terms of maximizing an expected discounted flow of virtual profits.

Proposition 1. *Let τ^Y be any $\mathbb{F}^Y$-stopping time and $\Pi[\tau^Y]$ be the supremum of profits achievable by IC-B contracts implementing τ^Y assuming truthful reporting by the expert. Then*

$$\Pi[\tau^Y] = \mathbb{E} \left[\int_0^{\tau^Y} e^{-\rho t} (\pi_t r_G - (1 - \pi_t) b) \, dt \right]. \tag{2}$$

Recall that $\pi_t = \mathbb{P}_t^Y\{\Lambda > t\}$ is the probability that the project's lifespan has not yet lapsed by time t, conditional on the history of output up to that time. Were the firm unable to employ an expert, π_t would also be the firm's posterior beliefs about the current state at time t. Of course, under truthful reporting by the expert the firm's posterior beliefs are degenerate and equal to $\mathbf{1}\{\Lambda > t\}$, and π_t possesses no inferential significance. Nonetheless, it may still be computed by the firm, and turns out to be of great practical importance for designing an optimal contract. Going forward, I will refer to the process π as the firm's *naive beliefs* about the state.

Proposition 1 establishes that the firm's optimal termination time does not directly maximize expected discounted flow profits, but instead optimizes expected discounted *virtual profits*, where virtual profits are r_G per unit time while the state of the project is Good, and $-b$ per unit time afterward. The following heuristic argument explains why. Consider any contract with termination policy τ^Y and bonus process as specified in Lemma 4. This contract operates the project until either τ^Y is reached or the state switches; in the latter case, the projected is halted at the time of the switch and the firm pays the expert his expected discounted flow of benefits from allowing the project to continue operating until τ^Y. The firm's expected profits under such a contract are therefore identical to a counterfactual setting with no expert in which the project operates until τ^Y but yields flow profits $-b$

rather than r_B when the state is Bad. And from an ex ante perspective, conditional on $\mathcal{F}_t^Y$ the fraction of the time the project is in the Good state is precisely π_t. Thus instantaneous expected profits at any time, conditional on $\mathcal{F}_t^Y$, are $\pi_t r_G - (1-\pi_t)b$. Integrating discounted flow profits over time and applying the law of conditional expectations yields equation (2).

Proposition 1 reduces the contracting problem with an expert to the solution of a particular virtual optimal stopping problem, closely related to the firm's problem without an expert (cf. equation (1)). In this virtual problem the firm learns about Λ as if no expert were available but incurs a reduced average flow cost of operating the project in the bad state of $-b$ rather than r_B.

Remark. *Public randomization is unnecessary to achieve the optimum of* $\Pi[\cdot]$.

This fact follows from the observation that, given the independence of the public randomization device from Y and Λ, any τ^Y employing randomization may be considered a distribution over stopping times τ_0^Y which don't condition on the public randomization device. Hence one may write $\Pi[\tau^Y] = \mathbb{E}[\Pi[\tau_0^Y]]$, where $\Pi[\tau_0^Y]$ is a random variable measurable with respect to the outcomes of the public randomization device. Then if τ^Y is to optimize $\Pi[\cdot]$, with probability 1 the randomization device must select a τ_0^Y optimizing $\Pi[\cdot]$ among all stopping times not using public randomization. As all such stopping times yield the same expected profits, τ^Y might as well be chosen not to employ randomization.

3.5 The optimal public deadline

Proposition 1 reduces the firm's optimal contracting problem to solving a single-person optimal stopping problem. In this section I discuss properties of the solution to this problem, and highlight qualitative insights for the design of an optimal public deadine.

Notice that the firm's virtual flow profits are increasing in π_t, and reach zero at the posterior odds ratio $\frac{\pi_t}{1-\pi_t} = b/r_G$, or equivalently at posterior beliefs $\pi_t = b/(b+r_G)$. Let $\tau^\dagger \equiv \inf\{t \ : \ \pi_t \leq b/(b+r_G)\}$ be the stopping time at which the firm's virtual flow profits first drop below zero. Certainly the firm optimally continues operating the project as long as flow profits are positive:

Remark. *If* τ^* *is an optimal public deadline, then* $\tau^* \geq \tau^\dagger$ *a.s.*

On average, π_t declines over time due to the expected arrival of a state switch. Indeed, for any times t and $s > t$, $\mathbb{E}_t[\pi_s \mid \mathcal{F}_t^Y] = \frac{1-H(s)}{1-H(t)}\pi_t \leq \pi_t$. Thus if the firm had to make a once and for all decision to stop or continue at $\tau^\dagger$, it would optimally stop. But because the firm may halt operations at any time, it retains a real option to continue the project temporarily and halt later if its beliefs continue to deteriorate. The optimization of $\Pi[\tau^Y]$ then amounts to calculating the value of this real option.

If the firm learned very little about the state from output, say because of high output variability, then the option to wait and learn would be worth very little, and $\tau^\dagger$ would be an approximately optimal stopping rule. (In the limit of no learning, $\tau^\dagger$ would be exactly optimal, as I showed in Section 4.1.) However, if output is sufficiently informative about the current state, π_t may often move upward. In this case the firm has an incentive to continue operating past $\tau^\dagger$, in the hopes of observing good runs of output that boost its beliefs about the state. The optimal public deadline is therefore a belief threshold $\underline{\pi}(t)$ sufficiently low that the value of waiting for possible good news is outweighed by the flow costs of operating at the current low beliefs. The following result establishes this fact rigorously.

Proposition 2. *There exists a function* $\underline{\pi}^* : \mathbb{R}_+ \to [0, b/(b + r_G)]$ *such that* $\tau^* = \inf\{t : \pi_t \leq \underline{\pi}^*(t)\}$ *is an optimal public deadline, and if* τ^{**} *is any other optimal public deadline, then* $\tau^{**} \geq \tau^*$ *a.s.*

The optimal belief threshold will typically be time-varying given the inhomogeneity of the state transition process. An important exception is when the state transition rate is homogeneous, i.e. $H(t) = H(0) + (1 - H(0))(1 - \exp(-\alpha t))$ for some transition rate $\alpha \geq 0$. In that case current beliefs are a sufficient statistic for the future evolution of the project, and the optimal threshold is time-invariant.

Calculating the optimal belief threshold $\underline{\pi}$ is a standard exercise in optimal stopping. As its solution involves no technical innovations, I defer further discussion to Appendix B. In the remainder of this subsection, I highlight the qualitative implications of Proposition 2 for the design of optimal deadlines which condition on past output. The key result will be a monotonicity result on the sensitivity of the deadline to runs of output as termination draws nearer. To illustrate this effect cleanly, I assume a constant hazard rate of state transitions $\alpha > 0$; this time-homogeneity allows for a consistent notion of project dynamics when the project is "far from" and "close to" termination. I also focus on the case of Brownian output $Y_t^\theta = r_\theta t + \sigma Z_t^\theta$ for some $\sigma > 0$, with no jumps in the output process. In Appendix C I provide a detailed analysis of the general model with jumps.

Any deadline which is a threshold rule in naive beliefs is not a deterministic cutoff date, but rather a random time which varies in response to good or bad runs of output. To study this effect, I analyze the stochastic process

$$\Delta_t \equiv \mathbb{E}_t^G[\tau^*] - t$$

which captures the expected time until project termination assuming the project does not expire before the public deadline. (I will assume that the optimal threshold $\underline{\pi}^* > 0$, so that this quantity is finite.) The quantity Δ_t is a natural forecast of the time until the firm terminates the project, subject to revision depending on future project performance. On average the expert expects Δ_t to diminish at rate 1, so long as the project is still in the Good state. Of course, since naive beliefs are sensitive to realizations of output, the expert also expects that Δ will fluctuate randomly due to these output fluctuations. This sensitivity will vary over time depending on the current state of beliefs, and the evolution of the output sensitivity of the deadline over the lifetime of the contract is a central design feature of the optimal deadline.

The expected time until termination can be expressed as $\Delta_t = \xi(\pi_t)$ for a strictly increasing value function $\xi : [\underline{\pi}^*, 1] \to \mathbb{R}_+$ characterized in Appendix C. Ito's lemma can then be used to show that

$$d\Delta_t = -dt + \frac{r_G - r_B}{\sigma}\pi_t(1-\pi_t)\xi'(\pi_t)\, d\overline{Z}_t^G,$$

where $\overline{Z}_t^G \equiv \sigma^{-1}(Y_t - r_G t)$ is a Brownian motion under the probability measure $\mathbb{P}^G$. The first term captures the fact that the forecasted time to termination drifts down at rate 1 on average, while the second term reflects responsiveness of the forecast to output surprises relative to expected levels.

The sensitivity of Δ to current project performance is captured by the function $\phi : [\underline{\pi}^*, 1] \to \mathbb{R}_+$ taking values

$$\phi(x) = \frac{r_G - r_B}{\sigma}x(1-x)\xi'(x).$$

The time-t sensitivity $\phi(\pi_t)$ is modulated by three factors: 1) the signal-to-noise ratio (SNR) $(r_G - r_B)/\sigma$, a measure of how distinguishable runs of output are in the Good and Bad state; 2) the Bayes' rule factor $\pi_t(1-\pi_t)$, capturing the fact that good and bad runs lead to larger belief updating the less definite are current beliefs about the state; and 3) $\xi'(\pi_t)$, linking changes in beliefs to changes in time to termination. The final factor ξ' varies with π_t, and captures increased sensitivity of the exact termination time to the current state when termination is close and short runs of output could mean the difference between termination or continuation.

It can be shown that ξ is a strictly increasing, strictly concave function, hence a given change in beliefs leads to a larger change in the deadline as termination draws nearer. However, whenever $\underline{\pi}^* < 1/2$, belief revisions tend to diminish close to termination, so signing the net effect requires careful balancing. The following lemma proves that the first effect dominates, so that sensitivity of Δ_t to output surprises is everywhere decreasing in π_t, i.e. is increasing the closer the project is to termination.

Lemma 6. *If $\underline{\pi}^* > 0$, then ϕ is a strictly decreasing function satisfying $\phi(1) = 0$.*

A key qualitative implication of the optimal termination rule characterized by Proposition 2 is therefore that sensitivity of the deadline to runs of output varies over the lifetime of the project, and increases the closer the project is to termination. In Appendix C I show that this result continues to hold in models with jumps, when runs of output are defined to be surprises in the continuous part of the output process. The response of the deadline to jumps displays richer dynamics, and can vary non-monotonically in π_t in general. Despite this complexity I establish that under a regularity condition, revisions of the deadline in response to jumps are also decreasing in π_t for sufficiently high beliefs.

3.6 Verifying incentive compatibility

Recall that the contract induced by optimizing $\Pi[\tau^Y]$ solves only the relaxed contracting problem which ignores all IC-G constraints. I now return to the problem of verifying that this solution satisfies full incentive compatibility and thus is an optimal contract. To do this, I develop a simple sufficient condition for incentive-compatibility which I then show is satisfied by the optimal IC-B contract.

The sufficient condition involves the expert's ex post utility process U. For a given contract (F, τ^Y), this process is defined to be

$$U_t \equiv \int_0^{t\wedge\tau^Y} e^{-\rho s} b\, ds + e^{-\rho(t\wedge\tau^Y)} F_{t\wedge\tau^Y}.$$

When $F = F^*$, where F^* is the bonus-minimizing payment process characterized in Lemma 4, this process can be equivalently written

$$U_t = \mathbb{E}_t^B\left[\int_0^{\tau^Y} e^{-\rho s} b\, ds\right].$$

U_t captures the ex post total utility of the expert supposing he reports a state switch at time t. The expert's ex ante utility from reporting policy Λ' is then just $\mathbb{E}[U_{\Lambda'}]$.

It turns out that incentive-compatibility holds so long as U drifts upward whenever the state of the project is Good. This fact may be stated formally using the following definitions.

Definition 8. *For any random variable X, let $\mathbb{E}^G[X] = \int X \, d\mathbb{P}^G$. For each $t \in \mathbb{R}_+$, let $\mathbb{E}_t^G[X]$ be the expectation of X under $\mathbb{P}^G$ conditional on $\mathcal{F}_t$.*

Analogously to $\mathbb{E}_t^B$, this definition formalizes a notion of expected value "conditional on the state never switching".

Definition 9. *Suppose X is an $\mathbb{F}$-adapted process. Then X is a* B-martingale *if $\mathbb{E}_t^B[X_s] = X_t$ for all $s > t$, and is a* G-martingale *if $\mathbb{E}_t^G[X_s] = X_t$ for all $s > t$. Super- and submartingales are defined analogously.*

Remark. *U is a B-martingale.*

The following lemma establishes a sufficient condition for incentive-compatibility.

Lemma 7. *Suppose τ^Y is an $\mathbb{F}^Y$-stopping time, and let F^* be as defined in Lemma 4. If U is a G-submartingale, then (F^*, τ^Y) is incentive-compatible.*

The proof of this lemma is very simple - if U is a G-submartingale, then the expert's ex post utility drifts upward over time so long as the state is Good. Thus the expected payoff from waiting until Λ to report a state switch must be at least as high as from reporting at any earlier time.

For general public deadlines, U is not guaranteed to be a G-submartingale. For instance, if τ^Y increases following bad runs of output and decreases following good runs, the fact that U is a B-martingale implies that U drifts downward while the state is Good. However, deadlines which follow a time-dependent threshold rule in naive beliefs do not behave this way. Because positive runs of output boost naive beliefs, any threshold rule will induce a positive association between τ^Y and past output. Therefore if U is a B-martingale, the increased incidence of good runs of output under $\mathbb{P}^G$ ensures that U will be a G-submartingale, implying full incentive-compatibility by Lemma 7. The following result verifies this intuition formally.

Proposition 3. *Let τ^Y be any $\mathbb{F}^Y$-stopping time such that $\tau^Y = \inf\{t : \pi_t \leq \underline{\pi}(t)\}$ for some function $\underline{\pi} : \mathbb{R}_+ \to [0, 1]$. Then U is a G-submartingale.*

Corollary. *The optimal public deadline τ^* characterized by Proposition 2 induces an incentive-compatible contract.*

This proposition and its corollary close the loop on the construction of an optimal contract, justifying my initial conjecture that none of the IC-G constraints bind for an optimal contract.

4 Economic fundamentals

Conceptually, the firm's contract design problem can be divided into two optimization problems for *scope* and *sensitivity*. The scope optimization sets the average lifespan of the project, while the sensitivity component calibrates how aggressively project scope responds to output surprises during operations.

4.1 Optimizing project scope

The scope component of the design problem can be illustrated by shutting down the firm's information channel and restricting contracts to condition only on the expert's report (but not the history of output). Without observing output, the firm relies entirely on the expert's report to decide when to terminate the project. In particular, truthful reporting can't be checked by looking for good or bad runs of output. As a result, the firm can induce truthful reporting only by compensating the expert for all fringe benefits lost if project termination is sped up by the report.

Concretely, suppose the firm imposes a deterministic deadline T for termination absent any reports, and the state switches at time $t < T$. If the firm responds to a truthful report of the switch by terminating the project at time $t' \in [t, T)$, the expert's lost benefits discounted from time t are $e^{-\rho(t'-t)} \frac{b}{\rho}(1 - e^{-\rho(T-t')})$. So the firm must make expected payments to the expert of at least this amount following a report at time t to achieve incentive-compatibility. Meanwhile the expected savings to the firm from terminating at time t' rather than T and avoiding losses after the state switch are $e^{-\rho(t'-t)} \frac{|r_B|}{\rho}(1 - e^{-\rho(T-t')})$. Since $|r_B| > b$ by assumption, these avoided losses are always greater than the associated incentive payment. In fact, the net savings is increasing in $T - t'$, so the firm optimally sets $t' = t$, i.e. terminates as soon as the expert has reported a state switch. If the firm pays the expert nothing until termination and then a lump sum equal to lost fringe benefits from termination until time T, then the expert receives the same total utility regardless of when he reports the state switch. Hence such a contract is incentive-compatible, and must surely be cost-minimizing among all IC contracts with termination deadline T.

The analysis of the previous paragraph reduces the firm's contracting problem to the choice of a single deadline T, at which the project is shut down absent a report from the expert that the state has switched. Any choice of $T \in \mathbb{R}_+ \cup \{\infty\}$ can be made incentive-compatible through sufficiently large payments upon termination. And the minimum required payment at time $t < T$ to achieve incentive-compatibility is $F_t = \frac{b}{\rho}(1 - e^{-\rho(T-t)})$, which is increasing in T. Hence the firm faces a tradeoff between output and payments - the higher T is set, the more output is collected when the project has a long lifespan, but the larger are payments to the expert when the project is short-lived.

Define a family of contracts $\mathcal{C}^T = (\Phi^T, \tau^T)$ for each $T \in \mathbb{R}_+ \cup \{\infty\}$ by $\tau^T = \Lambda' \wedge T$ and

$$\Phi_t^T = \frac{b}{\rho}\left(1 - e^{-\rho(T-\tau^T)}\right)\mathbf{1}\{t \geq \tau^T\}.$$

When Y is unobserved, the arguments above show that the firm's optimal contract must lie in this family of contracts for some choice of T. The following remark characterizes the optimal T.

Remark. *Let* $T = \inf\left\{t \ : \ \frac{1-H(t)}{H(t)} \leq \frac{b}{r_G}\right\}$, *with* $T = \infty$ *if* $\frac{1-H(\infty)}{H(\infty)} > \frac{b}{r_G}$. *Then if* Y *is unobserved by the firm,* $\mathcal{C}^T$ *is an optimal contract.*

To derive this result, consider the effect on profits of increasing T by dT and moving from contract $\mathcal{C}^T$ to $\mathcal{C}^{T+dT}$. Whenever $\Lambda > T$, additional output is obtained on the margin, yielding expected flow profits $r_G e^{-\rho T}\,dT$. Conversely, whenever $\Lambda \leq T$, the terminal payment to the expert must be increased by $\frac{b}{\rho}e^{-\rho(T-\Lambda)}\rho\,dT$, yielding a cost increase (discounted from time zero) of $be^{-\rho T}\,dT$ independent of the exact realization of Λ. The net change in expected profits from increasing T by dT is therefore $((1-H(T)r_G - H(T)b)e^{-\rho T}\,dT$. These incremental profits are positive so long as $(1-H(T))/H(T) \geq b/r_G$ and negative otherwise, yielding the optimal deadline in the remark.

This derivation illustrates the basic price-quantity tradeoff faced by the firm when setting an optimal scope: discounting from time T, increasing T collects marginal quantity $(1 - H(T))r_G$ at unit price $H(T)b$. As marginal output is declining while price is increasing, the benefits of extending the deadline diminish with T (and eventually turn negative, if $\Lambda = \infty$ is sufficiently unlikely).

