Expert advice and optimal project termination

Erik Madsen*

April 3, 2018

Abstract

I analyze how a firm should elicit advice from an expert on when to terminate a project with a stochastic lifespan. The firm cannot directly observe the project's lifespan, but imperfectly monitors its current state by observing incremental output. The expert directly observes the state of the project, but prefers to delay termination as much as possible. He possesses no capital and enjoys limited liability, so cannot be sold the project. The optimal long-term contract involves termination payments to the expert and a stochastic project deadline, which is calculated based on the counterfactual beliefs the firm would have held about the state of the project had it received no expert advice. I sharply characterize the efficiency impact of hiring an expert, and show how the optimal contract changes when the expert has initial capital, can be replaced, or can be assigned busywork on the job.

JEL Classification: D82, D86, J33, J41, L14, M55

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

Economic decision-makers commonly face the problem of deciding when to wind down temporary projects on the basis of imperfectly informative signals about the project's status. A classic example is a firm operating a depreciating plant which must eventually be retired or overhauled. The depreciation can manifest in various ways, such as a decline in average output quality or an increase in downtime or frequency of equipment failure. A key feature of this environment is that the firm cannot directly observe whether the plant has depreciated, but instead sees only a noisy indicator of its status. For instance, it might observe the quality of individual units of output; the market price of its output when demand fluctuates exogenously; or the incidence of individual breakdowns when malfunctions occur at random intervals. Other examples of this class of decision problem include private equity groups deciding when to spin off portfolio companies, firms evaluating when to wrap up consulting engagements, and sports teams determining when to cut aging or injured players.[1]

In these settings the decision-maker faces an optimal stopping problem of when to terminate the project based on the history of public signals. In practice, many such decisions are made with the assistance of experts who bring additional expertise to the table. For instance, the plant owner might hire a manager to oversee the plant's operation; by virtue of the manager's experience and proximity to day to day operations, she is likely to maintain superior knowledge of the plant's status. Similar insider information is possessed by the senior executives employed at a PE group's portfolio firms, the consultants engaged on a project, and athletes with intimate knowledge of their health and day to day performance.

The firm would clearly benefit from incorporating the expert's superior information into the project termination decision. However, experts attached to lucrative projects often possess misaligned incentives to prolong their involvement as long as possible, for instance due to empire-building concerns, on-the-job perquisites, or job search frictions following termination. They also typically possess capital constraints that prevent them from being sold the project to align incentives. The presence of these frictions creates a fundamental tradeoff for the firm between operational efficiency and fiscal economy. As the expert must be compensated for revealing news which instigates early project termination, the firm may choose to enforce a project deadline to economize on incentive payments.

In this paper I study how the firm should optimally incorporate the expert's information into its termination decision using long-term contracts with transfers, when both the expert

[1] A variant of this problem also appears in the operations research literature, in the context of detecting when a signal has been received by a radar station. This setting is often referred to as the problem of quickest detection, change-point detection, or the disorder problem. See Shiryaev (2010) for an overview.

and firm are risk-neutral and discount cashflows at the same rate. I allow output to evolve according to an arbitrary process with stationary independent increments in each state. In particular, my model nests settings in which output evolves as a Brownian motion or a Poisson process with upward or downward jumps. The Brownian case corresponds to environments in which profit fluctuates incrementally with each unit of output produced, while Poisson downward and upward jumps reflect environments with periodic "breakdowns" and "breakthroughs", respectively. 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.

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 for any lost benefits up to that deadline 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", which are the beliefs the firm would have held about the state of the project from monitoring output had it been operating the project without expert advice. If the expert advises that the project be terminated prior to this deadline, the firm does so, and compensates the expert with a termination payment large enough to leave him precisely indifferent between recommending termination and staying silent forever. If the project reaches the deadline, no termination payments are made.

I show that the belief threshold for termination is always lower than it would be if the firm were deciding when to terminate the project without an expert. Hence the expert's 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 an expert, it is now operated ex post efficiently. But if the project would have been shuttered too soon without an expert, its operational lifespan is prolonged, but not necessarily to the efficient termination point. In terms of comparative statics, the optimal termination threshold at each point in time is decreasing in the informativeness of project output about the state, increasing in the rate of state transitions, and increasing in the severity of the expert's incentive misalignment.

An important step in the construction of an optimal contract is the calculation of the firm's "virtual profit function", which eliminates incentive payments from the firm's objective and characterizes an optimal deadline as the solution to an optimal stopping problem maximizing virtual profits. This reduction is analogous to the classic Mirrleesian first-order approach of canonical mechanism design; in that setting, the payments required to implement an arbitrary allocation are first calculated via an envelope theorem argument, and the design problem is then reduced to an optimization over a function of the allocation alone. While my setting requires an entirely different technique to calculate optimal incentive-compatible payments, the spirit of my approach is the same. The virtual profit formulation in my problem is very flexible and can be readily extended in various ways. In particular, I show how to adapt my approach to environments 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. In all of these extensions the basic naive belief threshold rule structure of an optimal deadline survives, indicating the robustness of this result to alternative modeling assumptions.

1.1 Related literature

This paper contributes to the literature on dynamic mechanism design by developing techniques for settings with limited liability, private types which are not directly payoff-relevant to the agent, and imperfect public monitoring. These features depart significantly from the assumptions of most existing models. I briefly discuss this literature with an eye toward illustrating standard sets of assumptions and features of the associated optimal contracts.

One branch of the literature assumes fully flexible transfers and agent marginal valuations for allocations which are increasing in type.[2] 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. Baron and Besanko (1984) and Courty and Li (2000) consider two-period problems with a single agent, while Besanko (1985) and Battaglini (2005) study infinite-horizon settings in discrete time with one agent and special type processes. Pavan, Segal, and Toikka (2014) extend these results to an infinite horizon discrete-time setting

[2]This assumption is typically referred to as a "single crossing" or "strictly increasing differences" condition.

with many agents and general type processes. Williams (2011) analyzes an infinite-horizon problem in continuous time with one agent whose type evolves as an Ornstein-Uhlenbeck process, yielding additional tractability compared to the general discrete-time problem. Typical analyses in this setting are characterization of implementable allocations, including the possibility of efficiency, and optimal contracting.