4.2 Optimizing project sensitivity

Section 4.1 shows how adjustment of a deterministic deadline T allows the firm to trade off between allocative efficiency and payments to the expert. The firm has one additional tool to ameliorate its agency problem - it can condition T on the history of output, moving from a deterministic to a stochastic deadline.

Why is adding stochasticity to T helpful to the firm? Begin with a static deadline T, and construct a stochastic deadline τ by adding a mean-preserving spread to T. Formally, $\tau = T + \varepsilon$ where ε is a random variable with zero mean conditional on the state remaining Good until τ. This construction is useful because the firm can condition ε on the history of output, and in particular can choose ε to be positively correlated with high output histories. In this case whenever the expert delays reporting a state switch, he incurs a penalty to T due to the decline in average output under the Bad state as compared to the Good state. In other words, the mean of ε will be negative conditional on a delayed report by the expert. This effect lowers the expected stream of fringe benefits available from delaying a report, and so lowers the required payments to the expert to discourage late reporting. Adding stochasticy to the deadline therefore allows the firm to shrink incentive payments to the expert.

Unfortunately for the firm, adding randomness to the deadline doesn't come for free. Increasing the spread of ε incurs a cost to expected discounted output, due to the convexity of the discount factor $e^{-\rho t}$. If average output flow r_G is collected up to the termination time $\Lambda \wedge \tau$, then total expected discounted output is $\mathbb{E}\left[\frac{r_G}{\rho}\left(1 - e^{-\rho(\Lambda\wedge\tau)}\right)\right]$. As $1 - e^{-\rho(\Lambda\wedge\tau)}$ is concave, increasing the variability of τ lowers expected discounted output. Therefore the firm faces a price-quantity tradeoff in designing the variability of τ as well as its mean - the higher the spread, the lower the unit price of output but the less overall output is obtained.

The substance of the analysis in this paper is how to optimally design the distribution of ε. While the marginal distribution of ε can be chosen arbitrarily, its correlation with the underlying state is constrained by the structure of the output process in a complex way. In addition, the choice of T will affect the optimal choice of ε, and vice versa. My analysis untangles these interactions to construct an optimal contract.

5 Discussion

5.1 The impact of hiring an expert

Hiring an expert has a crisp, precisely characterizable impact on project dynamics versus a setting in which the firm receives no expert advice. Recall that without an expert, the firm's expected profits from a given termination policy τ^Y are

$$\mathbb{E}\left[\int_0^{\tau^Y} e^{-\rho t}(\pi_t r_G + (1-\pi_t) r_B)\, dt\right].$$

In this setting the process π reflects the firm's true beliefs about the state, but its distribution over paths is identical to the firm's naive beliefs in the problem with an expert. In other words, the firm "learns" about the state in the same way with or without the expert. Given the assumption that $|r_B| > b$, and in light of Proposition 1, the entire impact of the expert on the firm's optimal stopping problem is to lower the penalty for operating in the Bad state from r_B to $-b$.

In particular, it will still be the case without an expert that an optimal stopping rule is a time-dependent threshold in beliefs, as in Proposition 2. But this threshold will be higher at each moment in time without an expert versus with one, reflecting the increased cost of operating the project in the Bad state without an expert. Since the distribution of paths of beliefs is identical with and without an expert, this implies that the project is operated for a shorter time without an expert almost surely, conditional on no state change having occurred. In other words, the expert reduces the severity of Type I errors, in which the project is terminated before the state has actually switched. Of course, even with an expert such errors are typically not eliminated entirely, as termination due to reaching the public deadline is always inefficiently early ex post.

The presence of an expert exerts one additional major influence on project operation - he ensures that the project never actually operates into the Bad state. Thus an optimal contract with an expert does not exhibit Type II errors, in which the project continues operating after the state has switched. The impact of an expert can therefore be succinctly summarized as follows: he increases the ex post efficiency of project operation in every state of the world, and reduces the severity of Type I errors while completely eliminating Type II errors.

This result is a bit surprising, as the basic agency problem faced by the firm is the expert's desire to *increase* Type II errors. A first guess at the resolution of this conflict might therefore be that the expert reduces or eliminates Type I errors at the cost of more Type II errors. Yet the optimal contract actually exhibits the opposite asymmetry. This outcome crucially depends on the firm's ability to commit to terminal payments which compensate the expert for reporting bad news that ends the project.

Finally, note that hiring of an expert is *not* payoff-equivalent equivalent to simply changing r_B to $-b$ in the single-person problem. Such a change would both change the payoff structure *and* decrease the rate of learning, due to a reduction in the signal-to-noise ratio of the output process. In contrast, the comparison of project dynamics just performed relies crucially on the fact that the learning dynamics are identical across the two environments.

5.2 Comparative statics

Proposition 2 establishes that an optimal public deadline is a (generally time-dependent) threshold rule $\underline{\pi}(t)$ in naive beliefs. The threshold's characterization as the solution to an option value problem allows for a straightforward deduction of comparative statics.

All comparative statics with respect to b, ρ, and the distribution of Λ hold for arbitrary output processes. For clarity, when discussing comparative statics with respect to r_G, r_B, and volatility of output I take Y to be a Brownian motion with drift in each state of the world, with $Y^\theta = r_\theta t + \sigma Z^\theta$ for some (common) $\sigma > 0$ and independent Brownian motions Z^G and Z^B. Analogous comparative statics hold when Y is a Poisson process in each state of the world, and under some mild regularity conditions similar results can be derived for more general output processes.[12] To perform comparative statics with respect to the distribution of Λ, I define the hazard rate of a state switch by $\alpha(t) \equiv h(t)/(1 - H(t))$, where h is the derivative of H, which exists a.e. given that H is monotone. Whenever H is absolutely continuous, a shift in the hazard rate function α induces a new distribution over Λ via

$$H(t) = 1 - (1 - H(0)) \exp\left(-\int_0^t \alpha(s)\, ds\right).$$

If H is not absolutely continuous, then an increase in the hazard rate distribution can be taken as shorthand for a pointwise upward shift in H, or equivalently a downward shift in Λ in the FOSD sense.

Any parameter changes which raise the break-even virtual profit threshold will mechanically induce a higher optimal termination threshold, all else equal. Such changes include a decrease in r_G or an increase in b. Also, any parameter changes which diminish the option value of operating the project below break-even will increase the optimal termination threshold, all else equal. These changes include an increase in r_B, a decrease in r_G, or an increase in σ, all of which decrease the rate of learning; an increase in the discount rate, which increases the relative cost of operating below the break-even point; or an increase in the future path of α, which push down the distribution of future beliefs. Note that changing r_G affects both the break-even point and the option value of the project, but in the same direction. Table 1 summarizes these results.

Another important comparative statics exercise is how changes in parameter values impact the total profits of the firm under an optimal contract. An increase in b increases flow losses in the Bad state, decreasing firm profits under any contract and thus certainly under the optimal contract. Meanwhile an increase in ρ increases profits[13] under the optimal

[12] In particular, whenever the two states are not immediately distinguishable, there exists a $\sigma > 0$ and pure jump processes $\widetilde{Y}^\theta$ such each Y^θ is decomposable as $Y^\theta = \widetilde{r}_\theta t + \sigma Z^\theta + \widetilde{Y}^\theta$, with $\widetilde{r}_\theta = r_\theta - \mathbb{E}[\widetilde{Y}_1^\theta]$. A change in r_θ can be operationalized as a change in $\widetilde{r}_\theta$, holding fixed Z^θ and $\widetilde{Y}^\theta$. The comparative statics with respect to r_θ reported for the Brownian case hold whenever $\widetilde{r}_G > \widetilde{r}_B$, so that whenever r_G increases, Y^G and Y^B become more distinguishable, while whenever r_B increases, they become less distinguishable.

[13] For this comparative static, I examine the change in the normalized profits $\rho\Pi$ to net out the mechanical diminuation in flow profits at each period as ρ increases. The comparative static for the non-normalized

Parameter	$\Delta\underline{\pi}(t)/\Delta$Parameter
b	+
ρ	+
r_G	-
r_B	+
σ	+
$\{\alpha(s)\}_{s>t}$	+

Table 1: Comparative statics of the termination threshold

contract, as a higher discount rate puts more weight on flow profits earlier in the project, when the state is Good a higher fraction of the time. Similarly, an increase in the path of the state switching rate α decreases achievable profits by decreasing the amount of time the project spends in the Good state. Finally, an increase in r_G or a decrease in r_B or σ all increase optimal profits, as they speed learning and for r_G boost flow profits while the project is Good. Table 2 summarizes these results.

Parameter	$\Delta\Pi/\Delta$Parameter
b	-
ρ	+
r_G	+
r_B	-
σ	-
α	-

Table 2: Comparative statics of optimal profits

5.3 Dynamic verification

One important special case nested by my model is dynamic verification of a report about a persistent project state. This case corresponds to $H(t) = \pi_0$ for all t, where $\pi_0 \in (0,1)$ is the probability that a long-run project is worth undertaking forever. In this case the optimal termination threshold is a constant $\underline{\pi}$, and an optimal policy takes one of two forms: if $\pi_0 < \underline{\pi}$, then the project isn't worth undertaking at all (with or without an expert), and it is simply abandoned immediately. Otherwise, the expert is asked to report at time zero whether the project is worthwhile. If not, he is paid a lump-sum consulting fee and the

profit function would be reversed, as in that case a decrease in ρ would both diminish early flow profits (when the state is likely to be Good) and diminish the option value of learning about the state whenever flow profits are negative.

project is abandoned. Otherwise, the expert is employed and the project is operated until and unless naive beliefs drop below $\underline{\pi}$, at which point the project is terminated.

Note that in case the project is Good, the expert is paid no incentive bonuses - all his compensation comes from the stream of flow benefits accrued during project operation. The termination rule when the project is Good is therefore designed solely to limit the size of the consulting fee that must be paid in case the project is initially Bad. The optimal contract can be thought of as treating the expert's recommendation that the project should be undertaken with some skepticism, with the project's subsequent performance used to check the expert's report. This sort of dynamic verification is very similar to that of Varas (2017), with the difference that in that model, impatience of the agent relative to the principal limits the length of the optimal verification period, whereas in my model excessively long verification periods accrue inalienable rents to the expert and so offset the incentive effects of the verification.

6 Robustness

6.1 Many-state projects

The key to my derivation of an optimal contract is the identification of a cost-minimizing completion bonus structure which induces honest reporting by the expert. In this section I show how the key insights of that derivation are readily adapted to analyze projects exhibiting more complex operating profiles over their lifespan. In particular, I argue that when the firm offers delegation contracts to the expert, the presence of more than two states does not change the basic conclusions of the analysis, and in particular the optimal delegation contract is completely insensitive to the number of "bad" states and the rate of transition between them.

For simplicity, I will assume in this subsection that the project state is a Markov chain, with constant intensity of transitions between states over time. The project state space is $\Theta = \{\underline{n}, \underline{n} - 1, ..., \overline{n} - 1, \overline{n}\}$, where $\overline{n} > 0 \geq \underline{n}$ are both integers. The project's mean output per period is r_θ for $\theta \in \Theta$, with $r_\theta + b \geq 0$ for $\theta \geq 1$ and $r_\theta + b < 0$ for $\theta \leq 0$. Thus the project is efficient to operate exactly on the set of strictly positive states. State transitions are irreversible and always proceed from state n to state $n-1$, with transitions occurring at rate $\lambda_n > 0$ (which may depend on n). Define $\Lambda^{(n)} = \inf\{t \ : \ \theta_t = n\}$ to be the time of the state transition from state $n+1$ to n.

When the project has many states, the natural extension of the contract space in the baseline model is the set of dynamic delegation protocols, under which the expert is given limited discretion over the time at which the project is terminated. (At the end of this subsection, I return briefly to consider more complex mechanisms.) When designing a delegation protocol, the firm must decide when to rescind delegation and terminate the project automatically, and must give the expert incentives to terminate efficiently when he has discretion. As in the baseline model, the expert should be given incentives to terminate exactly when state $\theta = 0$ is reached. Any sooner, and operational efficiency is either reduced, or does not increase by enough to justify the required incentive payments. Any later, and the firm would gain by raising incentive payments to induce termination sooner.

Suppose the firm rescinds delegation and terminates at stopping time τ^Y if the expert has not yet taken any action. What are the cost-minimizing completion bonuses needed to induce efficient termination by the expert, supposing $\Lambda^{(0)}$ occurs prior to τ^Y? As in the baseline model, answering this question involves correctly guessing the binding incentive constraint for the expert. Recall that in the two-state case, the binding deviation involved the expert simply never terminating the project. In the many-state model, the binding deviation turns out to involve terminating the project at $\Lambda^{(-1)}$ rather than $\Lambda^{(0)}$. (Note that when there is a single bad state, $\Lambda^{(-1)} = \infty$ and so this reduces to the binding deviation in the baseline case.) Call a contract IC-0 if it is proof against such deviations. I will characterize the optimal IC-0 contract, then explain why the contract is fully incentive-compatible.

I first derive an analog of Lemma 4 for the many-state model. To state the result, I rely on an auxiliary probability measure $\mathbb{P}^{(0)}$, under which the project state is fixed at 0 forever. $\mathbb{E}^{(0)}$ takes expectations over uncertainty in output with respect to that measure, and $\mathbb{E}^{(0)}_t$ conditions on the history of output up to time t.

Lemma 8. *For any $\mathbb{F}^Y$-stopping time τ^Y, define a bonus process $F^{(0)}$ via*

$$F^{(0)}_t = \mathbb{E}^{(0)}_t \left[\int_{t \wedge \tau^Y}^{\tau^Y} e^{-\rho(s-t)} b \, ds \right].$$

Then $F^{(0)}$ is an $\mathbb{F}^Y$-adapted process, $(F^{(0)}, \tau^Y)$ satisfies IC-0, and $F^{(0)}_{\Lambda^{(0)} \wedge \tau^Y} \leq F_{\Lambda^{(0)} \wedge \tau^Y}$ a.s. for every F such that (F, τ^Y) satisfies IC-0. In particular, $(F^{(0)}, \tau^Y)$ maximizes expected profits among all IC-0 contracts (F, τ^Y).

The intuition for this result is that so long as the expert shutters the project no later than $\Lambda^{(-1)}$, the project state remains fixed at $\theta = 0$ throughout the interval of delay, which project output evolving accordingly. So to deter delayed reporting, the expert must be compensated for all lost benefits in the interval $[\Lambda^{(0)}, \Lambda^{(-1)}]$ using the correct measure over output. Further, the payment collected at time $\Lambda^{(-1)}$ under a deviation builds in an *additional* stream of benefits which would have accrued after $\Lambda^{(-1)}$ in the counterfactual world in which the project were in state 0 at that time. Summing this stream of benefits recursively shows that the expert must be compensated for at least the entire stream of benefits until exogenous project termination as if the project state never deteriorated beyond 0.

Inserting $F^{(0)}$ into the firm's profit function yields the reduced objective

$$\Pi[\tau^Y] = \widetilde{E}\left[\int_0^{\tau^Y} e^{-\rho t} \sum_{n=0}^{\overline{n}} \widetilde{\pi}_t(n)\widetilde{r}_n \, dt\right].$$

In this expression, $\widetilde{E}$ takes expectations with respect to a measure under which θ is a Markov process with transition rates λ_n from state n to $n-1$ for each $n \geq 1$, but with 0 an absorbing state. The processes $\widetilde{\pi}_t(n) = \widetilde{E}[\mathbf{1}\{\theta_t = n\} \mid (Y_s)_{s\leq t}]$ track the posterior distribution of states under the adjusted measure, and $\widetilde{r}_n = r_n$ for $n \geq 1$ while $\widetilde{r}_0 = -b$.

Under the relaxed problem considering only the IC-0 constraints, the firm's contracting problem reduces to an optimal stopping problem very similar to the problem without an expert. Recall that in the two-state model the adjustment to the objective from the expert's presence is fully captured by a change in the payoff in the Bad state. By contrast, with additional states the expert's presence changes the problem in two ways. First, analogous to the the two-state case, payoffs in the $\theta = 0$ state are improved from r_0 to $-b$. Second, the state process is truncated at 0 and never evolves beyond it. Thus the details of the state process past this point, including whether there are further "Bad" states, are irrelevant to designing an optimal contract!

Why is the relaxed optimal contract fully incentive-compatible? While the optimal termination policy is in general complex since beliefs are multi-dimensional, low (respectively, high) runs of output will drive beliefs over higher states down (up) and optimally move the project toward (away from) termination. So consider the expert's total utility process

$$U_t = \int_0^t e^{-\rho t} b \, dt + e^{-\rho t} F_t^{(0)}.$$

With the project state fixed at $\theta = 0$, this process is a martingale, which moves upward with high runs of output and downward with low runs due to the nature of the firm's optimal stopping rule. However, when the state has deteriorated further, low runs of output become more likely and high runs less likely, so the total utility process becomes a supermartingale. Thus whenever the state is negative, total utility deteriorates the longer the expert waits, meaning no further delay can be optimal. Conversely, whenever the state is positive total utility is a submartingale, and so early stopping is also never optimal. Thus the optimal relaxed contract is fully incentive-compatible. (This reasoning is very similar to the logic used to prove Proposition 3.)

In this subsection I have focused on characterizing an optimal delegation mechanism in which the expert is given authority to terminate the project and communicates no other information.[14] How would the optimal contract change if the firm used more general mechanisms which allowing multiple reports by the expert? When $\overline{n} \geq 2$, additional communication would allow the firm to dynamically tailor the responsiveness of the project deadline to output. In particular, the firm could offer a sliding scale of deadlines, with longer deadlines accompanied by more stringent output standards. As the project state deteriorated, the expert would opt for more restrictive deadlines with laxer output standards. I conjecture that offering such a contract would allow the firm to concentrate premature terminations in states where output is marginal, further improving profitability.

6.2 An expert with initial capital

So far I have considered an expert who arrives with no initial wealth to contribute to the firm. Suppose instead the expert possessed total wealth $W > 0$ which can be paid into the firm at any time. My analysis is readily adapted to incorporate this possibility.

Without loss, any payments from the expert are made up front, with the contract then respecting limited liability as in the benchmark model. For it will continue to be the case that all payments, either to or from the expert, can be made at the time of termination of the contract. Therefore the profit of any other IC contract could be replicated by charging the expert the discounted value of his largest possible (negative) termination charge up front, and then adding this amount, grown at the discount rate, onto all terminal payments, yielding a contract with the same incentive structure and no negative payments after time zero. This reduction also makes the contract robust to any dynamic IC constraints the firm might otherwise face on extracting payments from the expert ex post.