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.[3] 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 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.[4]

Garrett and Pavan (2012) retain fully flexible transfers but replace agent-payoff-relevant types with imperfect public monitoring, bringing their setting closer to this paper. In their model observable output is the sum of type and random noise as 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.[5] Unlike my model, the agent is not protected by limited liability but can extract rents via time-zero private information. This distinction,

[3]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).

[4]My model nests cases in which the project is known to be viable at time zero, and in which it may be nonviable initially with positive probability. In the latter case the expert possesses private information at time zero. However, the basic structure of an optimal contract is the same in both cases, and is not fundamentally driven by the possible presence of initial private information.

[5]Interestingly, in the presence of career concerns the linkage of output to payments leads disutility of effort to play the role typically served by consumption utility. See in particular their Proposition 4, which casts joint implementability of effort and allocation in terms of a single crossing condition in disutility of effort.

along with the presence of career concerns, 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 bad news from the agent.

A recent set of papers study dynamic delegation examine settings in which transfers are completely absent. These papers typically impose assumptions on preferences analogous to a single-crossing condition, with the agent's preferred allocation sensitive to their private type. Thus types can still be separated by conditioning allocations on reports, though the set of implementable allocations is restricted relative to the case with transfers and becomes more difficult to characterize. In particular, in the absence of transfers the possibility of dynamic information rents reappears when types change over time; papers in this literature tend to restrict attention to static types to retain tractability.

I highlight two recent papers in the dynamic delegation literature which complement my findings on design of optimal termination policies. Guo (2016) considers an environment with public experimentation by an agent who has private information about the true state and a bias toward experimentation. She derives an optimal policy very similar to mine: the principal delegates experimentation to the agent until a virtual belief about the state, which is updated as if the agent's actions were uninformative, reaches a threshold, after which experimentation is cut off forever. Grenadier, A. Malenko, and N. Malenko (2016) study elicitation of an agent's information about the optimal exercise time of a real option, when the agent has a bias for late exercise exactly equivalent to the flow of benefits specification of my model. They also predict an outcome with distinct similarities to the optimal policy in my model: the principal follows the agent's exercise recommendation until a threshold in the value of the underlying asset is reached, at which point the option is exercised immediately.

Finally, Varas (2017) and Green and Taylor (2016) are recent examples of work studying long-term contracting problems with transfers, dynamic arrival of private information, and limited liability. One key difference versus my setting is that in both papers the unobserved state is endogeneous and evolves depending on the level of costly effort privately exerted by the agent. 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 agent does not benefit from longer project operation, and the length of the verification period prior to making incentive payments is instead limited by an agent with 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 the 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 càdlàg stochastic processes satisfying three properties:

1. Y^G, Y^B, and Λ are mutually independent;
2. Each Y^θ has stationary, independent increments;
3. Each Y^θ satisfies $\mathbb{E}[Y^\theta_t] = r_\theta t$ for all t.

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.

This flexible specification nests the typical noise processes used in economic problems with dynamic learning, in particular the Brownian and Poisson good and bad news settings. The Brownian case corresponds to $Y^\theta_t = r_\theta t + \sigma Z^\theta_t$ with Z^θ a standard Brownian motion. And Poisson news is generated by $Y^\theta_t = rt - DN^\theta_t$, where N^θ is a Poisson counting process with rate λ^θ and r and D are chosen so that $r - D\lambda^\theta = r_\theta$. When $\lambda^G < \lambda^B$ this construction yields the bad news model, in which case $r, D > 0$, while $\lambda^G > \lambda^B$ is the good news model, in which case $r, D < 0$. More general signal structures are also permissible, such as compound

Poisson processes and mixtures of Brownian and Poisson signals.[6]

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. However, all results of this paper would go through in discrete time, at the cost of significantly more cumbersome notation.

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.[7] 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 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

[6]Formally, I have assumed that each Y^θ is a Lévy process, a class which nests each of these examples as well as even richer jump structures. The Lévy-Itô decomposition shows that any Lévy process can be decomposed into the sum of three independent processes: a Brownian motion with drift, a compound Poisson process with jumps of size greater than 1, and a compensated compound Poisson process with jumps of size less than 1. In the case that the Lévy process has finite mean, as in this paper, this decomposition may be rewritten as the sum of a Browian motion with drift and a compensated compound Poisson process.

[7]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.

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.[8] I assume that $r_B + b < 0$, so that it is jointly unprofitable for the project to be operated in the Bad state.[9] I also assume that 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

[8]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.

[9]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 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.

worth of output, nonpecuniary costs will never be imposed by the firm. In Section 5.6 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 5.4 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'.[10]

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.[11] To simplify formulae, I assume that a revelation contract makes no transfers subsequent to

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

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

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.[12]

3 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.

[12] In Section 5.5 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.

3.1 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).

3.2 Optimizing project sensitivity

Section 3.1 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 sochasticity 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, there's a catch - 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. In the next section I show how to untangle these interactions and construct an optimal contract.

4 Deriving an optimal contract

4.1 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*

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

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.

4.2 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 3.1 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 *termination fee*, 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 fee 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 termination fees 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.[13] Such contracts are summarizable by the pair (F, τ^Y). Crucially, both 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.

4.3 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 termination fee 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 fee 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

[13] 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.

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 $\mathbb{E}_t^B[X]$ be the expectation of X under $\mathbb{P}^B$ conditional on $\mathcal{F}_t$.*

Informally, one can think of $\mathbb{E}_t^B$ as satisfying $\mathbb{E}_t^B[X] = \mathbb{E}[X \mid (Y_s)_{s \leq t}, \Lambda \leq t]$. The latter expression, however, is not a well-defined random variable, and isn't meaningful if $\mathbb{P}\{\Lambda \leq t\} = 0$. 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 $\mathcal{F}_\infty^Y$-measurable, then $\mathbb{E}_t^B[X]$ is $\mathcal{F}_t^Y$-measurable for each t.*

In general, because $\mathbb{E}_t^B[X]$ may condition on the history of the indicator variable $\mathbf{1}\{\Lambda \leq t\}$, 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 $\mathcal{F}_t^Y$.

Remark. *Suppose X is a stochastic process and Λ' is an $\mathbb{F}$-stopping time satisfying $\Lambda' \geq \Lambda$. Then $\mathbb{E}_{\Lambda'}^B[X_{\Lambda'}] = \mathbb{E}_{\Lambda'}[X_{\Lambda'}]$ 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 $\mathbb{E}_t^B$ 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 termination fee process F^* via*

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

Then F^* is an $\mathbb{F}^Y$-adapted process, (F^*, τ^Y) satisfies IC-∞, and $F^*_{\Lambda \wedge \tau^Y} \leq F_{\Lambda \wedge \tau^Y}$ 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-∞ fee 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 $\mathbb{E}_t^B \left[\int_t^{\tau^Y} e^{-\rho(s-t)} b \, ds + e^{-\rho(\tau^Y - t)} F_{\tau^Y} \right]$ from withholding his report forever. IC-∞ therefore amounts to the requirement that

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

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-∞ fee process implementing τ^Y. Hence if F^* itself satisfies IC-∞, it must be profit-maximizing among all fee 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 fee process would not be guaranteed to satisfy IC-B. This would be true in particular if the state could switch back from Bad to Good. For in this case the expert's payoff from delaying a report includes the option value of waiting to see whether the state switches back before reporting. Since the optimal IC-∞ contract does not capture this option value, IC-B would be violated.

Similar complexities would arise if the expert received only an imperfect signal that the state has switched. For in this case by delaying the expert would retain an option to observe output and refine his belief about the true state before reporting. This option value is similarly omitted from the construction of F^*, and IC-B would fail to hold.

4.4 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 fee 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 fee 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.

4.5 Optimizing virtual profits

Proposition 1 reduces the firm's optimal contracting problem to solving a single-person optimal stopping problem. 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. Cast in the language of economics, the optimization of $\Pi[\tau^Y]$ 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 3.1.) 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 state transition rate $\alpha \geq 0$. In that case current beliefs are a sufficient statistic for the state of the system, and the optimal threshold is time-invariant.

For general output processes, which may exhibit rich jump distributions, the calculation of $\underline{\pi}(t)$ is complex.[14] I will forgo a general formulation and instead describe the solution to two important special cases which illustrate the forces shaping the optimal deadline.

First suppose that output follows a Poisson bad news process in each state: that is, $Y_t^\theta = rt - DN_t^\theta$ for Poisson counting processes N^θ with rates λ^θ satisfying $\lambda^B > \lambda^G \geq 0$, where $r, D > 0$ are characterized by the relations $r - D\lambda^\theta = r_\theta$. Standard arguments[15] show that posterior beliefs evolve according to the stochastic differential equation

$$d\pi_t = -\frac{\pi_{t-}\, dH(t)}{1 - H(t-)} + \frac{(\lambda^B - \lambda^G)\pi_{t-}(1-\pi_{t-})}{\lambda^G \pi_{t-} + \lambda^B(1-\pi_{t-})}\, d\overline{Z}_t,$$

where $\overline{Z}$ is an $\mathbb{F}^Y$-martingale known as the *innovation process* which satisfies

$$d\overline{Z}_t = D^{-1}\left(dY_t - (r_G\pi_{t-} + r_B(1-\pi_{t-}))\, dt\right).$$

The first term in the evolution of π_t captures the deterioration of beliefs due to the expected arrival of a state switch. Meanwhile the second term is the adjustment to beliefs from observation of the output flow. Because $\overline{Z}$ is a martingale, the average change in beliefs over time would be zero if not for the arrival of the state switch. The arrival of bad news always leads to a discrete downward adjustment of beliefs. On the other hand absence of bad news leads to an upward revision from the learning term, but a downward revision from the state switching term. Depending on the relative speeds of learning and state switching, as well as the current level of beliefs, this drift can take either sign.

The following lemma characterizes an optimal stopping rule when the state transition rate is sufficiently high. To state the result, I make use of the fact that H is a monotone function and is therefore differentiable a.e., with derivative h.

Lemma 6. *Suppose Y evolves as a Poisson bad news process in each state. Define $\alpha \equiv$ ess $\inf_{t\geq 0} \frac{h(t)}{1-H(t)}$. Then if $\alpha \geq \frac{\lambda^B - \lambda^G}{1 + b/r_G}$, $\tau^\dagger$ is an optimal public deadline.*

Under the lower bound on the hazard rate of state switching in the lemma statement, the state switches quickly enough that even absent bad news π drifts downward whenever $\pi_t/(1-\pi_t) > b/r_G$. Hence once the posterior likelihood drops below b/r_G, it can never again rise above this level. So there is no option value in waiting for news, meaning $\tau^\dagger$ is an optimal stopping policy. Note that beliefs may still drift upward absent news when the

[14]See Buonaguidi and Muliere (2016) for an example of solving an optimal stopping problem involving learning from Lévy processes.

[15]See Peskir and Shiryaev (2006), pg. 357, for a derivation under the homogeneous state transition distribution $H(t) = H(0) + (1-H(0))(1-\exp(-\alpha t))$. Their derivation extends to general H straightforwardly.

posterior likelihood is below b/r_G. But this upward drift is decreasing in π_t, and the bound in the lemma statement ensures that the drift becomes nonpositive by the time the likelihood reaches b/r_G.

Now suppose that output is Brownian in each state: that is, $Y_t^\theta = r_\theta\, dt + \sigma\, dZ_t^\theta$, where each Z^θ is a standard Brownian motion. In this setting beliefs evolve as[16]

$$d\pi_t = -\frac{\pi_{t-}\, dH(t)}{1 - H(t-)} + \left(\frac{r_G - r_B}{\sigma}\right) \pi_{t-}(1 - \pi_{t-})\, d\overline{Z}_t,$$

where

$$d\overline{Z}_t = \sigma^{-1} \left(dY_t - (r_G \pi_{t-} + r_B(1 - \pi_{t-}))\, dt\right)$$

is a standard Brownian motion with respect to $\mathbb{F}^Y$. Again there are two contributions to belief updating - a downward revision due to the expected arrival of a state switch, and a mean-zero revision due to learning about the state from output. The speed of learning is modulated by the signal-to-noise ratio $(r_G - r_B)/\sigma$. The larger the difference in drifts, or the smaller the noise in output, the larger are expected revisions to beliefs over time.