When all payments from the expert are extracted up front, the firm's problem can be decomposed into two parts. First, the firm charges the expert an upfront amount $W' \leq W$ to join the firm. Afterward, the firm solves an optimal contracting problem just as in the benchmark model, but under an additional participation constraint for the expert. This

[14]These mechanisms are very similar to the class of "no-communication" mechanisms studed in Kruse and Strack (2015).

constraint amounts to

$$\mathbb{E}\left[\int_0^{\tau^Y} e^{-\rho t} b\, dt\right] \geq W',$$

where the left-hand side is the total value of flow benefits plus termination payments anticipated by the expert under a given termination policy τ^Y. (Lemma 4 continues to characterize the optimal payments to the expert under a given termination policy.)

The solution to this problem depends on exactly how wealthy the expert is. If $W \leq \underline{W} \equiv \mathbb{E}\left[\int_0^{\tau^*} e^{-\rho t} b\, dt\right]$, with τ^* the optimal policy in the benchmark model, then the firm optimally charges the expert his entire endowment and operates the project just as in the benchmark model. On the other hand, if $W > \overline{W} = b/\rho$, then the firm optimally charges the expert exactly $\overline{W}$ and then operates the project efficiently, i.e. with $\tau^Y = \infty$.

The interesting case is when the expert has intermediate wealth. In this case the firm optimally charges the expert enough that the participation constraint binds, as otherwise it could increase profits by keeping the termination policy fixed and charging more up-front. However, the expert does not have enough wealth to pay the corresponding charge when $\tau^* = \infty$. It is therefore necessary to explicitly account for the participation constraint in the optimization problem. The firm's problem may be represented by the Lagrangian

$$\mathscr{L}(\tau^Y, W'; \lambda) = W' + \mathbb{E}\left[\int_0^{\tau^Y} e^{-\rho t}(\pi_t r_G - (1-\pi_t) b)\, dt\right] + \lambda\left(\mathbb{E}\left[\int_0^{\tau^Y} e^{-\rho t} b\, dt\right] - W'\right),$$

with λ the Lagrange multiplier on the participation constraint.

If $\lambda \geq 1$, then the optimizer of the Lagrangian is $(\tau^Y, W') = (\infty, 0)$, which is clearly not a maximizer of the true optimization problem. So $\lambda < 1$ for the saddle point corresponding to a maximum of the problem. Then $W' = W$ maximizes the Lagrangian, and dropping terms not depending on τ^Y from the Lagrangian leaves the reduced objective function

$$\mathbb{E}\left[\int_0^{\tau^Y} e^{-\rho t}(\pi_t(r_G + \lambda b) - (1-\pi_t) b(1-\lambda))\, dt\right].$$

The solution to this problem is a time-dependent threshold rule in naive beliefs, just as in the baseline problem. As flow profits in the Good state are higher and lower in the Bad state the higher is λ, increasing λ decreases the optimal threshold at all times and increases the optimal termination time $\tau^*(\lambda)$ in every state of the world. λ is then chosen to be the unique level, say $\lambda^*(W)$, such that the participation constraint just binds when $W' = W$. Clearly $\tau^*(\lambda^*(W))$ is increasing in W almost surely, with the corresponding belief threshold at each time decreasing.

The extension to an expert with capital therefore leads to the following changes to payoffs and project outcomes. From the point of view of project operation, a nonzero amount of wealth $\underline{W}$ is needed to yield any changes in project operation. Past this threshold level of wealth, higher wealth increases the optimal deadline almost surely, until wealth hits an upper threshold $\overline{W}$, past which the project is operated efficiency. Firm profits are strictly increasing in wealth up until $\overline{W}$, at which point they are flat. Finally, the expert's payoff (above and beyond his initial wealth) from participating in the project is decreasing in total wealth up until $\underline{W}$, zero between $\underline{W}$ and $\overline{W}$, and then increasing again beyond $\overline{W}$.

A related question is how much better off the expert is from a unit of additional capital, given the decreased rent extraction it entails. Between wealth levels 0 and $\underline{W}$, the expert is charged his entire wealth and receives a constant total amount of flow benefits from project operation, so he has zero marginal utility of wealth at these wealth levels. Meanwhile above $\underline{W}$ the expert's participation constraint binds, and so his net utility including initial wealth is exactly W and his marginal utility of wealth is 1. Thus while only the firm benefits from injections of capital at low wealth levels, at higher wealth levels both parties benefit.

6.3 The post-termination world

In my model I assume that the expert is crucial to the operation of the project, above and beyond his ability to identify the time of a state switch. Thus when he is terminated, the project must also be shuttered. It is easy to adapt my framework to deal with alternative post-termination options. In particular my framework can accommodate the hiring of a new expert or "going it alone" without one. Note that all the results of the benchmark setting reducing the contracting problem to the design of a public deadline continue to hold regardless of the firm's post-termination options. Therefore the only change that must be made to the analysis is the formulation of the firm's virtual profit function.

Suppose first that, upon terminating the expert, the firm may continue to operate the project on its own without expert advice. This problem may be solved in two steps, as follows. Let $\widetilde{\Pi}(t)$ be the firm's profits from optimally operating the project on its own, when Λ is distributed as $\widetilde{H}(s;t) \equiv \frac{H(s)-H(t)}{1-H(t)}$. This function satisfies

$$\widetilde{\Pi}(t) = \sup_{\tau^Y} \mathbb{E}^t \left[\int_0^{\tau^Y} e^{-\rho s}(\pi_s r_G + (1-\pi_s) r_B)\, ds \right],$$

where $\mathbb{E}^t$ takes expectations with respect to the probability measure under which $\Lambda \sim \widetilde{H}(\cdot;t)$.

It may be calculated just as optimal firm profits in the benchmark model are, with the optimal stopping rule a time-dependent threshold in beliefs.

With the auxiliary function $\widetilde{\Pi}(t)$ in hand, the optimal stopping problem characterized in Proposition 1 for obtaining an optimal contract may then be modified straightforwardly to incorporate this post-termination option, yielding

$$\Pi^* = \sup_{\tau^Y} \mathbb{E}\left[\int_0^{\tau^Y} e^{-\rho s}(\pi_s r_G - (1-\pi_s)b)\,ds + e^{-\rho\tau^Y}\pi_{\tau^Y}\widetilde{\Pi}(\tau^Y)\right].$$

The added term reflects the discounted continuation value of optimally operating the project without an expert, deflated by the probability that the project is still Good by the time τ^Y is reached.

This modification preserves the problem's basic recursive structure in the state variable $X_t = (\pi_t, t)$. In particular, a time-dependent termination threshold will continue to be optimal. The sole modification to the technique comes when computing the option value of continuing the project, where one must insert a termination payoff of $\pi_t\widetilde{\Pi}(t)$ rather than 0 as in the benchmark problem. Unsurprisingly, this positive termination payoff will push up the optimal termination threshold at all times compared to the setting with no ability to operate post-termination.

Now suppose instead that, upon terminating the expert, the firm may hire a new one at a cost $K > 0$. The expert may be replaced arbitrarily many times.[15] This problem may be written recursively as follows. Let $\Pi^*(t)$ be the profit of an optimal contract, with the option to replace the expert, when Λ is distributed as $\widetilde{H}(s;t) \equiv \frac{H(s)-H(t)}{1-H(t)}$. Then $\Pi^\dagger(t) = \max\{\Pi^*(t) - K, 0\}$ is the net continuation profit from firing the expert at time t when the state is still Good. Incorporating this post-termination option into the optimal stopping problem characterized in Proposition 1, $\Pi^*(t)$ must satisfy

$$\Pi^*(t) = \sup_{\tau^Y} \mathbb{E}^t\left[\int_0^{\tau^Y} e^{-\rho s}(\pi_s r_G - (1-\pi_s)b)\,ds + e^{-\rho\tau^Y}\pi_{\tau^Y}\Pi^\dagger(\tau^Y + t)\right],$$

Note that the firm's profits from an optimal contract of the full problem are in general not $\Pi^*(0)$, as that auxiliary problem conditions on the project not already being Bad at time zero. Rather, the profits of an optimal contract may be written $\Pi^{**} = \Pi^*(0-)$, where $\widetilde{H}(s;0-) = H(s)$.

Fixing a function $\Pi^\dagger(\cdot)$ on the rhs, the optimal stopping problem characterizing each

[15] I will assume in this setting that the project cannot be operated without an expert.

$\Pi^*(t)$ retains its basic recursive structure in the state variable $X_s = (\pi_s, s)$. In particular, a time-dependent termination threshold will continue to be optimal. The only change comes when computing the option value of continuing the project, where one must use a termination payoff of $\pi_s \Pi^\dagger(t+s)$ instead of 0 as in the benchmark problem. Unsurprisingly, this positive termination payoff will push up the optimal threshold at all times compared to the setting with no replacement. Also, the smaller is K, the higher this threshold will be.

The complex part of this exercise is solving what is essentially a fixed-point problem, whereby the continuation profit function $\Pi^\dagger(\cdot)$ must be chosen to induce a solution to the optimal stopping problem for each t consistent with the original choice of $\Pi^\dagger(\cdot)$. I will illustrate this fact for the special case $H(t) = 1 - \exp(-\alpha t)$, where the state is Good with probability 1 at time 0 and the state transition rate is homogeneous. In this case $\Pi^{**} = \Pi^*(0)$ and $\Pi^*(t) = \Pi^*(0)$ for all time. Define a function f by

$$f(x) \equiv \sup_{\tau^Y} \mathbb{E}\left[\int_0^{\tau^Y} e^{-\rho s}(\pi_s r_G - (1-\pi_s)b)\, ds + e^{-\rho \tau^Y} \pi_{\tau^Y} \max\{x - K, 0\}\right].$$

Then Π^{**} is a solution to the fixed point problem $x = f(x)$. More precisely, it should be the largest such fixed point in case there are several, but the following lemma ensures that there is exactly one:

Lemma 9. *There is exactly one solution to $x = f(x)$.*

An immediate corollary is that it is optimal to hire a replacement expert if and only if K is less than the firm's optimal profits $\Pi[\tau^*]$ in the benchmark problem without replacement. This can be seen simply by noting that when $K \geq \Pi[\tau^*]$, then $\Pi[\tau^*] = f(\Pi[\tau^*])$, while $f(\Pi[\tau^*]) > \Pi[\tau^*]$ when $K < \Pi[\tau^*]$.

When $K < \Pi[\tau^*]$, the optimal contract may be solved iteratively, as follows. First guess x and compute $f(x)$. If $f(x) > x$, then x was chosen too high, and the true value of Π^{**} must lie below x. And conversely if $f(x) < x$, the true value of Π^{**} lies above x. By repeatedly guessing x and readjusting, the full problem may be solved numerically.

This solution may be easy adapted to the case where $H(0) > 0$, by first solving for $\Pi^*(0)$ as the fixed point of $x = f(x)$, and then solving one more optimal stopping problem for Π^{**} using the true state transition distribution and $\Pi^*(0)$ as the termination payoff. In this case $\Pi^{**} < \Pi^*(0)$ given the cost of compensating the expert when the project is bad immediately. However, conditional on employing the initial expert at all, the optimal termination threshold for the first and all subsequent experts will be identical. The only possible difference in treatment is that if $H(0)$ is sufficiently large, the first expert will be asked to advise on whether the project is initially viable and then fired immediately no matter his response, in order to avoid costly incentive payments.

For the general inhomogeneous state transition setting, an analogous fixed point problem must be solved. However, in this case the entire function $\Pi^*(\cdot)$ must be guessed at once, and then checked against the resulting optimized profits at each time. The iterative procedure outlined above must then be replaced by more sophisticated techniques of value function iteration.

6.4 Busywork

Another important assumption of my model is that the expert's flow rents are unpledgeable, and in particular can't be dissipated by verifiable activity which is costly to the expert. I now relax this assumption and show how my analysis can be adapted to accommodate the presence of a dissipative "busywork" technology which imposes costs on both the firm and the expert.[16]

Suppose that the firm has access to a technology which can impose a utility cost of $k \in [0, b]$ on the expert at the expense of a reduction $C(k)$ to the firm's flow profits.[17] (The technology can be operated only while the project is active.) C is assumed to be twice continuously differentiable, strictly increasing, and strictly convex, with $C(0) = 0$. The firm can commit to a schedule of busywork along with payment and termination processes.

The virtual profit function derived in Proposition 1 is readily adapted to this setting. At any time t following any history in which the project has not yet expired, the firm receives flow profits r_G and incurs a busywork cost $C(k_t)$. Meanwhile in any history in which the project has expired, the firm incurs a virtual flow cost stemming from the terminal incentive payment, which must compensate the expert for any flow benefits b minus any busywork k_t that would have been imposed had the project continued.[18] Thus the firm's virtual profit function becomes

$$\mathbb{E}\left[\int_0^{\tau^Y} e^{-\rho t}(\pi_t(r_G - C(k_t)) - (1-\pi_t)(b-k_t))\,dt\right].$$

[16] I thank Jeff Ely for suggesting this analysis.

[17] I assume that the firm cannot impose more busywork on the expert at any moment in time than he receives in flow benefits. Otherwise the expert's participation constraint might be violated, complicating the analysis.

[18] The result that the firm stops the project immediately after the expert's report is robust to the busywork technology. For any IC-B contract imposing busywork after the report remains IC-B if no busywork is imposed while improving the firm's profits. So without loss busywork can be assumed to be imposed only prior to a report, in which case the reasoning for never stopping the project late continues to hold.

The optimal amount of busywork k_t^* at any time $t < \tau^Y$ can then be read off of the integrand:

$$k_t^* = \begin{cases} 0, & C'(0) \geq (1-\pi_t)/\pi_t \\ (C')^{-1}\left(\frac{1-\pi_t}{\pi_t}\right), & C'(b) > (1-\pi_t)/\pi_t > C'(0) \\ b, & (1-\pi_t)/\pi_t \geq C'(b). \end{cases}$$

Remark. *The optimal busywork process k^* can be chosen to be independent of the stopping time τ^Y, and for this process there exists a continuous, decreasing function $\kappa : (0,1] \to [0,b]$ such that $\kappa(1) = 0$ and $k_t^* = \kappa(\pi_t)$ for all time.*

Because the optimal amount of busywork is a function of current naive beliefs, the problem retains its recursive structure. Further, the proof of Proposition 2 can be adapted to show that an optimal termination rule continues to be a time-dependent threshold rule in naive beliefs.[19] And because imposition of busywork improves flow profits at every belief level, the optimal threshold is lower at each moment in time with busywork than without. Finally, note that busywork can lead to efficient project termination, i.e. a termination rule $\tau^Y = \infty$, if and only if $C(b) < r_G$. For when this inequality holds, the firm's virtual profits are always positive at the optimal busywork level, and otherwise they continue to be negative for sufficiently low beliefs.

7 Conclusion

In this paper I ask how a firm should optimally elicit expert advice on when to terminate a project which may eventually become unviable, if the expert accrues private benefits from prolonging the project as much as possible. Assuming the expert is capital-constrained and cannot buy the project, the firm must compensate him for reporting bad news which leads to early project termination. As a result, the firm prefers to commit to limit the lifespan of the project in order to economize on incentive payments.

I fully characterize the firm's optimal contract when it can imperfectly monitor the state of the project by observing its incremental output flow, under very general assumptions on the output and state transition processes. The optimal contract can be elegantly character-

[19]Interestingly, this is true even though virtual flow profits are no longer generally monotone in π under the optimal busywork rule κ. One way to understand this result is to note that by the envelope theorem, the derivative of virtual flow profits in π_t is $r_G - C(\kappa(\pi_t)) + b - \kappa(\pi_t)$, and this is positive whenever flow profits are non-negative. Thus virtual flow profits under the optimal busywork schedule cross zero exactly once. In other words, the option value of waiting for news about the project below the breakeven point diminishes the lower beliefs drop, implying optimality of a threshold rule.

ized - the expert is asked to report when the project should end, at which point the project is immediately terminated and a lump-sum termination payment is made to the expert. This payment is set to exactly compensate the expert for the private benefits he gives up by not hiding his knowledge and allowing the project to operate as long as possible. The firm also sets a stochastic public deadline, at which point the project is terminated even if the expert has not advised that it be. This deadline is optimally a time-dependent threshold rule in the firm's "naive beliefs", the beliefs the firm would have formed about the current state of the project had it learned only from output without the advice of the expert. An important implication of this deadline structure is that the forecasted time to termination exhibits time-varying sensitivity to output, with increased sensitivity to runs of output the closer the contract is to termination.

This characterization allows for a very clear analysis of the value of expert advice. With an expert, the firm completely eliminates "false negatives", i.e. operating the project past its expiration date, while partially mitigating "false positives", i.e. premature termination of the project. These gains lead to more efficient ex post project operation in every state of the world - no matter the actual state switch time and realization of output, the expert's advice yields higher net project output. My solution is also elegant and flexible enough to permit easy analysis of important extensions, including an expert with initial capital, replacement of experts, and the imposition of busywork to make the expert's position less cushy.

One interesting avenue for future work would be incorporating richer private information by the expert. In my model I assume that the expert gains private knowledge of the project's expiration only when it occurs. If the expert possessed some knowledge of the likely project lifetime earlier on, for instance at the start of the project, then the firm could attempt to elicit this information by offering contracts which verify reports by conditioning on output observed after a report is made. This design consideration, reminiscent of the verification contracts of Varas (2017), is absent in my setting.

Another important extension would be to combine the existing reporting problem with moral hazard. In particular, in many applications output may be generated by actions taken by the expert toward completing the project. If the expert shirks, then on the one hand current output drops, but on the other hand he leaves additional productive work for an opportune future time. The ability to strategically time the completion of work would interact with the variable sensitivity of the contract to output recommended by my analysis in deleterious ways, since the expert would benefit by delaying work from periods where the contract is less output-sensitive to periods where it is more sensitive. Including moral hazard would therefore add an interesting and realistic dimension to the design of optimal incentives.

References

Baron, D.P. and D. Besanko (1984). "Regulation, Asymmetric Information, and Auditing". In: *The RAND Journal of Economics* 15.4, pp. 447–470.

Battaglini, M. (2005). "Long-Term Contracting with Markovian Consumers". In: *American Economic Review* 95.3, pp. 637–658.

Besanko, D. (1985). "Multi-period contracts between principal and agent with adverse selection". In: *Economics Letters* 17.1, pp. 33–37.

Courty, P. and H. Li (2000). "Sequential Screening". In: *The Review of Economic Studies* 67.4, pp. 697–717.

Davis, Mark H. A. (2005). "Martingale Representation and All That". In: *Advances in Control, Communication Networks, and Transportation Systems: In Honor of Pravin Varaiya.* Ed. by E. H. Abed. Birkhäuser Boston, pp. 57–68.

Eso, Péter and Balázs Szentes (2017). "Dynamic contracting: an irrelevance theorem". In: *Theoretical Economics* 12.1, pp. 109–139.

Garrett, D.F. and Alessandro Pavan (2012). "Managerial Turnover in a Changing World". In: *Journal of Political Economy* 120.5, pp. 879–925.

Green, Brett and Curtis R. Taylor (2016). "Breakthroughs, Deadlines, and Self-Reported Progress: Contracting for Multistage Projects". In: *American Economic Review* 106.12, pp. 3660–99.

Grenadier, Steven R., Andrey Malenko, and Nadya Malenko (2016). "Timing Decisions in Organizations: Communication and Authority in a Dynamic Environment". In: *American Economic Review* 106.9, pp. 2552–81.

Guo, Yingni (2016). "Dynamic Delegation of Experimentation". In: *American Economic Review* 106.8, pp. 1969–2008.

Kallenberg, Olav (1997). *Foundations of Modern Probability.* Springer New York.

Karatzas, I. and S.E. Shreve (1991). *Brownian Motion and Stochastic Calculus.* Springer New York.

Kruse, T. and P. Strack (2015). "Optimal stopping with private information". In: *Journal of Economic Theory* 159, pp. 702–727.

Lazear, Edward P. (1983). "Pensions as Severance Pay". In: *Financial Aspects of the United States Pension System.* Ed. by Zvi Bodie and John B. Shoven. University of Chicago Press, pp. 57–90.

Milgrom, Paul and Ilya Segal (2002). “Envelope Theorems for Arbitrary Choice Sets”. In: *Econometrica* 70.2, pp. 583–601.

Myerson, Roger B. (1981). “Optimal auction design”. In: *Mathematics of Operations Research* 6.1, pp. 58–73.

Nualart, David and Wim Schoutens (2000). “Chaotic and predictable representations for Lévy processes”. In: *Stochastic Processes and their Applications* 90.1, pp. 109–122.

Pavan, Alessandro, Ilya Segal, and J. Toikka (2014). “Dynamic Mechanism Design: A Myersonian Approach”. In: *Econometrica* 82.2, pp. 601–653.

Varas, Felipe (2017). “Managerial Short-Termism, Turnover Policy, and the Dynamics of Incentives”. In: *The Review of Financial Studies*.

Williams, Noah (2011). “Persistent Private Information”. In: *Econometrica* 79.4, pp. 1233–1275.

Appendices

A Formal contract construction

The probability space $(\Omega, \mathcal{F}, \mathbb{P})$ on which the model is built encodes only the realizations of the exogenous processes Y and Λ. It doesn't accommodate the endogenous reported state process θ' generated by the expert's reporting policy. As a result, a contract (Φ, τ) cannot be formally considered a pair of random elements on the probability space. Rather, it is a mapping from reporting policies into pairs of random elements. The following definition recasts Definition 2 in this more rigorous formalism.

Definition A.1. *A* revelation contract $\mathcal{C} = (\Phi, \tau)$ *is a family of payment processes* $\Phi[T] \geq 0$ *and termination times* $\tau[T]$ *for each* $T \in \mathbb{R}_+ \cup \{\infty\}$*, such that:*

- *For every* $t \in \mathbb{R}_+$*, the maps* $(T, \omega) \mapsto \Phi[T]_t(\omega)$ *and* $(T, \omega) \mapsto \mathbf{1}\{\tau[T](\omega) \leq t\}$ *are* $\mathcal{B}(\mathbb{R}_+ \cup \{\infty\}) \otimes \mathcal{F}_t^Y$*-measurable,*
- *Each* $\Phi[T]$ *is right-continuous, increasing, and satisfies* $\Phi[T]_t = \Phi[T]_{\tau[T]}$ *for every* $t > \tau[T]$*,*
- *For every* $T, T' \in \mathbb{R}_+ \cup \{\infty\}$ *and* $t < \min\{T, T'\}$*,* $\Phi[T]_t = \Phi[T']_t$ *and* $\mathbf{1}\{\tau[T] \leq t\} = \mathbf{1}\{\tau[T'] \leq t\}$.

This definition characterizes how a contract maps deterministic reporting times into payment and termination policies. The mapping for a general reporting policy Λ' is then defined by

$$\Phi[\Lambda'](\omega) \equiv \Phi[\Lambda'(\omega)](\omega), \quad \tau[\Lambda'](\omega) \equiv \tau[\Lambda'(\omega)](\omega).$$

The joint measurability requirement in Definition A.1 ensures that this construction yields measurable, $\mathbb{F}$-adapted processes for any choice of Λ'. The final requirement in Definition A.1 ensures that (Φ, τ) is "$\mathbb{F}'$-adapted" in the sense of not conditioning payments or termination on a reported switch which hasn't yet arrived.

The notation $\mathbb{E}^{\Lambda'}$ used in the body of the paper is shorthand for expectations wrt uncertainty induced by the processes $\Phi[\Lambda']$ and $\tau[\Lambda']$ wherever Φ and τ appear in the interior of the expectation. For instance,

$$\mathbb{E}^{\Lambda'}\left[\int_0^{\tau} e^{-\rho t}(b\,dt + d\Phi_t)\right] = \mathbb{E}\left[\int_0^{\tau[\Lambda']} e^{-\rho t}(b\,dt + d\Phi[\Lambda']_t)\right].$$

B The optimal belief threshold for termination

Calculating the optimal belief threshold for termination of the project amounts to solving a relatively standard optimal stopping problem. In this Appendix I briefly outline the standard procedure for characterizing the optimal threshold.

The first step is to observe that the firm's continuation profit process

$$\Pi_t = \mathbb{E}_t^Y \left[\int_t^{\tau^*} e^{-\rho t} (\pi_t r_G - (1 - \pi_t) b) \, dt \right]$$

is recursive in the state variable (t, π_t), and so can be written $\Pi_t = V(\pi_t, t)$ for some value function V. Then V is conjectured to be sufficiently smooth that a version of Ito's lemma applies, from which an HJB equation characterizing V can be derived.

To apply Ito's lemma a differential representation for the evolution of π_t is needed. Standard results from filtering theory give us the desired form, as follows. Let $\Delta Y_t \equiv Y_t - Y_{t-}$ be the jump part of Y and $Y_t^c \equiv Y_t - \sum_{s \leq t} \Delta Y_s$ be its continuous part, and for each $\theta \in \{G, B\}$ define $\widetilde{r}_\theta \equiv r_\theta - \sum_{i=1}^n d_i \lambda_i^\theta$. Let

$$\overline{Z}_t \equiv \sigma^{-1} \left(Y_t^c - \int_0^t (\pi_{s-} \widetilde{r}_G + (1 - \pi_{s-}) \widetilde{r}_B) \, ds \right)$$

and

$$\overline{N}_i(t) \equiv \sum_{s \leq t} \mathbf{1}\{\Delta Y_s = d_i\} - \int_0^t (\pi_{s-} \lambda_i^G + (1 - \pi_{s-}) \lambda_i^B) \, ds$$

be the innovation processes; a standard result from filtering theory is that with respect to the filtration $\mathbb{F}^Y$, $\overline{Z}$ is a standard Brownian motion, while each $\overline{N}_i$ is a compensated Poisson process with rate process $\overline{\lambda}_i(\pi_t) \equiv \pi_t \lambda_i^G + (1 - \pi_t) \lambda_i^B$. Then π evolves according to the SDE

$$d\pi_t = -\frac{\pi_{t-}}{1 - H(t-)} dH(t) + \frac{\widetilde{r}_G - \widetilde{r}_B}{\sigma} \pi_{t-} (1 - \pi_{t-}) \, d\overline{Z}_t + \sum_{i=1}^n \frac{\lambda_i^G - \lambda_i^B}{\overline{\lambda}_i(\pi_{t-})} \pi_{t-} (1 - \pi_{t-}) \, d\overline{N}_i(t).$$

To ensure a smooth value function, I will assume that H is an absolutely continuous function with continuous hazard rate function $\alpha(t) \equiv H'(t)/(1 - H(t))$.

Ito's lemma can then be used to derive the HJB equation

$$\begin{aligned}\rho V(x,t) = x r_G - (1-x)b - \left(\alpha(t)x + x(1-x)\sum_{i=1}^{n}(\lambda_i^G - \lambda_i^B)\right) V_x(x,t) \\ + \frac{1}{2}\left(\frac{\widetilde{r}_G - \widetilde{r}_B}{\sigma}\right)^2 x^2(1-x)^2 V_{xx}(x,t) \\ + \sum_{i=1}^{n}\overline{\lambda}_i(x)\left(V\left(\frac{\lambda_i^G x}{\overline{\lambda}_i(x)}, t\right) - V(x,t)\right) + V_t(x,t).\end{aligned}$$

Along the optimal stopping boundary $\underline{\pi}(t)$, the value function should satisfy the value matching equation $V(\underline{\pi}(t),t) = 0$ and the smooth pasting condition $V_x(\underline{\pi}(t),t) = 0$. Characterizing the optimal termination policy then reduces to the purely mathematical problem of solving a free boundary problem, by finding a function $\underline{\pi}$ for which the HJB equation with associated boundary conditions has a solution. Note that in the time-homogeneous case with $\alpha(t) = \alpha$ a constant, the value function is independent of time and all conditioning on t in the HJB equation, along with the term V_t, may be dropped.

C News sensitivity of the optimal deadline

In this Appendix I show how to generalize the characterization of the new sensitivity of the optimal public deadline to accommodate output processes with jumps. The use of the expected time to termination is a less convenient forecasting measure in jump models, because this measure does not drift downward at rate 1 in general. Rather, the drift rate is adjusted to account for the information provided by lack of jumps over a given time interval, with this additional drift compensated by occasional jumps up or down in response to output jumps.

To adjust for this fact, I will define the *forecasted* time to termination Δ for the general model by

$$\Delta_t = \widetilde{E}_t^G[\tau^*] - t,$$

where $\mathbb{E}^G$ takes expectations over output paths wrt a measure $\widetilde{\mathbb{P}}^G$ under which the state is always Good and $N_i^G = 0$ for all i. Essentially, this measure conditions on the state never switching and no jumps ever arriving. In the absence of jumps, Δ drifts downward at rate 1, with random variation due to fluctuations in the continuous production of output. While Δ is not the expected time to termination except in the Brownian case, it is a reasonable benchmark termination date which is subject to revision in response to the arrival of an output jump. When jumps in the aggregate are good news about the state, Δ will tend to experience net upward revisions over the lifetime of the contract, while when jumps in the aggregate are bad news the opposite is true.[20]

To characterize Δ, let $\Delta Y_t \equiv Y_t - Y_{t-}$ be the jump part of Y and $Y_t^c \equiv Y_t - \sum_{s\le t} \Delta Y_s$ be its continuous part, and define $\widetilde{r}_\theta \equiv r_\theta - \sum_{i=1}^n d_i \lambda_i^\theta$. Let

$$\overline{Z}_t \equiv \sigma^{-1}\left(Y_t^c - \int_0^t (\pi_{s-}\widetilde{r}_G + (1-\pi_{s-})\widetilde{r}_B)\, ds\right)$$

and

$$\overline{N}_i(t) \equiv \sum_{s\le t} \mathbf{1}\{\Delta Y_s = d_i\} - \int_0^t (\pi_{s-}\lambda_i^G + (1-\pi_{s-})\lambda_i^B)\, ds$$

be the innovation processes for each component of output. Under the measure $\mathbb{P}$ and filtration $\mathbb{F}^Y$, $\overline{Z}$ is a standard Brownian motion and each $\overline{N}_i$ is an inhomogeneous compensated Poisson counting process with rate $\overline{\lambda}_i(\pi_t) \equiv \pi_t\lambda_i^G + (1-\pi_t)\lambda_i^B$. By standard results from optimal filtering, π evolves according to the SDE

$$d\pi_t = -\alpha\pi_{t-}\, dt + \frac{\widetilde{r}_G - \widetilde{r}_B}{\sigma}\pi_{t-}(1-\pi_{t-})\, d\overline{Z}_t + \sum_{i=1}^n \frac{\lambda_i^G - \lambda_i^B}{\overline{\lambda}_i(\pi_{t-})}\pi_{t-}(1-\pi_{t-})\, d\overline{N}_i(t).$$

This SDE may be rewritten

$$\begin{aligned} d\pi_t = &\left(-\alpha\pi_{t-} + \frac{\widetilde{r}_G - \widetilde{r}_B}{\sigma}\pi_{t-}(1-\pi_{t-})^2 - \sum_{i=1}^n (\lambda_i^G - \lambda_i^B)\pi_{t-}(1-\pi_{t-})\right) dt \\ &+ \frac{\widetilde{r}_G - \widetilde{r}_B}{\sigma}\pi_{t-}(1-\pi_{t-})\, d\overline{Z}_t^G + \sum_{i=1}^n \frac{\lambda_i^G - \lambda_i^B}{\overline{\lambda}_i(\pi_{t-})}\pi_{t-}(1-\pi_{t-})\, dN_i(t), \end{aligned}$$

where $Z_t^G \equiv \sigma^{-1}(Y_t^c - r_G t)$ and $N_i(t) \equiv \sum_{s\le t} \mathbf{1}\{\Delta Y_s = d_i\}$. (This expression also accommodates pure jump processes with no Brownian term. In that case the state is immediately detectable unless $\widetilde{r}_G = \widetilde{r}_B$, as otherwise the drift rates in the absence of jumps would differ i the two states. So the non-trivial pure jump process case can be handled by setting the signal-to-noise ratio $(\widetilde{r}_G - \widetilde{r}_B)/\sigma = 0$ in the expressions above.)

Under the measure $\widetilde{\mathbb{P}}^G$, Z^G is a standard Brownian motion while each N_i is identically zero. A standard hitting time calculation shows that Δ satisfies $\Delta_t = \xi(\pi_t)$ for a function $\xi : [0,1]$ such that $\xi(x) = 0$ for $x < \underline{\pi}^*$, while on $[\underline{\pi}^*, 1]$ ξ is a C^2 function satisfying the value

[20] If jumps in the aggregate are very bad news about the state, then it may be that $\Delta = \infty$ due to large upward belief revisions over intervals when jumps are absent. To avoid this possibility, I will assume that $\alpha > -\sum_{i=1}^n (\lambda_i^G - \lambda_i^B)$.

function ODE

$$0 = 1 + \left(-\alpha x + \left(\frac{\widetilde{r}_G - \widetilde{r}_B}{\sigma} \right)^2 x(1-x)^2 - \sum_{i=1}^{n} (\lambda_i^G - \lambda_i^B) x(1-x) \right) \xi'(x) + \frac{1}{2} \left(\frac{\widetilde{r}_G - \widetilde{r}_B}{\sigma} \right)^2 x^2 (1-x)^2 \xi''(x)$$

with boundary condition $\xi(\underline{\pi}^*) = 0$. Ito's lemma then implies that

$$d\Delta_t = -\,dt + \frac{\widetilde{r}_G - \widetilde{r}_B}{\sigma} \pi_{t-}(1-\pi_{t-}) \xi'(\pi_{t-})\, d\overline{Z}_t^G + \sum_{i=1}^{n} \left\{ \xi \left(\pi_{t-} + \frac{\lambda_i^G - \lambda_i^B}{\overline{\lambda}_i(\pi_{t-})} \pi_{t-}(1-\pi_{t-}) \right) - \xi(\pi_{t-}) \right\} dN_i(t).$$

The following lemma generalizes the sensitivity monotonicity result of Lemma 6 to models with jumps.

Lemma C.1. *Suppose* $\underline{\pi}^* > 0$ *and* $(\widetilde{r}_G - \widetilde{r}_B)/\sigma \neq 0$. *Then the function* $\phi : [\underline{\pi}^*, 1] \to \mathbb{R}_+$ *defined by* $\phi_0(x) = x(1-x)\xi'(x)$ *is strictly decreasing and satisfies* $\phi_0(1) = 0$.

Proof. We first establish that ξ is strictly monotone. Note that for every $x' > x \geq \underline{\pi}^*$,

$$\xi(x') = \widetilde{\mathbb{E}}^G[\tau(x' \to x)] + \xi(x),$$

where $\tau(x' \to x)$ is the first time posterior beliefs drop below x for the posterior belief process started at x'. Under the measure $\widetilde{\mathbb{P}}^G$, the posterior belief process is pathwise continuous a.s. Hence $\tau(x' \to x) > 0$ a.s., implying $\xi(x') > \xi(x)$.

The remainder of the proof is a straightforward generalization of the proof of Lemma 6. The value function ODE can be used to show that

$$\phi_0'(x) = \frac{\left(\alpha x + \left(\Delta\lambda - \frac{1}{2} SNR^2 x(1-x)\right)\right) \xi'(x) - 1}{\frac{1}{2} SNR^2 x(1-x)},$$

where $SNR = (\widetilde{r}_G - \widetilde{r}_B)/\sigma$ and $\Delta\lambda = \sum_{i=1}^{n} (\lambda_i^G - \lambda_i^B)$. This derivative is trivially strictly negative if $\Delta\lambda < \frac{1}{2} SNR^2$ and $x \leq 1 - \frac{\alpha}{\frac{1}{2} SNR^2 - \Delta\lambda}$. So the interesting cases are $x > 1 - \frac{\alpha}{\frac{1}{2} SNR^2 - \Delta\lambda}$ for $\Delta\lambda < \frac{1}{2} SNR^2$, and all x for $\Delta\lambda \geq \frac{1}{2} SNR^2$. In either case, define

$$\zeta(x) = \frac{1}{\alpha x + \left(\Delta\lambda - \frac{1}{2} SNR^2\right) x(1-x)}.$$

Some algebra shows that

$$\begin{aligned} &\zeta'(x) - \frac{\left(\alpha x - SNR x(1-x)^2 + \Delta\lambda x(1-x)\right)\zeta'(x) - 1}{\frac{1}{2}SNR^2 x(1-x)} \\ =& \frac{\alpha}{x(1-x)(\alpha + \left(\Delta\lambda - \frac{1}{2}SNR^2\right)(1-x))^2} > 0, \end{aligned}$$

proving that ζ is a subsolution to the value function ODE satisfying $\zeta(1) = \xi'(1) = 1/\alpha$ and thus that $\xi'(x) < \zeta(x)$ on the desired interval, establishing $\phi_0'(x) < 0$. □

Thus the forecasted time to termination has a sensitivity to variation in continuous runs of output which is increasing as termination draws near, just as in the Brownian case. Of course, the *direction* of the sensivity depends on the distribution of jumps - it is possible that $\widetilde{r}_G < \widetilde{r}_B$ if positive jumps are more likely in the Good state and negative jumps are more likely in the Bad state. In that case high continuous runs of output are actually *bad* news about the state, because they are associated with fewer future upward output jumps and more future downward jumps.