This problem is recursive in the state variable $X_t = (\pi_t, t)$, with time included in the state due to the inhomogeneity of the transition process for θ. The value of the real option at time t may therefore be denoted $V(\pi_t, t)$, where V is referred to as the problem's value function. For simplicity, I will assume that H is continuously differentiable with derivative h, and define $\alpha(t) \equiv h(t)/(1 - H(t))$ to be the conditional hazard rate for arrival of the state switch. If α is of bounded variation, then V is sufficiently smooth that it satisfies the HJB equation

$$\rho V(\pi, t) = \pi r_G - (1 - \pi) b - \alpha(t) \pi V_\pi(\pi, t) + \frac{1}{2} \left(\frac{r_G - r_B}{\sigma}\right)^2 \pi^2 (1 - \pi)^2 V_{\pi\pi}(\pi, t) + V_t(\pi, t)$$

for beliefs $\pi > \underline{\pi}(t)$, where $\underline{\pi}(t) \leq b/(b + r_G)$ is the firm's optimal termination threshold. This threshold is calculated by solving the free boundary problem posed by the HJB equation along with the value-matching and smooth-pasting conditions $V(\underline{\pi}(t), t) = 0$ and $V_\pi(\underline{\pi}(t), t) = 0$ at the termination boundary. Transversality conditions at $\pi = 1$ (where the ODE becomes singular) and $t = \infty$ are also needed to eliminate spurious exploding solutions. As the true value function lies in the interval $[0, r_G/\rho]$, it is sufficient to require that V be a bounded function. Appendix B formally establishes that a solution to this problem characterizes the optimal public deadline.

[16] See Peskir and Shiryaev (2006), pg. 309-10, for a derivation when the state transition rate is homogeneous.

The HJB equation captures the costs and benefits to the firm of waiting for news about the project's state. The left-hand side is the total value of waiting, normalized by ρ to be in flow units. On the right-hand side are the flow profits received by the firm, plus the expected revisions to the project's option value due to the passage of time. The second term accounts for the fact that beliefs drift downward on average over time, diminishing the value of waiting. Because flow profits are increasing in beliefs, V is monotone increasing in the state and so this term is always negative. The third term captures the option value of waiting, due to stochasticity in the belief process which might improve future beliefs. The value function can be shown to be convex in beliefs due to this optionality (see the proof of Proposition 2), so this term is always positive, and is increasing in the signal-to-noise ratio $(r_G - r_B)/\sigma$ of the process. The final term captures the inhomogeneity of the state transition process. This term can be either positive or negative, depending on whether state transitions speed up or slow down as time passes. In the case where $\alpha(t) = \alpha$ for all time and state transitions are homogeneous, $V_t(\pi, t) = 0$.

4.6 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 fee-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 $\mathbb{E}^G_t[X]$ be the expectation of X under $\mathbb{P}^G$ conditional on $\mathcal{F}_t$.*

Analogously to $\mathbb{E}^B_t$, 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 $\mathbb{E}^B_t[X_s] = X_t$ for all $s > t$, and is a* G-martingale *if $\mathbb{E}^G_t[X_s] = X_t$ 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 under a slight specialization of the model to output processes whose jumps have sizes drawn from a finite set.[17] For any process X, define $\Delta X_t \equiv X_t - X_{t-}$ to be its jump process, which is nonzero only when X is discontinuous.

[17] I conjecture that the result holds true as well for general output processes, which may follow any Lévy process with general jump size distributions. However, the martingale representation theorem and belief updating rule I use to formally prove the result are difficult to work with in the general case.

Proposition 3. *Suppose there exists a finite set $\mathcal{D} \subset \mathbb{R}\backslash\{0\}$ such that ΔY^θ is $\mathcal{D}\cup\{0\}$-valued for each θ. 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. *If Y satisfies the hypothesis of Proposition 3, then 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.

5 Discussion and extensions

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 characterizes an optimal public deadline as a (generally time-dependent) threshold rule $\underline{\pi}(t)$ in naive beliefs. Even without calculating $\underline{\pi}(t)$ explicitly, its characterization as the solution to an option value problem allows for an easy discussion of comparative statics.

I begin by precisely defining what is meant by several of the comparative statics exercises discussed below. First, whenever each Y^θ has a Brownian component with common[18] volatility $\sigma > 0$, comparative statics on the drift of output r_θ will be interpreted as a change in the drift rate of Y^θ, leaving its volatility and jump distribution constant. Such a shift may be accomplished by decomposing Y^θ as $Y^\theta = \widetilde{r}_\theta t + \sigma Z^\theta + \widetilde{Y}^\theta$, where $\widetilde{Y}^\theta$ is a pure jump process, with $\widetilde{r}_\theta = r_\theta - \mathbb{E}[\widetilde{Y}^\theta_1]$. A change in r_θ will then be formally defined as a change in $\widetilde{r}_\theta$, holding fixed Z^θ and $\widetilde{Y}^\theta$. When performing this comparative static, I will assume that $\widetilde{r}_G > \widetilde{r}_B$, so

[18] If the volatilities were not constant, the processes would be instantaneously distinguishable and the comparative statics exercise would be uninteresting.

that whenever r_G increases, Y^G and Y^B become more distinguishable, while whenever r_B increases, they become less distinguishable. This holds whenever r_G is sufficiently large, r_B is sufficiently small, or when each Y^θ is continuous so that $\mathbb{E}[\widetilde{Y}^\theta] = 0$.

If $\sigma = 0$, then a literal shift of the drift rate in the fashion just outlined would lead to immediately distinguishable processes and be uninteresting. However, when Y^θ has jumps of a finite number of sizes, so that jumps of a given size d have a well-defined arrival intensity $\lambda^\theta(d)$ under Y^θ, then all comparative statics holding for an increase in r_θ continue to hold in the following sense. An increase in r_G can be taken to be an increase in $\lambda^G(d)$ for jumps of some size $d > 0$, or a decrease in $\lambda^G(d)$ for jumps of some size $d < 0$, for any d such that $d(\lambda^G(d) - \lambda^B(d)) > 0$. And similarly, an increase in r_B can be taken to be an increase in $\lambda^B(d)$ for jumps of some size $d > 0$, or a decrease in $\lambda^B(d)$ for jumps of some size $d < 0$, for any d such that $d(\lambda^G(d) - \lambda^B(d)) < 0$. The final condition in each of these definitions ensures that when r_G is increased, Y^G and Y^B become more distinguishable, while whenever r_B is increased, they become less distinguishable. This mimics the requirement that $\widetilde{r}_G > \widetilde{r}_B$ imposed when $\sigma > 0$. In particular, under a Poisson good news process increasing r_θ corresponds to increasing λ^θ, while under a Poisson bad news process increasing r_θ corresponds to decreasing λ^θ.