Finally, I consider the reaction of the forecasted deadline to jumps. From the Ito's lemma expansion of Δ_t, a jump of size d_i at time t leads to a revision of Δ of size

$$\Delta_t - \Delta_{t-} = \xi\left(\pi_{t-} + \frac{\lambda_i^G - \lambda_i^B}{\overline{\lambda}_i(\pi_{t-})}\pi_{t-}(1-\pi_{t-})\right) - \xi(\pi_{t-}).$$

Naturally, jumps more likely to occur when the state is Good lead to upward revisions of the deadline, while jumps more likely to occur when the state is Bad lead to downward revisions. (Jumps that arrive at the same rate in each state do not lead to belief revisions are not conditioned on in an optimal contract.)

The sensitivity of the deadline to jumps has a more complex relationship to current beliefs than the responsiveness to continuous runs of output. In particular, consider the polar extremes of perfectly revealing jumps. Suppose there exists a jump size which occurs only when the state is Good, i.e. $\lambda_i^B = 0$. In that case $\Delta_t - \Delta_{t-} = \xi(1) - \xi(\pi_{t-})$ whenever the jump occurs, and thus the magnitude of the deadline revision is decreasing the further the project is from the deadline. Conversely, suppose the jump occurs only when the state is Bad, i.e. $\lambda_i^G = 0$. In that case $\Delta_t - \Delta_{t-} = -\xi(\pi_{t-})$, and so the magnitude of the deadline revision is increasing the further the project is from the deadline. In general for jumps which don't perfectly reveal the state, deadline revisions when the jump occurs may be non-monotone in the current state of the project.

The following lemma characterizes a sufficient condition for jump sizes to be decreasing in distance to the deadline provided the deadline is sufficiently far away. The lemma requires that jumps in the aggregate not be excessively good news about the state, in the sense that the aggregate arrival rate of jumps $\sum_{i=1}^{n} \lambda_i^\theta$ not be too much larger when $\theta = G$ than when $\theta = B$.

The lemma statement uses the following pieces of notation. For a jump of size d_i, define

$$\phi_i(x) = \left| \xi\left(x + \frac{\lambda_i^G - \lambda_i^B}{\overline{\lambda}_i(x)} x(1-x) \right) - \xi(x) \right|$$

to be the magnitude of the revision to the forecasted termination time upon arrival of a jump of size d_i when beliefs prior to the jump were x. Let $\Delta\lambda \equiv \sum_{i=1}^{n}(\lambda_i^G - \lambda_i^B)$ be the difference in the aggregate arrival rate of jumps across states.

Lemma C.2. *Suppose that $\alpha > \Delta\lambda$. Then for each i, $\phi_i'(x) < 0$ whenever $\frac{x}{1-x} \geq \sqrt{\lambda_i^B/\lambda_i^G}$.*

Proof. To prove the lemma, we first show that under either of the conditions in the lemma statement, the value function ξ satisfies $\xi''(x) < 0$ for all $x < 1$. So conditional on a belief revision of a given size $\Delta\pi$, the change to the forecasted deadline

$$\xi(\pi + \Delta\pi) - \xi(\pi) = \int_{\pi}^{\pi+\Delta\pi} \xi'(x)\, dx$$

is smaller the larger is π. We then show that the size of a belief revision from arrival of a jump diminishes in π whenever $\pi/(1-\pi) \geq \sqrt{\lambda_i^B/\lambda_i^G}$. As ξ is strictly increasing (see the proof of Lemma C.1), the magnitude of the deadline revision following an output jump is therefore strictly decreasing in beliefs over the range claimed.

To establish concavity of the value function, first consider the case $SNR = (\widetilde{r}_G - \widetilde{r}_B)/\sigma = 0$. In that case the value function ODE implies

$$\xi'(x) = \frac{1}{\alpha x + \Delta\lambda x(1-x)},$$

where $\Delta\lambda \equiv \sum_{i=1}^{n}(\lambda_i^G - \lambda_i^B)$. If $\Delta\lambda \geq 0$, then this expression is always strictly positive and well-defined. If $\Delta\lambda < 0$, then the additional assumption that $\alpha > -\Delta\lambda$ (see footnote 20) ensures the same result.

Differentiating yields

$$\xi''(x) = -\frac{\alpha + \Delta\lambda(1-2x)}{(\alpha x + \Delta\lambda x(1-x))^2}.$$

Suppose first that $\Delta\lambda \leq 0$. Then the numerator of this expression is minimized at $x = 0$, where it reduces to $\alpha + \Delta\lambda$. By assumption $\alpha > -\Delta\lambda$, so the numerator is always strictly positive, meaning $\xi''(x) < 0$. On the other hand if $\Delta\lambda > 0$, then the numerator is maximized at $x = 1$, where it reduces to $\alpha - \Delta\lambda$. By assumption this difference is strictly positive, so in all cases $\xi''(x) < 0$.

So assume $SNR > 0$. note that for $x < 1$ the value function ODE can be re-arranged to obtain

$$\xi''(x) = \frac{(\alpha x - SNR^2 x(1-x)^2 + \Delta\lambda x(1-x))\xi'(x) - 1}{\frac{1}{2}SNR^2 x^2(1-x)^2}.$$

This expression is trivially strictly negative whenever $\beta(x) = \alpha - SNR^2(1-x)^2 + \Delta\lambda(1-x) \leq 0$. So if we can show that

$$\xi'(x) < \zeta(x) = \frac{1}{\alpha x - SNR^2 x(1-x)^2 + \Delta\lambda x(1-x)}$$

whenever $\beta(x) > 0$, we'll have established the desired result. Note that β is a strictly concave function, and $\beta(1) = \alpha > 0$, so the subset of $[0,1]$ on which β is strictly positive is an interval $(x^*, 1]$, with $x^* = 0$ if it is positive everywhere on $(0,1]$. It is therefore enough to establish $\xi'(x) < \zeta(x)$ on the interval $(x^*, 1)$.

I prove this inequality by showing that ζ is a subsolution to the value function ODE on $[x^*, 1]$. For this it is sufficient that $\zeta'(x) < 0$ for $x \in (x^*, 1)$, as the remaining terms in the value function ODE collectively vanish when ζ is inserted. Computing this derivative yields

$$\zeta'(x) = -\frac{\alpha - SNR^2(1-x)(1-3x) + \Delta\lambda(1-2x)}{(\alpha x - SNR^2 x(1-x)^2 + \Delta\lambda x(1-x))^2}.$$

The numerator $\gamma(x) = \alpha - SNR^2(1-x)(1-3x) + \Delta\lambda(1-2x)$ is a strictly concave function satisfying $\gamma(1) = \alpha - \Delta\lambda$, which by assumption is strictly positive. If $\gamma(x^*) \geq 0$, then γ is strictly positive on the interval $(x^*, 1)$, establishing that $\zeta'(x) < 0$ on this interval as desired. Note that γ may be written

$$\gamma(x) = \beta(x) + x(2SNR^2(1-x) - \Delta\lambda).$$

If $x^* = 0$, then $\gamma(x^*) \geq \beta(x^*) \geq 0$, so assume $x^* > 0$. In this case $\gamma(x^*) = x^*(2SNR^2(1 - x^*) - \Delta\lambda)$, and it suffices to prove that $1 - x^* \geq \Delta\lambda/(2SNR^2)$. Recall that x^* is a positive root of the quadratic equation $\beta(x^*) = 0$, which when solved yields the unique positive root

$$1 - x^* = \frac{\Delta\lambda + \sqrt{(\Delta\lambda)^2 + 4SNR^2\alpha}}{2SNR^2} > \frac{\Delta\lambda}{2SNR^2}.$$

So $\gamma(x^*) \geq 0$ in all cases, establishing $\zeta'(x) < 0$ on $(x^*, 1)$ and thus showing ζ is a subsolution to the value function ODE on that interval. Then as $\zeta(1) = \xi'(1) = 1/\alpha$, it follows that $\xi'(x) < \zeta(x)$ on $(x^*, 1)$, as desired.

Having established strict concavity of the value function, we now return to calculating the size of the belief revision following arrival of a jump. The magnitude of the revision is

$$\kappa(x) = \frac{x(1-x)}{\bar{\lambda}_i(x)} = \frac{x(1-x)}{\lambda_i^G x + \lambda_i^B (1-x)}.$$

The derivative of this expression is

$$\kappa'(x) = \frac{-x^2 \lambda_i^G + (1-x)^2 \lambda_i^B}{\bar{\lambda}_i(x)^2}.$$

The magnitude of the belief revision is therefore decreasing in x whenever $x/(1-x) \geq \sqrt{\lambda_i^B/\lambda_i^G}$, as desired. □

D Proofs of results from the text

D.1 Proof of Lemma 1

IC-G and IC-B are clearly implied by incentive-compatibility. For the converse result, suppose a contract $\mathcal{C} = (\Phi, \tau)$ satisfies IC-G and IC-B, and fix an arbitrary $\mathbb{F}$-stopping time Λ'. Let $\underline{\Lambda'} \equiv \Lambda' \wedge \Lambda$ and $\overline{\Lambda'} \equiv \Lambda' \vee \Lambda$. Then by IC-G,

$$
\begin{aligned}
&\mathbb{E}^{\Lambda}\left[\int_0^\tau e^{-\rho t}\,(b\,dt + d\Phi_t)\right] \\
&\geq \mathbb{E}^{\underline{\Lambda'}}\left[\int_0^\tau e^{-\rho t}\,(b\,dt + d\Phi_t)\right] \\
&= \mathbb{E}^{\underline{\Lambda'}}\left[\mathbf{1}\{\Lambda' \leq \Lambda\}\int_0^\tau e^{-\rho t}\,(b\,dt + d\Phi_t)\right] + \mathbb{E}^{\underline{\Lambda'}}\left[\mathbf{1}\{\Lambda' > \Lambda\}\int_0^\tau e^{-\rho t}\,(b\,dt + d\Phi_t)\right].
\end{aligned}
$$

Now, the interior of the first expectation on the last line is identical under the policies Λ' and $\underline{\Lambda'}$, as on the set of states $\{\Lambda' \leq \Lambda\}$ the two policies coincide and so induce the same τ and Φ. Similarly, the interior of the second expectation is identical under the policies Λ and $\underline{\Lambda'}$, as on the set of states $\{\Lambda' > \Lambda\}$ the two policies coincide. Hence this inequality may be written

$$
\begin{aligned}
&\mathbb{E}^{\Lambda}\left[\int_0^\tau e^{-\rho t}\,(b\,dt + d\Phi_t)\right] \\
&\geq \mathbb{E}^{\Lambda'}\left[\mathbf{1}\{\Lambda' \leq \Lambda\}\int_0^\tau e^{-\rho t}\,(b\,dt + d\Phi_t)\right] + \mathbb{E}^{\Lambda}\left[\mathbf{1}\{\Lambda' > \Lambda\}\int_0^\tau e^{-\rho t}\,(b\,dt + d\Phi_t)\right].
\end{aligned}
$$

Subtracting the final term from both sides yields

$$
\mathbb{E}^{\Lambda}\left[\mathbf{1}\{\Lambda' \leq \Lambda\}\int_0^\tau e^{-\rho t}\,(b\,dt + d\Phi_t)\right] \geq \mathbb{E}^{\Lambda'}\left[\mathbf{1}\{\Lambda' \leq \Lambda\}\int_0^\tau e^{-\rho t}\,(b\,dt + d\Phi_t)\right].
$$

A very similar argument using $\overline{\Lambda'}$ and the IC-B constraint yields

$$
\mathbb{E}^{\Lambda}\left[\mathbf{1}\{\Lambda' > \Lambda\}\int_0^\tau e^{-\rho t}\,(b\,dt + d\Phi_t)\right] \geq \mathbb{E}^{\Lambda'}\left[\mathbf{1}\{\Lambda' > \Lambda\}\int_0^\tau e^{-\rho t}\,(b\,dt + d\Phi_t)\right].
$$

Summing these two inequalities results in

$$
\mathbb{E}^{\Lambda}\left[\int_0^\tau e^{-\rho t}\,(b\,dt + d\Phi_t)\right] \geq \mathbb{E}^{\Lambda'}\left[\int_0^\tau e^{-\rho t}\,(b\,dt + d\Phi_t)\right].
$$

Hence $\mathcal{C}$ is incentive-compatible.

D.2 Proof of Lemma 2

Fix an IC-B contract $\mathcal{C} = (\Phi, \tau)$. Define an $\mathbb{F}^Y$-adapted process ϕ by

$$\phi_t = \mathbb{E}_t^B \left[\int_{t \wedge \tau[t]}^{\tau[t]} e^{-\rho(s-t)} (b\, ds + d\Phi[t]_s) \right],$$

where $\mathbb{E}_t^B$ is as defined in Definition 7 and $\Phi[t], \tau[t]$ are as defined in Appendix A. Construct a new contract $\mathcal{C}' = (\Phi', \tau')$ by setting $\tau' = \tau \wedge \Lambda'$ and

$$\Phi'_t = \begin{cases} \Phi_t, & t < \tau', \\ \Phi_{\tau'} + \phi_{\Lambda'} \mathbf{1}\{\tau > \Lambda'\}, & t \geq \tau'. \end{cases}$$

Fix a reporting strategy $\Lambda' \geq \Lambda$, and let $U[\Lambda']$ be the expert's payoffs under (Φ, τ) and Λ', and similarly $U'[\Lambda']$ be his payoff under (Φ', τ') and Λ'. Then

$$\begin{aligned} U'[\Lambda'] &= \mathbb{E}^{\Lambda'} \left[\int_0^{\tau'} e^{-\rho t} (b\, dt + d\Phi'_t) \right] \\ &= \mathbb{E}^{\Lambda'} \left[\int_0^{\Lambda' \wedge \tau} e^{-\rho t} (b\, dt + d\Phi_t) + e^{-\rho \Lambda'} \phi_{\Lambda'} \mathbf{1}\{\tau > \Lambda'\} \right] \\ &= \mathbb{E} \left[\int_0^{\Lambda' \wedge \tau[\Lambda']} e^{-\rho t} (b\, dt + d\Phi[\Lambda']_t) + e^{-\rho \Lambda'} \phi_{\Lambda'} \mathbf{1}\{\tau[\Lambda'] > \Lambda'\} \right], \end{aligned}$$

where the last line makes explicit the dependence of τ and Φ on the reporting policy. Now, $\Lambda' \geq \Lambda$ means that by the definition of $\mathbb{E}_t^B$

$$\phi_{\Lambda'} = \mathbb{E} \left[\int_{\Lambda' \wedge \tau[\Lambda']}^{\tau[\Lambda']} e^{-\rho(t-\Lambda')} (b\, dt + d\Phi[\Lambda']_t) \;\middle|\; \mathcal{F}_{\Lambda'} \right].$$

Then applying the law of iterated expectations to the previous representation of $U'[\Lambda']$,

$$\begin{aligned} U'[\Lambda'] &= \mathbb{E} \left[\int_0^{\Lambda' \wedge \tau[\Lambda']} e^{-\rho t} (b\, dt + d\Phi[\Lambda']_t) + e^{-\rho \Lambda'} \mathbf{1}\{\tau[\Lambda'] > \Lambda'\} \int_{\Lambda' \wedge \tau[\Lambda']}^{\tau[\Lambda']} e^{-r(t-\Lambda')} (b\, dt + d\Phi[\Lambda']_t) \right] \\ &= \mathbb{E} \left[\int_0^{\Lambda' \wedge \tau[\Lambda']} e^{-\rho t} (b\, dt + d\Phi[\Lambda']_t) \right] = U[\Lambda']. \end{aligned}$$

So every $\Lambda' \geq \Lambda$ yields the same expected utility to the expert under $\mathcal{C}'$ as under $\mathcal{C}$, meaning $\mathcal{C}'$ is an IC-B contract.

Meanwhile, the firm's profits under $\mathcal{C}'$ and truthful reporting are

$$\begin{aligned}\Pi[\mathcal{C}'] &= \mathbb{E}\left[\int_0^{\Lambda\wedge\tau[\Lambda]} e^{-\rho t}(r_G\,dt - d\Phi'[\Lambda]_t) - e^{-\rho\Lambda}\phi_\Lambda \mathbf{1}\{\tau[\Lambda] > \Lambda\}\right] \\ &= \mathbb{E}\left[\int_0^{\Lambda\wedge\tau[\Lambda]} e^{-\rho t}(r_G\,dt - d\Phi[\Lambda]_t) - \int_{\Lambda\wedge\tau[\Lambda]}^{\tau[\Lambda]} e^{-\rho t}(b\,dt + d\Phi[\Lambda]_t)\right],\end{aligned}$$

where I have used the representation of $\phi_{\Lambda'}$ derived above and the law of iterated expectations to move from the first expression to the second. By comparison, the firm's profits under $\mathcal{C}$ are

$$\Pi[\mathcal{C}] = \mathbb{E}\left[\int_0^{\Lambda\wedge\tau[\Lambda]} e^{-\rho t} r_G\,dt + \int_{\Lambda\wedge\tau[\Lambda]}^{\tau[\Lambda]} e^{-\rho t} r_B\,dt - \int_0^{\tau[\Lambda]} e^{-\rho t}\,d\Phi[\Lambda]_t\right].$$

Thus

$$\Pi[\mathcal{C}'] - \Pi[\mathcal{C}] = -\mathbb{E}\left[\int_{\Lambda\wedge\tau[\Lambda]}^{\tau[\Lambda]} e^{-\rho t}(r_B + b)\,dt\right].$$

Then $r_B + b < 0$ implies $\Pi[\mathcal{C}'] \geq \Pi[\mathcal{C}]$, and this inequality is strict if $\mathbb{P}\{\tau[\Lambda] > \Lambda\} = \mathbb{P}^\Lambda\{\tau > \Lambda\} > 0$.

D.3 Proof of Lemma 3

Fix an IC-B contract $\mathcal{C} = (\Phi, \tau)$ satisfying $\tau \leq \Lambda'$ and $\mathbb{P}^\Lambda\{\tau < \infty\} = 1$. First note that τ may be decomposed as $\tau[T] = \tau[\infty] \wedge T$ for each $T \in \mathbb{R}_+ \cup \{\infty\}$, where $\tau[T]$ is as defined in Appendix A. For either the contracts terminates at the time of the report, i.e. $\tau[T](\omega) = T$, or it terminates prior to this time, in which case the eventual time of the report does not impact the termination time and $\tau[T](\omega) = \tau[\infty](\omega)$. The τ^Y in the lemma statement may then be taken to be $\tau^Y = \tau[\infty]$.

Define a new payment process $\widetilde{\Phi}$ by

$$\widetilde{\Phi}[T]_t = \begin{cases} \Phi[T]_t, & \min\{t, T\} < \tau[\infty] \\ \min\{\Phi[\infty]_{\tau[\infty]}, \Phi[\tau[\infty]]_{\tau[\infty]}\}, & t, T \geq \tau[\infty]. \end{cases}$$

This payment process modifies Φ so that whenever $\tau[\infty]$ is reached without a prior reported state switch, the terminal payment $\Delta\Phi_{\tau[\infty]}$ does not depend on whether the expert reports a state switch at $\tau[\infty]$. The modification of this terminal payoff is chosen so that $\widetilde{\Phi}[\Lambda']_{\tau[\Lambda']} \leq$

$\Phi[\Lambda']_{\tau[\Lambda']}$ for any reporting policy $\Lambda' \geq \Lambda$.