Next, I will 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.

Finally, I will refer to the signal-to-noise ratio (SNR) of the output process when considering changes in the speed at which the underlying true state is extracted from observation of the signal process. I will define this measure only when each Y^θ has a Brownian component, and will take it to be the inverse of the (common) volatility parameter σ of the Brownian component of output. (If the volatility parameters are not equal in the two states, then the states are immediately distinguishable and all comparative statics are trivial.)

I now consider how a parameter change impacts the level of the threshold. Any parameter changes which raise the break-even virtual profit threshold will mechanically induce a higher optimal termination threshold. Such changes include a decrease in r_G or an increase in

b. Also, any parameter changes which diminish the option value of continuing the project will increase the optimal termination threshold. These changes include an increase in r_B, a decrease r_G, or a decrease in the SNR, all of which decrease fluctuations in beliefs; 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 by assumption in the same direction. Table 1 summarizes these results.

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

Table 1: Comparative statics of the termination threshold

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 decreases profits under the optimal contract straightforwardly due to the resulting diminution of flow profits at each moment in time.[19] Next, an increase in the path of the state switching rate α straightforwardly decreases achievable profits by decreasing the amount of time the project spends in the Good state.[20] Finally, an increase in r_G or the SNR or a decrease in r_B 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.

Unsurprisingly, the set of parameters that increase optimal profits are precisely the ones which decrease the termination threshold. This equivalence reflects the fact that increased optionality both increases optimal profits and decreases the threshold at which the project should be terminated.

[19] This reasoning is valid only when comparing the optimal contract under different discount rates. For a fixed contract, it's possible for an increase in ρ to increase expected profits, if the contract suffers large losses from operating long past a state switch. However, an optimal contract never operates this way and so cannot benefit from an increase in the discount rate.

[20] As for ρ, for a fixed contract an increase in α might actually increase profits. In particular, if a contract is sensitive to output early on but after a trial period operates forever, then faster switching early on can lead the contract to terminate more often instead of operating forever with flow losses much of the time. But an optimal contract will never exhibit such behavior.

Parameter	Sign of ΔProfits/ΔParameter
r_G	+
SNR	+
b	-
r_B	-
ρ	-
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 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.

5.4 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 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.

5.5 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

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

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 $X_s = (\pi_s, s)$. 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 $\pi_s \widetilde{\Pi}(t+s)$ 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.[21] This problem may be written recursively as follows. Let $\Pi^*(t)$ be the profit of an optimal contract, with the

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

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 profit from replacing 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 $\Pi^*(t)$ 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 8. *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.

5.6 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.[22]

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.[23] (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

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

[23] 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.

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.[24] Thus the firm's virtual profit function becomes

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

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.[25] 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

[24] 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.

[25] 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.

are always positive at the optimal busywork level, and otherwise they continue to be negative for sufficiently low beliefs.

6 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 characterized - 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.

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 important unexplored avenue is a richer model of state transitions, for instance more than two states or a project that fluctuates between Good and Bad. Solving such a model would require a much more delicate characterization of incentive-compatibility constraints, as any misreport by the expert at one stage of the project would have complex implications for the optimal timing of reporting in future stages. (By contrast, in the literature on dynamic mechanism design without limited liability the expert can be charged up front for all future information rents, in return for efficient usage of all orthogonalized information received after time zero.)

Another interesting avenue for future work would be coarsening the information structure for the expert. One possibility would be to give the expert only an imperfect signal of the state. This would force the expert to form his own private beliefs about the state, and would embed an additional option value problem into the environment, since the expert would possess optionality to learn more about the state before making a report. The calculation of optimal incentive payments would then become more challenging, and the optimal deadline would be less crisply characterizable. Another interesting information structure would involve an expert who observes a state switch only with a delay. This modification would preserve the structure of optimal incentive payments but lead to an optimal stopping problem for the firm involving a higher-dimensional belief space, as the firm would need to form beliefs not only about whether the state has switched, but also about whether the expert has observed the switch.

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.

Buonaguidi, B. and P. Muliere (2016). "Bayesian sequential testing for Lévy processes with diffusion and jump components". In: *Stochastics* 88.7, pp. 1099–1113.

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.

Jacka, S. D. (1991). “Optimal Stopping and the American Put”. In: *Mathematical Finance* 1.2, pp. 1–14.

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.

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.

Peskir, Goran (2005). “A Change-of-Variable Formula with Local Time on Curves”. In: *Journal of Theoretical Probability* 18.3, pp. 499–535.

Peskir, Goran and Albert N. Shiryaev (2006). *Optimal Stopping and Free-Boundary Problems*. Birkhäuser Verlag.

Shiryaev, Albert N. (2010). “Quickest Detection Problems: Fifty Years Later”. In: *Sequential Analysis* 29.4, pp. 345–385.

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 10. *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* $\mathcal{B}(\mathbb{R}_+ \cup \{\infty\}) \otimes \mathcal{F}^Y_t$*-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 10 ensures that this construction yields measurable, $\mathbb{F}$-adapted processes for any choice of Λ'. The final requirement in Definition 10 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 public deadline under Brownian output

As discussed in Section 4.5, when output is Brownian and the state transition hazard rate α is continuous and BV, the firm's value function solves the free boundary problem

$$\rho V(\pi,t) = \pi r_G - (1-\pi)b - \alpha(t)\pi V_\pi(\pi,t) + \frac{1}{2}\left(\frac{r_G - r_B}{\sigma}\right)^2 \pi^2(1-\pi)^2 V_{\pi\pi}(\pi,t) + V_t(\pi,t)$$

subject to boundedness of V and the boundary conditions $V(\underline{\pi}(t),t) = 0$ and $V_\pi(\underline{\pi}(t),t) = 0$ for some $\underline{\pi}(t) \leq b/(b+r_G)$. The following verification theorem establishes that an appropriately smooth solution to this problem yields an optimal public deadline.