I first claim that $\widetilde{\mathcal{C}} = (\widetilde{\Phi}, \tau)$ is an IC-B contract. For any reporting policy Λ' let $U[\Lambda']$ and $\widetilde{U}[\Lambda']$ be the payoffs for the expert under $\mathcal{C}$ and $\widetilde{\mathcal{C}}$, respectively. Suppose that $\widetilde{U}[\Lambda'] > \widetilde{U}[\Lambda]$ for some reporting policy $\Lambda' \geq \Lambda$. On $\{\Lambda \geq \tau[\infty]\}$, $\widetilde{\Phi}[\Lambda'] = \widetilde{\Phi}[\Lambda]$ by construction and $\tau[\Lambda'] = \tau[\infty] = \tau[\Lambda]$, and therefore

$$\int_0^{\tau[\Lambda']} e^{-\rho t}(b\, dt + d\widetilde{\Phi}[\Lambda']_t) = \int_0^{\tau[\Lambda]} e^{-\rho t}(b\, dt + d\widetilde{\Phi}[\Lambda]_t).$$

In other words, the ex post payoff to the expert is the same under Λ and Λ' whenever $\Lambda \geq \tau[\infty]$. Then by hypothesis there must exist a positive-measure set $S \subset \{\Lambda < \tau[\infty]\}$ on which

$$\int_0^{\tau[\Lambda']} e^{-\rho t}(b\, dt + d\widetilde{\Phi}[\Lambda']_t) > \int_0^{\tau[\Lambda]} e^{-\rho t}(b\, dt + d\widetilde{\Phi}[\Lambda]_t).$$

But $\widetilde{\Phi}[\Lambda] = \Phi[\Lambda]$ on $\{\Lambda < \tau[\infty]\}$, and $\widetilde{\Phi}[\Lambda']_t = \Phi[\Lambda']_t$ for $t < \tau[\Lambda']$ while $\Phi[\Lambda']_{\tau[\Lambda']} \geq \widetilde{\Phi}[\Lambda']_{\tau[\Lambda']}$. Therefore

$$\int_0^{\tau[\Lambda']} e^{-\rho t}(b\, dt + d\Phi[\Lambda']_t) > \int_0^{\tau[\Lambda]} e^{-\rho t}(b\, dt + d\Phi[\Lambda]_t)$$

on S. Finally, construct a new reporting policy Λ'' by $\Lambda'' = \Lambda$ on $\Omega \setminus S$ and $\Lambda'' = \Lambda'$ on S. Then $U[\Lambda''] > U[\Lambda]$, contradicting IC-B. So $\widetilde{\mathcal{C}}$ must be IC-B.

I next claim that $\Pi[\widetilde{\mathcal{C}}] \geq \Pi[\mathcal{C}]$. This follows immediately from the fact that $\widetilde{\Phi}[\Lambda]_t = \Phi[\Lambda']_t$ for $t < \tau[\Lambda']$ while $\widetilde{\Phi}[\Lambda]_{\tau[\Lambda]} \leq \Phi[\Lambda']_{\tau[\Lambda]}$, so that total discounted payments to the expert are weakly lower under $\widetilde{\mathcal{C}}$ than $\mathcal{C}$.

Now define a completion bonus process F by

$$F_t = e^{\rho t} \int_0^t e^{-\rho s}\, d\widetilde{\Phi}[t]_s.$$

As $(\widetilde{\Phi}[t]_s)_{s \leq t}$ is $\mathcal{F}_t^Y$-adapted for each t, F is $\mathbb{F}^Y$-adapted. Define a new payment process Φ' by $\Phi'_t = F_\tau \mathbf{1}\{t \geq \tau\}$ and let $\mathcal{C}' = (\Phi', \tau)$. Then the expert's payoff $U'[\Lambda']$ under $\mathcal{C}'$ and reporting policy Λ' is

$$\begin{aligned} U'[\Lambda'] &= \mathbb{E}^{\Lambda'}\left[\int_0^\tau e^{-\rho t} b\, dt + \mathbf{1}\{\tau < \infty\} e^{-\rho\tau} F_\tau\right] \\ &= \mathbb{E}\left[\int_0^{\tau[\Lambda']} e^{-\rho t} b\, dt + \mathbf{1}\{\tau[\Lambda'] < \infty\} \int_0^{\tau[\Lambda']} e^{-\rho t} d\widetilde{\Phi}[\tau[\Lambda']]_t\right], \end{aligned}$$

where in the second line I have made the dependence of τ and Φ on the reporting policy explicit. Now, by construction $\widetilde{\Phi}[\tau[\Lambda']] = \widetilde{\Phi}[\Lambda']$. This is trivially true on $\{\Lambda' \leq \tau[\infty]\}$, and otherwise $\widetilde{\Phi}$ is constructed to be independent of the exact timing of the report; in particular $\widetilde{\Phi}[\tau[\infty]] = \widetilde{\Phi}[\infty] = \widetilde{\Phi}[\Lambda']$ on $\{\Lambda' > \tau[\infty]\}$. Hence

$$\begin{aligned} U'[\Lambda'] &= \mathbb{E}\left[\int_0^{\tau[\Lambda']} e^{-\rho t} b\, dt + \mathbf{1}\{\tau[\Lambda'] < \infty\} \int_0^{\tau[\Lambda']} e^{-\rho t} d\widetilde{\Phi}[\Lambda']_t\right] \\ &\leq \mathbb{E}\left[\int_0^{\tau[\Lambda']} e^{-\rho t} b\, dt + \int_0^{\tau[\Lambda']} e^{-\rho t} d\widetilde{\Phi}[\Lambda']_t\right] = \widetilde{U}[\Lambda'], \end{aligned}$$

with equality whenever $\mathbb{P}\{\tau[\Lambda'] < \infty\} = 1$. In particular, by assumption $\mathbb{P}\{\tau[\Lambda] < \infty\} = 1$ and so $U'[\Lambda] = \widetilde{U}[\Lambda']$. Then the fact that $\widetilde{\mathcal{C}}$ is IC-B implies $\mathcal{C}'$ is as well.

Finally, the firm's expected profits under $\mathcal{C}'$ are

$$\begin{aligned} \Pi[\mathcal{C}'] &= \mathbb{E}\left[\int_0^{\tau[\Lambda]} e^{-\rho t} r_G\, dt - \mathbf{1}\{\tau[\Lambda] < \infty\} e^{-\rho\tau[\Lambda]} F_\tau\right] \\ &= \mathbb{E}\left[\int_0^{\tau[\Lambda]} e^{-\rho t} r_G\, dt - \mathbf{1}\{\tau[\Lambda] < \infty\} e^{-\rho\tau[\Lambda]} \int_0^{\tau[\Lambda']} e^{-\rho t} d\widetilde{\Phi}[\tau[\Lambda]]_t\right] \\ &= \mathbb{E}\left[\int_0^{\tau[\Lambda]} e^{-\rho t} r_G\, dt - \mathbf{1}\{\tau[\Lambda] < \infty\} e^{-\rho\tau[\Lambda]} \int_0^{\tau[\Lambda']} e^{-\rho t} d\widetilde{\Phi}[\Lambda]_t\right] \\ &\geq \mathbb{E}\left[\int_0^{\tau[\Lambda]} e^{-\rho t} r_G\, dt - e^{-\rho\tau[\Lambda]} \int_0^{\tau[\Lambda']} e^{-\rho t} d\widetilde{\Phi}[\Lambda]_t\right] \\ &= \Pi[\widetilde{\mathcal{C}}]. \end{aligned}$$

Then as $\Pi[\widetilde{\mathcal{C}}] \geq \Pi[\mathcal{C}]$, $\Pi[\mathcal{C}'] \geq \Pi[\mathcal{C}]$ as well.

D.4 Proof of Lemma 4

Fix a contract (F, τ^Y). Let $S \subset \Omega$ be the set of states of the world on which

$$\mathbb{E}\left[\int_{\Lambda\wedge\tau^Y}^{\tau^Y} e^{-\rho(s-\Lambda\wedge\tau^Y)} b\, ds + e^{-\rho(\tau^Y-\Lambda\wedge\tau^Y)} F_{\tau^Y} \,\middle|\, \mathcal{F}_{\Lambda\wedge\tau^Y}\right] > F_{\Lambda\wedge\tau^Y}.$$

I first claim that (F, τ^Y) satisfies IC-∞ iff $\mathbb{P}S = 0$. Suppose first that $\mathbb{P}S > 0$, and define Λ' by $\Lambda'(\omega) = \Lambda(\omega)$ on $\omega \in \Omega \setminus S$ and $\Lambda'(\omega) = \infty$ on $\omega \in S$. Then the expert's expected profits under reporting policy Λ' are

$$U(\Lambda') = \mathbb{E}\left[\mathbf{1}_S\left(\int_0^{\tau^Y} e^{-\rho t} b\,dt + e^{-\rho\tau^Y} F_{\tau^Y}\right) + \mathbf{1}_{\Omega\setminus S}\left(\int_0^{\Lambda\wedge\tau^Y} e^{-\rho t} b\,dt + e^{-\rho(\Lambda\wedge\tau^Y)} F_{\Lambda\wedge\tau^Y}\right)\right].$$

Note that $S \in \mathcal{F}_{\Lambda\wedge\tau^Y}$, so by the law of iterated expectations and the assumption that $\mathbb{P}S > 0$,

$$\begin{aligned}
U(\Lambda') &> \mathbb{E}\left[\mathbf{1}_S\left(\int_0^{\Lambda\wedge\tau^Y} e^{-\rho t} b\,dt + e^{-\rho(\Lambda\wedge\tau^Y)} F_{\Lambda\wedge\tau^Y}\right) + \mathbf{1}_{\Omega\setminus S}\left(\int_0^{\tau^Y\wedge\Lambda} e^{-\rho t} b\,dt + e^{-\rho(\tau^Y\wedge\Lambda)} F_{\tau^Y\wedge\Lambda}\right)\right] \\
&= \mathbb{E}\left[\int_0^{\tau^Y\wedge\Lambda} e^{-\rho t} b\,dt + e^{-\rho(\tau^Y\wedge\Lambda)} F_{\tau^Y\wedge\Lambda}\right] \\
&= U(\Lambda).
\end{aligned}$$

Hence such a contract would violate IC-∞.

In the other direction, suppose $\mathbb{P}S = 0$, and consider any reporting policy Λ' such that $\Lambda'(\omega) \in \{\Lambda(\omega), \infty\}$ for each ω. Let $S' = \{\Lambda' > \Lambda\}$. Then the expert's profits under Λ' are

$$\begin{aligned}
U(\lambda') &= \mathbb{E}\left[\mathbf{1}_{S'}\left(\int_0^{\tau^Y} e^{-\rho t} b\,dt + e^{-\rho\tau^Y} F_{\tau^Y}\right) + \mathbf{1}_{\Omega\setminus S'}\left(\int_0^{\Lambda\wedge\tau^Y} e^{-\rho t} b\,dt + e^{-\rho(\Lambda\wedge\tau^Y)} F_{\Lambda\wedge\tau^Y}\right)\right] \\
&= \mathbb{E}\left[\mathbf{1}_{S'}\mathbf{1}_{\Omega\setminus S}\left(\int_0^{\tau^Y} e^{-\rho t} b\,dt + e^{-\rho\tau^Y} F_{\tau^Y}\right) + \mathbf{1}_{\Omega\setminus S'}\left(\int_0^{\Lambda\wedge\tau^Y} e^{-\rho t} b\,dt + e^{-\rho(\Lambda\wedge\tau^Y)} F_{\Lambda\wedge\tau^Y}\right)\right] \\
&\le \mathbb{E}\left[\mathbf{1}_{S'}\mathbf{1}_{\Omega\setminus S}\left(\int_0^{\Lambda\wedge\tau^Y} e^{-\rho t} b\,dt + e^{-\rho(\Lambda\wedge\tau^Y)} F_{\Lambda\wedge\tau^Y}\right) + \mathbf{1}_{\Omega\setminus S'}\left(\int_0^{\tau^Y\wedge\Lambda} e^{-\rho t} b\,dt + e^{-\rho(\tau^Y\wedge\Lambda)} F_{\tau^Y\wedge\Lambda}\right)\right] \\
&= \mathbb{E}\left[\int_0^{\tau^Y\wedge\Lambda} e^{-\rho t} b\,dt + e^{-\rho(\tau^Y\wedge\Lambda)} F_{\tau^Y\wedge\Lambda}\right] \\
&= U(\Lambda).
\end{aligned}$$

So (F, τ^Y) satisfies IC-∞ if $\mathbb{P}S = 0$.

Note that $\mathbb{P}S = 0$ along with $F \ge 0$ implies

$$\begin{aligned}
F_{\Lambda\wedge\tau^Y} &\ge \mathbb{E}\left[\int_{\Lambda\wedge\tau^Y}^{\tau^Y} e^{-\rho(s-\Lambda\wedge\tau^Y)} b\,ds + e^{-\rho(\tau^Y-\Lambda\wedge\tau^Y)} F_{\Lambda\wedge\tau^Y} \,\middle|\, \mathcal{F}_{\Lambda\wedge\tau^Y}\right] \\
&\ge \mathbb{E}\left[\int_{\Lambda\wedge\tau^Y}^{\tau^Y} e^{-\rho(s-\Lambda\wedge\tau^Y)} b\,ds \,\middle|\, \mathcal{F}_{\Lambda\wedge\tau^Y}\right]
\end{aligned}$$

a.s. I next claim that

$$F^*_{\Lambda\wedge\tau^Y} = \mathbb{E}\left[\int_{\Lambda\wedge\tau^Y}^{\tau^Y} e^{-\rho(s-\Lambda\wedge\tau^Y)} b\,ds \,\middle|\, \mathcal{F}_{\Lambda\wedge\tau^Y}\right]$$

a.s. On $\{\tau^Y \leq \Lambda\}$ this is trivially true as both sides are zero, so consider the set of states $\{\Lambda < \tau^Y\}$. In this case the definition of $\mathbb{E}^B_t$ implies that the left- and right-hand sides coincide.

It follows that $F_{\Lambda\wedge\tau^Y} \geq F^*_{\Lambda\wedge\tau^Y}$ a.s. for every IC-∞ contract (F, τ^Y). It remains only to show that (F^*, τ^Y) is itself IC-∞. But $F^*_{\tau^Y} = 0$ by construction, so $\mathbb{P}S = 0$ boils down to

$$F^*_{\Lambda\wedge\tau^Y} \geq \mathbb{E}\left[\int_{\Lambda\wedge\tau^Y}^{\tau^Y} e^{-\rho(s-\Lambda\wedge\tau^Y)} b\,ds \,\middle|\, \mathcal{F}_{\Lambda\wedge\tau^Y}\right]$$

a.s., and I just showed that the rhs is equal to the lhs a.s. So indeed (F^*, τ^Y) is IC-∞.

Finally, to see that F^* is $\mathbb{F}^Y$-adapted simply note that $\int_{t\wedge\tau^Y}^{\tau^Y} e^{-\rho(s-t)} b\,ds$ is $\mathcal{F}^Y_\infty$-measurable given that τ^Y is an $\mathbb{F}^Y$-stopping time, and invoke the remark made following Definition 7.

D.5 Proof of Lemma 5

Fix any reporting policy $\Lambda' \geq \Lambda$. The expected payoff to the expert of Λ' under (F^*, τ^Y) is

$$U[\Lambda'] = \mathbb{E}\left[\int_0^{\tau^Y\wedge\Lambda'} e^{-\rho t} b\,dt + e^{-\rho(\tau^Y\wedge\Lambda')} F^*_{\tau^Y\wedge\Lambda'}\right].$$

I claim that

$$F^*_{\tau^Y\wedge\Lambda'} = \mathbb{E}\left[\int_{\tau^Y\wedge\Lambda'}^{\tau^Y} e^{-\rho(t-\tau^Y\wedge\Lambda')} b\,dt \mid \mathcal{F}_{\tau^Y\wedge\Lambda'}\right].$$

On $\{\Lambda' \leq \tau^Y\}$ this identity follows from the definition of $\mathbb{E}^B_t$ and the fact that $\Lambda' \geq \Lambda$. And on $\{\Lambda' > \tau^Y\}$ the lhs and rhs are both zero, hence the identity holds everywhere. Then substitute this identity into the previous expression for $U[\Lambda']$ and use the law of iterated expectations to obtain

$$U[\Lambda'] = \mathbb{E}\left[\int_0^{\tau^Y} e^{-\rho t} b\,dt\right].$$

In other words, the expert's payoff is independent of his reporting policy, and in particular $U[\Lambda'] = U[\Lambda]$. Thus IC-B is satisfied.

D.6 Proof of Proposition 1

Fix an $\mathbb{F}^Y$-stopping time τ^Y, and let F^* be the bonus process defined in Lemma 4. Then Lemma 5 implies

$$\Pi[\tau^Y] = \mathbb{E}\left[\int_0^{\Lambda\wedge\tau^Y} e^{-\rho t} r_G\, dt - e^{-\rho(\Lambda\wedge\tau^Y)} F^*_{\Lambda\wedge\tau^Y}\right].$$

Recall from the proof of Lemma 5 that

$$F^*_{\tau^Y\wedge\Lambda'} = \mathbb{E}\left[\int_{\tau^Y\wedge\Lambda'}^{\tau^Y} e^{-\rho(t-\tau^Y\wedge\Lambda')} b\, dt \mid \mathcal{F}_{\tau^Y\wedge\Lambda'}\right].$$

Inserting this identity into the previous representation of $\Pi[\tau^Y]$ and applying the law of iterated expectations yields

$$\Pi(\tau^Y) = \mathbb{E}\left[\int_0^{\Lambda\wedge\tau^Y} e^{-\rho t} r_G\, dt - \int_{\Lambda\wedge\tau^Y}^{\tau^Y} e^{-\rho t} b\, dt\right] = \mathbb{E}\left[\int_0^{\Lambda\wedge\tau^Y} e^{-\rho t}(r_G + b)\, dt - \int_0^{\tau^Y} e^{-\rho t} b\, dt\right].$$

Another application of the law of iterated expectations reduces the first term on the rhs to

$$\begin{aligned}
\mathbb{E}\left[\int_0^{\Lambda\wedge\tau^Y} e^{-\rho t}(r_G + b)\, dt\right] &= \mathbb{E}\left[\int_0^{\infty} \mathbf{1}\{t \le \Lambda\}\mathbf{1}\{t \le \tau^Y\} e^{-\rho t}(r_G + b)\, dt\right] \\
&= \mathbb{E}\left[\int_0^{\infty} \mathbb{E}_t^Y[\mathbf{1}\{t \le \Lambda\}]\mathbf{1}\{t \le \tau^Y\} e^{-\rho t}(r_G + b)\, dt\right] \\
&= \mathbb{E}\left[\int_0^{\tau^Y} \pi_t e^{-\rho t}(r_G + b)\, dt\right].
\end{aligned}$$

Thus

$$\Pi(\tau^Y) = \mathbb{E}\left[\int_0^{\tau^Y} e^{-\rho t}(\pi_t r_G - (1-\pi_t) b)\, dt\right].$$

D.7 Proof of Proposition 2

I will assume that $H(t) < 1$ for every $t \in \mathbb{R}_+$. The remaining case may be treated by a slight modification to the proof considering times only up to $\inf H^{-1}(1)$.