Lemma 9. *Suppose there exists a continuous BV function $\underline{\pi} : \mathbb{R}_+ \to [0, b/(b+r_G)]$ and a bounded, non-negative C^2 function V on $\mathcal{O} = \{(\pi,t) \in [0,1) \times \mathbb{R}_+ \ : \ \pi \geq \underline{\pi}(t)\}$ such that V satisfies the HJB equation on $\mathcal{O}$ and $V(\underline{\pi}(t),t) = 0$ and $V_\pi(\underline{\pi}(t),t) = 0$ for all t. Then $\tau^* = \inf\{t \ : \ \pi_t \leq \underline{\pi}(t)\}$ is an optimal public deadline.*

Proof. Extend V to $[0,1) \times \mathbb{R}_+$ by setting $V(\pi,t) = 0$ for $\pi < \underline{\pi}(t)$. On this extended domain V is C^2 everywhere except on the set $\partial\mathcal{O} = \{(\pi,t) \ : \ \pi = \underline{\pi}(t)\}$, where V_π exists and is continuous by smooth pasting. Notice also that on the extended domain, V satisfies

$$\rho V(\pi,t) \geq \pi r_G - (1-\pi)b - \alpha(t)\pi V_\pi(\pi,t) + \frac{1}{2}\left(\frac{r_G - r_B}{\sigma}\right)^2 \pi^2(1-\pi)^2 V_{\pi\pi}(\pi,t) + V_t(\pi,t)$$

(except on $\partial\mathcal{O}$), with equality on $\mathcal{O}$.

Now fix an arbitrary $\mathbb{F}^Y$-stopping time, and suppose $H(0) > 0$, so that $\pi_0 < 1$. Let $\tau^n = \tau^Y \wedge \inf\{t \ : \ \pi_t \geq \tau^n\}$. Because $\{(\omega,t) \ : \ X_t(\omega) \in \partial\mathcal{O}\}$ has measure zero and V_π is continuous on $\partial\mathcal{O}$, the extension of Ito's lemma given in Theorem 2.1 of Peskir (2005) implies that for each t,

$$\begin{aligned}
&e^{-\rho(t\wedge\tau^n)} V(\pi_{t\wedge\tau^n}, t\wedge\tau^n) \\
&= V(\pi_0, 0) + \int_0^{t\wedge\tau^n} e^{-\rho s}\left(\rho V(\pi_s,s) - \alpha(s)\pi_s V_\pi(\pi_s,s)\right. \\
&\qquad\qquad \left. + \frac{1}{2}\left(\frac{r_G - r_B}{\sigma}\right)^2 \pi_s^2(1-\pi_s)^2 V_{\pi\pi}(\pi_s,s) + V_t(\pi_s,s)\right) ds \\
&\quad + \int_0^{t\wedge\tau^n} e^{-\rho s}\frac{r_G - r_B}{\sigma}\pi_s(1-\pi_s)V_\pi(\pi_s,s)\, d\overline{Z}_s.
\end{aligned}$$

As $V_\pi(\pi_s,s)$ is bounded on $[0,\tau^n]$, the final term is a martingale. Then taking expectations and invoking the HJB equation yields

$$V(\pi_0, 0) \geq \mathbb{E}\left[\int_0^{t\wedge\tau^n} e^{-\rho s}(\pi_s r_G - (1-\pi_s)b)\, ds + e^{-\rho(t\wedge\tau^n)} V(\pi_{t\wedge\tau^n}, t\wedge\tau^n)\right],$$

with equality if $\tau^n \leq \tau^*$. Now take $t, n \to \infty$. As the interior of the expectation on the rhs is uniformly bounded for all t and n, the bounded convergence theorem allows us to swap the expectation and limits. And as $\mathbb{P}^G$ and $\mathbb{P}^B$ induce equivalent measures over output paths, $\pi_t < 1$ for all t almost surely, hence $\tau^n \to \tau^Y$ pointwise a.s. So

$$V(\pi_0, 0) \geq \mathbb{E}\left[\int_0^{\tau^Y} e^{-\rho s}(\pi_s r_G - (1-\pi_s)b)\, ds + e^{-\rho\tau^Y} V(\pi_{\tau^Y}, \tau^Y)\right],$$

with equality if $\tau^Y \leq \tau^*$, where the final term is taken to be zero if $\tau^Y = \infty$. In the case that $\tau^Y = \tau^*$, the final term is zero a.s. given $V(\pi, t) = 0$ for $\pi \leq \underline{\pi}(t)$. Thus

$$V(\pi_0, 0) = \mathbb{E}\left[\int_0^{\tau^*} e^{-\rho s}(\pi_s r_G - (1-\pi_s)b)\, ds\right].$$

And as $V \geq 0$ by assumption,

$$V(\pi_0, 0) \geq \mathbb{E}\left[\int_0^{\tau^Y} e^{-\rho s}(\pi_s r_G - (1-\pi_s)b)\, ds\right]$$

for arbitrary τ^Y. Thus $\Pi[\tau^*] \geq \Pi[\tau^Y]$, i.e. τ^* is an optimal stopping rule.

Finally, suppose $H(0) = 0$. Let $\underline{t} = \inf\{t \,:\, H(t) > 0\}$. Then $\pi_t = 1$ for all $t \leq \underline{t}$ while $\pi_t < 1$ for all $t > \underline{t}$ a.s., and the previous argument may be slightly modified by beginning the Ito expansion after $\underline{t}$ to show that

$$\mathbb{E}\left[\int_t^{\tau^*\vee t} e^{-\rho s}(\pi_s r_G - (1-\pi_s)b)\, ds\right] \geq \mathbb{E}\left[\int_t^{\tau^Y\vee t} e^{-\rho s}(\pi_s r_G - (1-\pi_s)b)\, ds\right]$$

for every $t > \underline{t}$ and τ^Y. Now take $t \downarrow \underline{t}$ and exchange limits and expectations by the bounded convergence theorem to obtain

$$\mathbb{E}\left[\int_{\underline{t}}^{\tau^*\vee\underline{t}} e^{-\rho s}(\pi_s r_G - (1-\pi_s)b)\, ds\right] \geq \mathbb{E}\left[\int_{\underline{t}}^{\tau^Y\vee\underline{t}} e^{-\rho s}(\pi_s r_G - (1-\pi_s)b)\, ds\right].$$