I begin by defining a family of optimal stopping problems indexed by the initial belief about θ and the starting time for the state transition distribution. For each $(x, t) \in [0,1]\times\mathbb{R}_+$, define a probability measure $\mathbb{P}^{(x,t)}$ on $(\Omega, \mathcal{F})$ satisfying:

- $\Lambda \sim H^{(x,t)}$, where $H^{(x,t)}(s) = x + \frac{H(s+t)-H(t)}{1-H(t)}(1-x)$,
- Y is identical in law to $Y^G_{s\wedge\Lambda} + (Y^B_s - Y^B_{s\wedge\Lambda})$,
- The public randomization device has the same distribution as under $\mathbb{P}$ and is independent of Y and Λ.

In the optimal stopping problem indexed by (x,t), the initial probability that the state is Bad is x while the conditional state transition rate $\frac{dH^{(x,t)}(s)}{1-H^{(x,t)}(s-)}$ equals $\frac{dH(s+t)}{1-H((s+t)-)}$ for all s. The objective function $\Pi[\cdot]$ corresponds to $(x,t) = (H(0),0)$, and in fact $\mathbb{P}^{(H(0),0)} = \mathbb{P}$. Also by construction, $\mathbb{P}^{(x,t)} = x\mathbb{P}^{(1,t)} + (1-x)\mathbb{P}^{(0,t)}$ for every x,t. I will write $\mathbb{E}^{(x,t)}$ for the expectation wrt $\mathbb{P}^{(x,t)}$. Note that for any $\mathcal{F}^Y_\infty$-measurable random variable X, $\mathbb{E}^{(0,t)}[X] = \mathbb{E}^B[X]$ for any t. Also in the context of this proof, in the problem indexed by (x,t) the filtration $\mathbb{F}^Y$ will be taken to be the $\mathbb{P}^{(x,t)}$-augmentation of the filtration generated by Y and the public randomization device.

Let $\mathcal{T}$ be the set of $\mathbb{F}^Y$-stopping times. Define the value function $v : [0,1]\times\mathbb{R}_+ \to [0, r_G/\rho]$ for the family of stopping problems by

$$v(x,t) = \sup_{\tau^Y\in\mathcal{T}} \mathbb{E}^{(x,t)}\left[\int_0^{\tau^Y} e^{-\rho s}(\pi^{(x,t)}_s r_G - (1-\pi^{(x,t)}_s)b)\,ds\right],$$

where $\pi^{(x,t)}_s \equiv \mathbb{E}^{(x,t)}[\mathbf{1}\{\Lambda > s\} \mid \mathcal{F}^Y_s]$. Using the reasoning in the proof of Proposition 1, the value function may be equivalently written

$$v(x,t) = \sup_{\tau^Y\in\mathcal{T}} \mathbb{E}^{(x,t)}\left[\int_0^{\tau^Y} e^{-\rho s}(\mathbf{1}\{\Lambda > s\}r_G - \mathbf{1}\{\Lambda \le s\}b)\,ds\right].$$

Fix any $\tau^Y \in \mathcal{T}$. Using the fact that $\mathbb{P}^{(x,t)} = x\mathbb{P}^{(1,t)} + (1-x)\mathbb{P}^{(0,t)}$, the payoff $v^{\tau^Y}(x,t)$ of this strategy is

$$\begin{aligned} v^{\tau^Y}(x,t) = {} & x\mathbb{E}^{(1,t)}\left[\int_0^{\tau^Y} e^{-\rho s}(\mathbf{1}\{\Lambda > s\}r_G - \mathbf{1}\{\Lambda \le s\}b)\,ds\right] \\ & + (1-x)\mathbb{E}^{(0,t)}\left[\int_0^{\tau^Y} e^{-\rho s}(\mathbf{1}\{\Lambda > s\}r_G - \mathbf{1}\{\Lambda \le s\}b)\,ds\right], \end{aligned}$$

or equivalently,

$$v^{\tau^Y}(x,t) = x\mathbb{E}^{(1,t)}\left[\int_0^{\tau^Y} e^{-\rho s}(\pi^{(1,t)}_s r_G - (1-\pi^{(1,t)}_s)b)\,ds\right] - (1-x)\mathbb{E}^B\left[\int_0^{\tau^Y} e^{-\rho s}b\,ds\right].$$

Hence the value function may be written

$$v(x,t) = \sup_{\tau^Y \in \mathcal{T}} \left\{ x\mathbb{E}^{(1,t)} \left[\int_0^{\tau^Y} e^{-\rho s}(\pi_s^{(1,t)} r_G - (1-\pi_s^{(1,t)})b)\, ds \right] - (1-x)\mathbb{E}^B \left[\int_0^{\tau^Y} e^{-\rho s} b\, ds \right] \right\}.$$

Now fix t and consider the family of problems ranging over x. Note that this problem is well-defined for any $x \in \mathbb{R}$. As $v(\cdot, t)$ is a supremum over affine functions of x, it is convex in x and continuous on $\mathbb{R}$; in particular, on $[0,1]$. Further, $v(x,t) \geq 0$ for all x and $v(0,t) = 0$, so $v(\cdot, t)$ is an increasing function on $[0,1]$.

Define $\underline{\pi}^*(t) \equiv \sup\{x \in [0,1] \ : \ v(x,t) = 0\}$. By continuity and monotonicity, $v(x,t) = 0$ for $x \leq \underline{\pi}^*(t)$, and $v(x,t) > 0$ for $x > \underline{\pi}^*(t)$. Define $\tau^{(x,t)} \equiv \inf\{s \ : \ \pi_s^{(x,t)} \leq \underline{\pi}^*(t+s)\}$.

I first show that $\tau^{(x,t)} \in \mathcal{T}$. Note that if $\tau^{(x,t)} < s$ for some s, then $\pi_{s'}^{(x,t)} \leq \underline{\pi}^*(t+s')$ for some $s' < s$. Hence $\{\tau^{(x,t)} < s\} \in \mathcal{F}_s^Y$ for every s. But then $\{\tau^{(x,t)} \leq s\} = \bigcap_{s'>s}\{\tau^{(x,t)} < s'\} \in \mathcal{F}_{s+}^Y$, so if $\mathbb{F}^Y$ is right-continuous, then $\{\tau^{(x,t)} \leq s\} \in \mathcal{F}_s^Y$ for all s. In other words, $\tau^{(x,t)} \in \mathcal{T}$.

It is not immediately obvious that $\mathbb{F}^Y$ is right-continuous, as Y is not a strong Markov process. However, the 2-dimensional process X defined by $X_s = (\pi_s^{(x,t)}, s)$ *is* a strong Markov process, and so is (Y, X). And as X is $\mathbb{F}^Y$-adapted, the augmentation of the filtration generated by (Y, X) and the public randomization device must be the same as the augmentation of the filtration generated by Y and the public randomization device itself. So $\mathbb{F}^Y$ is right-continuous by Proposition 2.7.7 of Karatzas and Shreve (1991).

The next portion of the proof is dedicated to establishing that $\tau^{(x,t)}$ is the smallest optimal stopping time in the problem indexed by (x,t). This will imply in particular that $\tau^{(H(0),0)} = \inf\{t \ : \ \pi_t \leq \underline{\pi}^*(t)\}$ is the smallest maximizer of $\Pi[\cdot]$.

To establish this result, fix the problem indexed by (x,t). I first prove that any $\tau^Y \in \mathcal{T}$ such that $\mathbb{P}^{(x,t)}\{\tau^Y < \tau^{(x,t)}\} > 0$ can be strictly improved upon by another stopping time bounded below by $\tau^{(x,t)}$. Define $F \equiv \{\tau^Y < \tau^{(x,t)}\}$, and suppose $\mathbb{P}^{(x,t)}F > 0$. Then $v\left(\pi_{\tau^Y}^{(x,t)}, \tau^Y\right) > 0$ on F by definition of $\tau^{(x,t)}$. Define

$$w(s) \equiv \mathbf{1}\{\Lambda > s\} r_G - \mathbf{1}\{\Lambda \leq s\} b.$$

Given the strong Markov structure of the optimal stopping problem, there exists a $\tau' \in \mathcal{T}$ such that $\tau' = \tau^Y$ on $\Omega \setminus F$, while $\tau' > \tau^Y$ and

$$\mathbb{E}^{(x,t)} \left[\int_{\tau^Y}^{\tau'} e^{-\rho(s-\tau^Y)} w(s)\, ds \,\middle|\, \mathcal{F}_{\tau^Y}^Y \right] > \frac{1}{2} v\left(\pi_{\tau^Y}^{(x,t)}, \tau^Y\right) > 0$$

on F. Then the payoff of τ' is

$$\begin{aligned}
v^{\tau'}(x,t) &= \mathbb{E}^{(x,t)}\left[\int_0^{\tau'} e^{-\rho s} w(s)\,ds\right] \\
&= \mathbb{E}^{(x,t)}\left[\mathbf{1}_{\Omega\setminus F}\int_0^{\tau^Y} e^{-\rho s} w(s)\,ds + \mathbf{1}_F\left(\int_0^{\tau^Y} e^{-\rho s} w(s)\,ds + e^{-\rho\tau^Y}\int_{\tau^Y}^{\tau'} e^{-\rho(s-\tau^Y)} w(s)\,ds\right)\right] \\
&> \mathbb{E}^{(x,t)}\left[\mathbf{1}_{\Omega\setminus F}\int_0^{\tau^Y} e^{-\rho s} w(s)\,ds + \mathbf{1}_F\int_0^{\tau^Y} e^{-\rho s} w(s)\,ds\right] = v^{\tau^Y}(x,t).
\end{aligned}$$

So τ' yields a strictly higher payoff than τ^Y.

Next I show that any $\tau^Y \in \mathcal{T}$ can be modified to be bounded above by $\tau^{(x,t)} + \varepsilon$ for any $\varepsilon > 0$ while weakly improving payoffs. Fix $\tau^Y \in \mathcal{T}$. By definition of $\tau^{(x,t)}$, for any $\varepsilon > 0$ there exists a $\widetilde{\tau} \in \mathcal{T}$ such that $\tau^{(x,t)} \le \widetilde{\tau} \le \tau^{(x,t)} + \varepsilon$ and $v\left(\pi_{\widetilde{\tau}}^{(x,t)}, \widetilde{\tau}\right) = 0$ on $\{\widetilde{\tau} < \infty\}$. Let $E \equiv \{\tau^Y > \widetilde{\tau}\}$. Then the payoff of τ^Y is

$$\begin{aligned}
& v^{\tau^Y}(x,t) \\
&= \mathbb{E}^{(x,t)}\left[\int_0^{\tau^Y} e^{-\rho s} w(s)\,ds\right] \\
&= \mathbb{E}^{(x,t)}\left[\mathbf{1}_{\Omega\setminus E}\int_0^{\tau^Y} e^{-\rho s} w(s)\,ds + \mathbf{1}_E\left(\int_0^{\widetilde{\tau}} e^{-\rho s} w(s)\,ds + e^{-\rho\widetilde{\tau}}\int_{\widetilde{\tau}}^{\tau^Y} e^{-\rho(s-\widetilde{\tau})} w(s)\,ds\right)\right].
\end{aligned}$$

Given that the optimal stopping problem is strongly Markovian in (x,t),

$$\mathbb{E}^{(x,t)}\left[\int_{\widetilde{\tau}}^{\tau^Y} e^{-\rho(s-\widetilde{\tau})} w(s)\,ds \;\middle|\; \mathcal{F}_{\widetilde{\tau}}^Y\right] \le v\left(\pi_{\widetilde{\tau}}^{(x,t)}, \widetilde{\tau}\right).$$

Also, by assumption $v\left(\pi_{\widetilde{\tau}}^{(x,t)}, \widetilde{\tau}\right) = 0$ given that $E \subset \{\widetilde{\tau} < \infty\}$. Hence

$$v^{\tau^Y}(x,t) \le \mathbb{E}^{(x,t)}\left[\mathbf{1}_{\Omega\setminus E}\int_0^{\tau^Y} e^{-\rho s} w(s)\,ds + \mathbf{1}_E\int_0^{\widetilde{\tau}} e^{-\rho s} w(s)\,ds\right].$$

So define $\tau' \in \mathcal{T}$ by $\tau' = \widetilde{\tau}$ on E and $\tau' = \tau^Y$ on $\Omega\setminus E$. Then by construction τ' yields weakly higher payoffs than τ^Y, $\tau' \le \tau^Y$, and $\tau' \le \tau^{(x,t)} + \varepsilon$. In particular, note that if $\tau^Y \ge \tau^{(x,t)}$, then also $\tau' \ge \tau^{(x,t)}$.

I'm now ready to show that $\tau^{(x,t)}$ is an optimal stopping time. Choose a sequence $\tau^1, \tau^2, \ldots$ in $\mathcal{T}$ such that $v^{\tau^n}(x,t) > v(x,t) - 1/n$ for each n. Modify this sequence to a new sequence

$\widetilde{\tau}^1, \widetilde{\tau}^2, ..$ in $\mathcal{T}$ such that $\tau^{(x,t)} \leq \widetilde{\tau}^n \leq \tau^{(x,t)} + 1/n$ and $v^{\widetilde{\tau}^n}(x,t) > v(x,t) - 1/n$ and n. The results just proven show that such a modification is possible. By the squeeze theorem $\widetilde{\tau}^n \to \tau^{(x,t)}$ pointwise. Then by the bounded convergence theorem, $v^{\widetilde{\tau}^n}(x,t) \to v^{\tau^{(x,t)}}(x,t)$. Hence $v(x,t) \leq v^{\tau^{(x,t)}}$. But as $v(x,t)$ is the supremum of payoffs of all elements of $\mathcal{T}$, it must be that $v^{\tau^{(x,t)}} = v(x,t)$, so $\tau^{(x,t)}$ is an optimal stopping time. Further, $\tau^{(x,t)}$ is the pointwise essential infimum of all optimal stopping times. For suppose $\tau^* \in \mathcal{T}$ is another optimal stopping time. Then in light of previous results, $\tau^* \geq \tau^{(x,t)}$ almost surely.

Finally, note that $v\left(\pi^{(x,t)}_{\tau^{(x,t)}}, \tau^{(x,t)}\right) = 0$ a.s. For as demonstrated earlier, any stopping time which halts when the continuation value is positive with positive probability can be extended to yield strictly higher profits, which would contradict the fact that $\tau^{(x,t)}$ is an optimal stopping time. So suppose that $\underline{\pi}^*(t) > b/(b + r_G)$ for some t. Then $v(x,t) = 0$ for some $x > b/(b + r_G)$. Define $\widetilde{\tau}^{(x,t)} \equiv \inf\{t \ : \ \pi^{(x,t)} \leq b/(b + r_G)\}$. By right-continuity of $\pi^{(x,t)}$, $\widetilde{\tau}^{(x,t)} > 0$. Hence $v^{\widetilde{\tau}^{(x,t)}}(x,t) > 0$, a contradiction. So $\underline{\pi}^*(t) \leq b/(b + r_G)$ for all time.

D.8 Proof of Lemma 6

We begin by establishing strict monotonicity of ξ. This follows from the fact that for every $x' > x \geq \underline{\pi}^*$,

$$\xi(x') = \mathbb{E}^G[\tau(x' \to x)] + \xi(x),$$

where $\tau(x' \to x)$ is the first hitting time at x for the posterior belief process started at x'. Since the posterior belief process is continuous a.s., $\tau(x' \to x) > 0$ a.s., implying $\xi(x') > \xi(x)$. The remainder of the proof establishes the desired properties of ϕ.

Recall that $\xi : [\underline{\pi}^*, 1] \to \mathbb{R}_+$ is a C^2 function whose derivative ξ' satisfies the value function ODE

$$0 = 1 + \left(-\alpha x + SNR^2 x(1-x)^2\right) f(x) + \frac{1}{2} SNR^2 x^2 (1-x)^2 f'(x)$$

on $[\underline{\pi}^*, 1]$, where $SNR = (r_G - r_B)/\sigma$. Note in particular that $\xi''(1)$ is bounded, hence $\xi'(1) = 1/\alpha$ and $\phi(1) = 0$. Now, differentiating ϕ yields

$$\phi'(x) = SNR \left(x(1-x)\xi''(x) + (1-2x)\xi'(x)\right).$$

We will establish that $\phi'(x) < 0$ for every $x \in [\underline{\pi}^*, 1)$. Eliminating ξ'' using the differential equation satisfied by ξ' yields

$$\phi'(x) = SNR\frac{\left(\alpha x - \frac{1}{2}SNR^2 x(1-x)\right)\xi'(x) - 1}{\frac{1}{2}SNR^2 x(1-x)}.$$

Recall that $\xi'(x) \leq 0$ for all x by strict monotonicity. Thus whenever $\alpha \leq \frac{1}{2}SNR^2(1-x)$, $\phi'(x) < 0$. In particular, this is true whenever $x \leq 1 - \frac{\alpha}{\frac{1}{2}SNR^2}$. So the non-trivial case is $x > 1 - \frac{\alpha}{\frac{1}{2}SNR^2}$, in which case $\alpha x - \frac{1}{2}SNR^2 x(1-x) > 0$. We're done if we can prove that

$$\Delta'(x) < \zeta(x) \equiv \frac{1}{\alpha x - \frac{1}{2}SNR^2 x(1-x)}$$

for $x \in \left(1 - \frac{\alpha}{\frac{1}{2}SNR^2}, 1\right)$.

To establish this inequality, we will show that ζ constitutes a subsolution to the value function ODE satisfied by Δ' on $\left(1 - \frac{\alpha}{\frac{1}{2}SNR^2}, 1\right)$. Then as $\Delta'(1) = \zeta(1) = 1/\alpha$, it follows that $\zeta(x) > \Delta'(x)$ for $x \in \left(1 - \frac{\alpha}{\frac{1}{2}SNR^2}, 1\right)$, as desired. A subsolution is a function g satisfying

$$g'(x) < \frac{(\alpha x - SNR^2 x(1-x)^2)\, g(x) - 1}{\frac{1}{2}SNR^2 x^2 (1-x)^2}.$$

Substituting ζ into this equation, the lhs becomes

$$lhs = -\frac{1}{(\alpha x - \frac{1}{2}SNR^2 x(1-x))^2}\left(\alpha - \frac{1}{2}SNR^2(1-2x)\right),$$

while the rhs becomes

$$rhs = \frac{2x-1}{x(1-x)(\alpha x - \frac{1}{2}SNR^2 x(1-x))}.$$

Taking the difference of these two expressions yields

$$rhs - lhs = \frac{\alpha}{x(1-x)(\alpha - \frac{1}{2}SNR^2(1-x))^2} > 0,$$

completing the proof.

D.9 Proof of Lemma 7

By Lemma 1, to prove incentive-compatibility it suffices to establish IC-G. Fix any $\Lambda' \leq \Lambda$. As U is a G-submartingale, so is the stopped process U^Λ. Therefore for each t,

$$U_{\Lambda' \wedge t} \leq \mathbb{E}^G_{\Lambda' \wedge t}\left[U_{\Lambda \wedge t}\right].$$

As $U_{\Lambda \wedge t}$ is $\mathcal{F}_\Lambda$-measurable, and the stopped output process Y^Λ is identical in law to $(Y^G)^\Lambda$, it must be that $\mathbb{E}^G_{\Lambda' \wedge t}\left[U_{\Lambda \wedge t}\right] = \mathbb{E}_{\Lambda' \wedge t}\left[U_{\Lambda \wedge t}\right]$. Then taking unconditional expectations and using the law of iterated expectations,

$$\mathbb{E}\left[U_{\Lambda' \wedge t}\right] \leq \mathbb{E}\left[U_{\Lambda \wedge t}\right].$$

Now note that U is a bounded process taking values in $[0, b/\rho]$. So take $t \to \infty$ and use the bounded convergence theorem to exchange limits and expectations, yielding

$$\mathbb{E}\left[U_{\Lambda'}\right] \leq \mathbb{E}\left[U_\Lambda\right],$$

where $U_\infty = b/\rho$ in case $\Lambda' = \infty$ or $\Lambda = \infty$. But the lhs is the expert's ex ante payoff under Λ', while the rhs is his ex ante payoff under Λ. Hence the contract satisfies IC-G.

D.10 Proof of Proposition 3

I prove the proposition under the assumption that $\sigma > 0$, with the $\sigma = 0$ case being an easy modification of what follows.

Let $\Delta Y_t \equiv Y_t - Y_{t-}$ be the jump part of Y and $Y^c_t \equiv Y_t - \sum_{s \leq t} \Delta Y_s$ be its continuous part, and define $\widetilde{r}_\theta \equiv r_\theta - \sum_{i=1}^n d_i \lambda_i^\theta$. Let

$$\overline{Z}_t \equiv \sigma^{-1}\left(Y^c_t - \int_0^t (\pi_{s-}\widetilde{r}_G + (1 - \pi_{s-})\widetilde{r}_B)\, ds\right)$$

and

$$\overline{N}_i(t) \equiv \sum_{s \leq t} \mathbf{1}\{\Delta Y_s = d_i\} - \int_0^t (\pi_{s-}\lambda_i^G + (1 - \pi_{s-})\lambda_i^B)\, ds$$

be the usual innovation processes. By a standard calculation, π evolves according to the SDE

$$d\pi_t = -\frac{\pi_{t-}}{1 - H(t-)} dH(t) + \frac{\widetilde{r}_G - \widetilde{r}_B}{\sigma}\pi_{t-}(1 - \pi_{t-})\, d\overline{Z}_t + \sum_{i=1}^n \frac{(\lambda_i^G - \lambda_i^B)\pi_{t-}(1 - \pi_{t-})}{\pi_{t-}\lambda_i^G + (1 - \pi_{t-})\lambda_i^B}\, d\overline{N}_i(t).$$

Also, for each $\theta \in \{G, B\}$ let $\widetilde{Z}_t^\theta \equiv \sigma^{-1}(Y_t^c - \widetilde{r}_\theta t)$ and $\widetilde{N}_t^\theta \equiv \sum_{s \le t} \mathbf{1}\{\Delta Y_s = d_i\} - \lambda_i^\theta t$. Under $\mathbb{P}^\theta$, $\widetilde{Z}^\theta$ is a standard Brownian motion, $\widetilde{N}^\theta$ is a compensated Poisson process with rate parameter λ_i, and $\widetilde{Z}^\theta, \widetilde{N}_1^\theta, ..., \widetilde{N}_n^\theta$ are mutually independent.

Note that the updating rule for π may be rewritten

$$\begin{aligned} d\pi_t = &-\frac{\pi_{t-}}{1 - H(t-)} dH(t) + \frac{\widetilde{r}_G - \widetilde{r}_B}{\sigma} \pi_{t-}(1 - \pi_{t-})\, d\widetilde{Z}_t^B + \sum_{i=1}^n \frac{(\lambda_i^G - \lambda_i^B)\pi_{t-}(1 - \pi_{t-})}{\pi_{t-}\lambda_i^G + (1 - \pi_{t-})\lambda_i^B}\, d\widetilde{N}_i^B(t) \\ &- \left(\left(\frac{\widetilde{r}_G - \widetilde{r}_B}{\sigma} \right)^2 \pi_{t-}(1 - \pi_{t-}) + \sum_{i=1}^n \frac{(\lambda_i^G - \lambda_i^B)^2 \pi_{t-}(1 - \pi_{t-})}{\pi_{t-}\lambda_i^G + (1 - \pi_{t-})\lambda_i^B} \right) \pi_{t-}\, dt, \end{aligned}$$

so the process X defined by $X_t \equiv (\pi_t, t)$ is a strong Markov process under $\mathbb{P}^B$. Then as τ^Y is a function only of the path of X and X_t is $\mathcal{F}_t^Y$-measurable for all t, there must exist a function $u : [0, 1] \times \mathbb{R}_+ \to [0, b/\rho]$ such that for each t, $U_t = u(X_t)$ $\mathbb{P}^B$-a.s. on $\{t \le \tau^Y\}$.

Further, $u(p, t)$ must be increasing in p for fixed t. This is because by Bayes' rule, the posterior belief $\pi_t = \mathbb{E}[\mathbf{1}\{\Lambda > s\} \mid \mathbb{F}_s^Y]$ must be monotone increasing in the prior $\mathbb{E}[\mathbf{1}\{\Lambda > s\} \mid \mathbb{F}_t^Y] = \frac{1-H(s)}{1-H(t)}\pi_t$ for every $s > t$, conditional on $(Y_s)_{s \ge t}$. Thus τ^Y is increasing in π_t conditional on $(Y_s)_{s \ge t}$. And under $\mathbb{P}^B$, $(Y_s)_{s \ge t}$ is independent of π_t, since $(Y_s)_{s \ge t}$ is identical in law to $(Y_s^B)_{s \ge t}$ no matter the value of Λ under $\mathbb{P}^B$. Hence the distribution of τ^Y under $\mathbb{P}^B$ is increasing in the FOSD sense as π_t increases. This implies that $u(p, t)$ is increasing in p.

Now I invoke the generalized martingale representation theorem for Lévy processes developed in Nualart and Schoutens (2000), as specialized to the finite jump support case by Davis (2005) (see pg. 66). As U is a B-martingale and is adapted to the filtration generated by Y given the lack of randomization in τ^Y, there exist $\mathbb{F}^Y$-predictable processes $\phi_0, \phi_1, ..., \phi_n$, satisfying $\mathbb{E}^B\left[\int_0^\infty \phi_i^2(t)\, dt\right] < \infty$ for every i, such that

$$U_t = U_0 + \int_0^t \phi_0(s)\, d\widetilde{Z}_s^B + \sum_{i=1}^n \int_0^t \phi_i(s)\, d\widetilde{N}_i^B(s).$$

for all t.

It must be that $(\widetilde{r}_G - \widetilde{r}_B)\phi_0 \ge 0$ and $(\lambda_i^G - \lambda_i^B)\phi_i \ge 0$ for all $i \ge 1$ $\mathbb{P}^B$-a.e. On $\{(t, \omega) : t > \tau^Y(\omega)\}$ this is trivial, as U is a constant process for $t \ge \tau^Y$ and so every ϕ_i must be zero a.e. on this set. So consider the claim on $\{(t, \omega) : t \le \tau^Y(\omega)\}$. Recall that in the updating rule for π_t stated earlier, the loadings on the $(\widetilde{r}_G - \widetilde{r}_B)d\widetilde{Z}^B$ and $(\lambda_i^G - \lambda_i^B)\, d\widetilde{N}_i^B(t)$ terms are all positive. If any of these loadings had the opposite sign in the martingale expansion for dU_t on a positive measure of times and states, then there would exist a time t and positive-measure subsets $A, B \subset \{t \leq \tau^Y\} \subset \Omega$ such that $\pi_t(\omega) \geq \pi_t(\omega')$ and $U_t(\omega) < U_t(\omega')$ for every $\omega \in A, \omega' \in B$. Informally, states in A (respectively, B) correspond to realizations of output with feature at least one of the following:

1. High (low) continuous output runs whenever $\phi_0(s) < 0$ and $\widetilde{r}_G - \widetilde{r}_B > 0$,
2. Low (high) continuous output runs whenever $\phi_0(s) > 0$ and $\widetilde{r}_G - \widetilde{r}_B < 0$,
3. Many (few) jumps of size d_i whenever $\phi_i(s) < 0$ and $\lambda^G > \lambda^B$,
4. Few (many) jumps of size d_i whenever $\phi_i(s) > 0$ and $\lambda^B > \lambda^G$.

But this contradicts the fact that $U_t = u(\pi_t, t)$ $\mathbb{P}^B$-a.s. on $\{t \leq \tau^Y\}$, for a function u which is increasing in its first argument. So $(\widetilde{r}_G - \widetilde{r}_B)\phi_0 \geq 0$ and $(\lambda_i^G - \lambda_i^B)\phi_i \geq 0$ for all $i \geq 1$ $\mathbb{P}^B$-a.e. Without loss, modify the ϕ_i if necessary so these inequalities hold everywhere.

Now rewrite the martingale representation of U as

$$U_t = U_0 + \int_0^t \phi_0(s)\, d\widetilde{Z}_s^G + \sum_{i=1}^n \int_0^t \phi_i(s)\, d\widetilde{N}_i^G(s) + \int_0^t \frac{\widetilde{r}_G - \widetilde{r}_B}{\sigma}\phi_0(s)\, ds + \sum_{i=1}^n \int_0^t (\lambda_i^G - \lambda_i^B)\phi_i(s)\, ds.$$

As the jumps of Y^G and Y^B have bounded support, the Radon-Nikodym derivative for the change of measure over output paths from $\mathbb{P}^G$ to $\mathbb{P}^B$ is $\mathbb{P}^B$-square-integrable. The Cauchy-Schwartz inequality then guarantees that $\mathbb{E}^G\left[\int_0^t \phi_i^2(s)\, ds\right] < \infty$ for every i and t. Hence

$$\mathbb{E}_{t'}^G[U_t] = U_{t'} + \mathbb{E}_{t'}^G\left[\int_{t'}^t \frac{\widetilde{r}_G - \widetilde{r}_B}{\sigma}\phi_0(s)\, ds + \sum_{i=1}^n \int_{t'}^t (\lambda_i^G - \lambda_i^B)\phi_i(s)\, ds\right]$$

for every t and $t' < t$. Since the interior of the expectation on the rhs is non-negative, U is a G-submartingale.

D.11 Proof of Lemma 8

Fix a contract (F, τ^Y). This contract satisfies IC-0 iff

$$F_t \geq \mathbb{E}\left[\int_t^{\Lambda^{(-1)} \wedge \tau^Y} e^{-\rho(s-t)} b\, ds + e^{-\rho(\Lambda^{(-1)} \wedge \tau^Y - t)} F_{\Lambda^{(-1)} \wedge \tau^Y} \,\middle|\, (Y_s)_{s \leq t}, \theta_t = 0\right],$$

when $t \leq \Lambda^{(0)} \wedge \tau^Y$ a.s. This constraint may be equivalently written

$$F_t \geq \mathbb{E}_t^{(0)} \left[\int_t^{\sigma^{(1)} \wedge \tau^Y} e^{-\rho(s-t)} b\, ds + e^{-\rho(\sigma^{(1)} \wedge \tau^Y - t)} F_{\sigma^{(1)} \wedge \tau^Y} \right]$$

for all $t \leq \tau^Y$, where the state space has been augmented to support a family of random variables $\sigma^{(n)}$ for $n = 1, 2, ...$, and $\mathbb{P}^{(0)}$ has been extended so that each $\sigma^{(n)}$ is distributed as the nth tick of a Poisson clock with rate λ_0, and the Poisson clock and Y are independently distributed. The expectation $\mathbb{E}_t^{(0)}$ now also conditions on the event that $\sigma^{(1)} > t$. (This conditioning does not impact the value of the expectation whenever the interior of the expectation does not involve any $\sigma^{(n)}$.)

I first prove that $F^{(0)}$ satisfies these constraints. This may be checked by substituting in the definition of $F^{(0)}$ into the rhs of the constraint and using the law of iterated expectations to find that the rhs becomes $\mathbb{E}_t^{(0)} \left[\int_t^{\tau^Y} e^{-\rho(s-t)} b\, ds \right]$, which is just F_t^0. So indeed $(F^{(0)}, \tau^Y)$ satisfies IC-0.

I next prove that for any IC-0 bonus scheme F, the formula

$$F_t \geq \mathbb{E}_t^{(0)} \left[\int_t^{\sigma^{(n)} \wedge \tau^Y} e^{-\rho(s-t)} b\, ds + e^{-\rho(\sigma^{(n)} \wedge \tau^Y - t)} F_{\sigma^{(n)} \wedge \tau^Y} \right]$$

holds for each $n = 1, 2, ...$ and every $t \leq \tau^Y$. The proof is by induction. The base case has already been established, so it remains only to perform the inductive step. First note that at any time t, the random variable

$$\xi_t = \mathbb{E}_t^{(0)} \left[\int_t^{\sigma^{(1)} \wedge \tau^Y} e^{-\rho(s-t)} b\, ds + e^{-\rho(\sigma^{(1)} \wedge \tau^Y - t)} F_{\sigma^{(1)} \wedge \tau^Y} \right].$$

is equal to the expected discounted flow of benefits until the first tick of a Poisson clock with rate λ_0 which is independent of Y, plus the discounted payment F at the time of the clock tick. Conditional on the information at time $t = \sigma^{(n)}$, the time $\sigma^{(n+1)}$ has the correct distribution, implying that

$$\xi_{\sigma^{(n)} \wedge \tau^Y} = \mathbb{E}_{\sigma^{(n)} \wedge \tau^Y}^{(0)} \left[\int_{\sigma^{(n)} \wedge \tau^Y}^{\sigma^{(n+1)} \wedge \tau^Y} e^{-\rho(s - \sigma^{(n)} \wedge \tau^Y)} b\, ds + e^{-\rho(\sigma^{(n+1)} \wedge \tau^Y - \sigma^{(n)} \wedge \tau^Y)} F_{\sigma^{(n+1)} \wedge \tau^Y} \right]$$

for each n. Hence the base case implies that

$$F_{\sigma^{(n)}\wedge\tau^Y} \geq \mathbb{E}^{(0)}_{\sigma^{(n)}\wedge\tau^Y}\left[\int_{\sigma^{(n)}\wedge\tau^Y}^{\sigma^{(n+1)}\wedge\tau^Y} e^{-\rho(s-\sigma^{(n)}\wedge\tau^Y)} b\, ds + e^{-\rho(\sigma^{(n+1)}\wedge\tau^Y-\sigma^{(n)}\wedge\tau^Y)} F_{\sigma^{(n+1)}\wedge\tau^Y}\right]$$

for all n. So assume the inductive formula has been proven up to case n. Then the inequality just derived combined with the law of iterated expectations implies

$$F_t \geq \mathbb{E}^{(0)}_t\left[\int_t^{\sigma^{(n+1)}\wedge\tau^Y} e^{-\rho(s-t)} b\, ds + e^{-\rho(\sigma^{(n+1)}\wedge\tau^Y-t)} F_{\sigma^{(n+1)}\wedge\tau^Y}\right],$$

proving case $n+1$. So the formula holds for all n.

Finally, take $n\to\infty$ and use the fact that $\sigma^{(n)}\to\infty$ a.s. Fatou's lemma implies

$$F_t \geq \mathbb{E}^{(0)}_t\left[\int_t^{\tau^Y} e^{-\rho(s-t)} b\, ds + e^{-\rho(\tau^Y-t)} F_{\tau^Y}\right] = F^{(0)}_t + \mathbb{E}^{(0)}_t\left[e^{-\rho(\tau^Y-t)} F_{\tau^Y}\right].$$

Hence $F \geq F^{(0)}$ for any completion bonus scheme satisfying IC-0, establishing the optimality of $F^{(0)}$ among IC-0 contracts.

D.12 Proof of Lemma 9

For $x \leq K$, $f(x) = \underline{f}$ for some constant $\underline{f} > 0$. Meanwhile for $x \geq K$, the arguments in the proof of Theorem 2 show that there exists an optimal stopping rule $\tau^*(x)$. Then the conditions of the generalized envelope theorem stated in Theorem 3 of Milgrom and Segal (2002) hold, and f is differentiable for every $x > K$ with $f'(x) = \mathbb{E}[e^{-\rho\tau^*(x)}\pi_{\tau^*(x)}] \leq 1$, with the same result holding as a right-derivative at $x = K$. As the left-derivative of f also exists at $x = K$, f is therefore continuous everywhere. Note further that if $f(x) > x - K$ then necessarily $\tau^*(x) > 0$ with positive probability, hence $f'(x) < 1$. Also, whenever $x \geq K + r_G/\rho$, an optimal stopping rule is trivially $\tau^*(x) = 0$ and so $f(x) = x - K$.

Suppose first that $\underline{f} \geq K$. Recall that $f'(x) \leq 1$ for all $x \geq K$, with the equality strict whenever $f(x) > x - K$. Thus if $f(x) = x$ for some $x \geq K$, then by the fundamental theorem of calculus $f(x') < x'$ for all $x' > x$. So at most one fixed point can exist for $x \geq K$. Further, $f(K) = \underline{f} > K$ by assumption while $f(x) = x - K < x$ for $x \geq K + r/\rho$. So by the intermediate value theorem must exist a unique fixed point $f(x) = x$ on $[K, K + r_G/\rho]$, and since $f(x) = \underline{f} > K$ for $x < K$, this must be the unique fixed point for all x.

Suppose instead that $\underline{f} < K$. Then automatically $f(x) < K$ for all $x \geq K$ by the fundamental theorem of calculus, so trivially f has the single fixed point $\underline{f}$.