Finally, add $\mathbb{E}\left[\int_0^{\underline{t}} e^{-\rho s} r_G \, ds\right]$ to both sides to obtain

$$\mathbb{E}\left[\int_0^{\tau^* \vee \underline{t}} e^{-\rho s}(\pi_s r_G - (1-\pi_s)b) \, ds\right] \geq \mathbb{E}\left[\int_0^{\tau^Y \vee \underline{t}} e^{-\rho s}(\pi_s r_G - (1-\pi_s)b) \, ds\right].$$

Given $\tau^* > \underline{t}$ by construction, the lhs is exactly $\Pi[\tau^*]$. Meanwhile $\int_{\tau^Y}^{\tau^Y \vee \underline{t}} e^{-\rho t} r_G \, dt \geq 0$, so the rhs is an upper bound on $\Pi[\tau^Y]$. So $\Pi[\tau^*] \geq \Pi[\tau^Y]$ for every τ^Y, as desired. □

The remaining step is to establish that this free boundary problem in fact has a solution. This can be accomplished using a variant of the approach used to characterize American put option prices. See, e.g., Jacka (1991), in particular Proposition 2.6, for a detailed treatment.

In the time-homogeneous case where $\alpha(t)$ is constant and V and $\underline{\pi}$ are time-independent, the following lemma provides a constructive existence proof, exhibiting a solution involving Tricomi's confluent hypergeometric function $U(a, b, z)$.

Lemma 10. *Let $k \equiv \frac{r_G - r_B}{\sigma}$ and $\beta \equiv \frac{k^2 + 2\alpha + \sqrt{(k^2+2\alpha)^2 + 8k^2\rho}}{2k^2}$. Then $\beta > 1$ and there exist constants $C > 0$ and $\underline{\pi} \in (0, b/(b + r_G)]$ such that*

$$v(\pi) = \frac{r_G + b}{\rho + \alpha}\pi - \frac{b}{\rho} + C\pi^\beta(1-\pi)^{1-\beta} U\left(\beta - 1, 2\beta - \frac{2\alpha}{k^2}, \frac{2\alpha}{k^2}\frac{\pi}{1-\pi}\right)$$

is a bounded, non-negative C^2 function on $[\underline{\pi}, 1)$ satisfying

$$\rho v(\pi) = \pi r_G - (1-\pi)b - \alpha\pi v'(\pi) + \frac{1}{2}\left(\frac{r_G - r_B}{\sigma}\right)^2 \pi^2(1-\pi)^2 v''(\pi)$$

on $[\underline{\pi}, 1)$ and $v(\underline{\pi}) = v'(\underline{\pi}) = 0$.

Proof. I begin by deriving a general solution to the ODE

$$\rho v(x) = x r_G - (1-x)b - \alpha x v'(x) + \frac{k^2}{2}x^2(1-x)^2 v''(x).$$

This is an inhomogeneous second-order linear ODE, whose solution can be found by conjecturing a particular solution and then solving the associated homogeneous equation. A natural conjecture is linear in x; inserting $v_0(x) = c_1 x + c_0$ and matching coefficients reveals that

$$v_0(x) = \frac{r_G + b}{\rho + \alpha}x - \frac{b}{\rho}$$

is a particular solution to the ODE. The problem of solving the ODE then reduces to solving the associated homogeneous equation

$$\rho v_H(x) = -\alpha x v_H'(x) + \frac{k^2}{2} x^2 (1-x)^2 v_H''(x).$$

Now I make the transformation $z \equiv \frac{x}{1-x}$, obtaining the transformed ODE

$$\rho \hat{v}_H(z) = \left(-\alpha z + k^2 - \alpha - \frac{k^2}{z+1} \right) z \hat{v}_H'(z) + \frac{k^2}{2} z^2 \hat{v}_H''(z),$$

where $\hat{v}_H(z) \equiv v_H(z/(z+1))$. Next I guess that $w_H(z) \equiv \frac{z^\beta}{1+z} \hat{v}_H(z)$ satisfies a simpler ODE than $\hat{v}_H$ itself for some positive power of β. Inserting into the ODE yields

$$\begin{aligned}
&\left(\alpha(\beta-1)z + \rho + \alpha + (\alpha - k^2)(\beta - 1) - \frac{1}{2} k^2 (\beta-1)(\beta-2) \right) w_H(z) \\
&= (k^2\beta - \alpha - \alpha z) z w_H(z) + \frac{k^2}{2} z^2 w_H''(z).
\end{aligned}$$

The right choice of β is therefore a solution to

$$\rho + \alpha + (\alpha - k^2)(\beta - 1) - \frac{1}{2} k^2 (\beta-1)(\beta-2) = 0,$$

which is the quadratic

$$\beta^2 - \left(1 + \frac{2\alpha}{k^2}\right)\beta - 2\frac{\rho}{k^2} = 0.$$

It is straightforward to verify that a unique positive solution to this equation exists. Taking this choice of β, the ODE for w_H reduces to

$$\alpha(\beta-1) w_H(z) = (k^2\beta - \alpha - \alpha z) w_H'(z) + \frac{k^2}{2} z w_H'(z).$$

Finally, make the substitution $t \equiv \frac{2\alpha}{k^2} z$. I arrive at the ODE

$$(\beta - 1)\hat{w}_H(t) = \left(2\beta - \frac{2\alpha}{k^2} - t \right) \hat{w}_H'(t) + \hat{w}_H''(t),$$

where $\hat{w}_H(t) \equiv w_H\left(\frac{k^2}{2\alpha} t\right)$. This is Kummer's differential equation, which has general solution

$$\hat{w}_H(t) = C_1 U(m, n, t) + C_2 M(m, n, t),$$

where U and M are Tricomi's and Kummer's confluent hypergeometric functions and $m \equiv$

$\beta - 1$ and $n \equiv 2\beta - \frac{2\alpha}{k^2}$, with $m > 0$ and $n > m + 2$ given $\beta > 1 + \frac{2\alpha}{k^2}$. Transforming back to the original variables, a general solution to the homogeneous equation is

$$v_H(x) = x^{m+1}(1-x)^{-m}\left(CU\left(m, n, \frac{2\alpha}{k^2}\frac{x}{1-x}\right) + DM\left(m, n, \frac{2\alpha}{k^2}\frac{x}{1-x}\right)\right).$$

Now, Kummer's function $M(m, n, t)$ diverges as $t \to \infty$. As the leading term of v_H also diverges in this limit, all bounded solutions satisfy $D = 0$. This leaves solutions to the ODE of the form

$$v(x) = \frac{r_G + b}{\rho + \alpha}x - \frac{b}{\rho} + Cx^{m+1}(1-x)^{-m}U\left(m, n, \frac{2\alpha}{k^2}\frac{x}{1-x}\right)$$

as candidates for solving the desired boundary value problem.

Let $w(x) \equiv x^{m+1}(1-x)^{-m}U\left(m, n, \frac{2\alpha}{k^2}\frac{x}{1-x}\right)$. Since $U(m, n, \cdot)$ is a strictly positive function on $(0, \infty)$ when $m > 0$, w is strictly positive and the function $\Gamma(x) \equiv \log w(x)$ is well-defined on $(0, 1)$. As $U(m, n, \cdot)$ is an analytic function on $(0, \infty)$, w and Γ are C^2 functions on $(0, 1)$.

Lemma 11. $\lim_{x\to 1} w(x) = \left(\frac{k^2}{2\alpha}\right)^{\beta-1}$, $\lim_{x\to 1} \Gamma'(x) = -\frac{\rho}{\alpha}$ *and* $\liminf_{x\to 0} \Gamma'(x) = -\infty$.

Proof. Let $\gamma \equiv \frac{k^2}{2\alpha}$ and $z \equiv \gamma^{-1}\frac{x}{1-x}$. Then $w(x)$ may be written

$$w(x) = x\gamma^m z^m U(m, n, z).$$

Now, U has the third-order asymptotic expansion

$$U(m, n, z) = z^{-m}\left(a_0 - a_1 z^{-1} + \frac{1}{2}a_2 z^{-2} + \Phi(z)\right),$$

where $a_0 = 1$, $a_1 = m(m-n+1)$, $a_2 = m(m+1)(m-n+1)(m-n+2)$, and $\Phi(z) \sim O(z^{-3})$. In other words, $\lim_{z\to\infty} z^N\Phi(z) = 0$ when $N < 3$. As U is analytic, so is Φ, and L'hopital's rule implies $\Phi'(z) \sim O(z^{-4})$ and $\Phi''(z) \sim O(z^{-5})$.

Now insert this expansion into w to obtain

$$w(x) = x\gamma^m\left(a_0 - a_1 z^{-1} + \frac{1}{2}a_2 z^{-2} + \Phi(z)\right).$$

Taking $x \to 1$ implies $\lim_{x\to 1} w(x) = a_0\gamma^m$. Next, differentiate the asymptoptic expansion wrt x, noting that $\frac{dz}{dx} = \frac{\gamma^{-1}}{(1-x)^2} = \gamma^{-1}(1 + \gamma z)^2$. The result is

$$w'(x) = \gamma^m\left(a_0 - a_1 z^{-1} + \frac{1}{2}a_2 z^{-2} + \Phi(z)\right) + x\gamma^{m-1}(1+\gamma z)^2\left(a_1 z^{-2} - a_2 z^{-3} + \Phi'(z)\right).$$

Thus

$$\lim_{x\to 1} w'(x) = \gamma^m a_0 + \gamma^{m+1} a_1 = \left(\frac{k^2}{2\alpha}\right)^m + \left(\frac{k^2}{2\alpha}\right)^{m+1} m(m-n+1).$$

Written explicitly in terms of model parameters, this is

$$\lim_{x\to 1} w'(x) = \left(\frac{k^2}{2\alpha}\right)^{\beta-1} \left(1 + \frac{k^2}{2\alpha}(\beta-1)\left(-\beta + \frac{2\alpha}{k^2}\right)\right) = \left(\frac{k^2}{2\alpha}\right)^{\beta-1} \left(\frac{k^2}{2\alpha}\left(1+\frac{2\alpha}{k^2}\right)\beta - \frac{k^2}{2\alpha}\beta^2\right).$$

Using the quadratic formula characterizing β, this simplifies to

$$\lim_{x\to 1} w'(x) = -\left(\frac{k^2}{2\alpha}\right)^{\beta-1} \frac{\rho}{\alpha},$$

The limiting expression for $\Gamma'(x) = w'(x)/w(x)$ as $x \to 1$ may then be obtained by combining this expression with the limiting expression for $w(x)$.

As for the second limit, rewrite $w(x)$ as

$$w(x) = (1-x)\gamma^m z^{m+1} U(m,n,z).$$

When $m > 0$ and $n > m+2$, $\lim_{z\to 0} z^{m+1}U(m,n,z) = \infty$, so $\lim_{x\to 0} w(x) = \infty$ and $\lim_{x\to 0}\Gamma(x) = \infty$. Now suppose by way of contradiction that $\liminf_{x\to 0}\Gamma'(x) > -\infty$. In this case there exists an $x > 0$ and an $M < \infty$ such that $\Gamma'(x) \geq -M$ for all $y \in (0,x]$. The fundamental theorem of calculus then implies

$$\Gamma(y) = \Gamma(x) - \int_y^x \Gamma'(t)\,dt \leq \Gamma(x) + M(x-y) \leq \Gamma(x) + Mx.$$

Thus $\Gamma(y)$ is bounded above on $(0,x]$, contradicting $\lim_{x\to 0}\Gamma(x) = \infty$. $\square$

The fact that w is finite in the limit as $x \to 1$ establishes that for any choice of $\underline{\pi}$ and C, v is bounded on $[\underline{\pi}, 1)$. The remaining limits of Γ' from the previous lemma provide the necessary tools to demonstrate the existence of constants $\underline{\pi} \in (0,1)$ and $C > 0$ such that v satisfies $v(\underline{\pi}) = v'(\underline{\pi}) = 0$. These boundary conditions may be equivalently written $Cw(\underline{\pi}) = -\frac{r_G+b}{\rho+\alpha}\underline{\pi} + \frac{b}{\rho}$ and $Cw'(\underline{\pi}) = -\frac{r_G+b}{\rho+\alpha}$. Dividing the second equation through by the first, existence of an appropriate $\underline{\pi}$ is equivalent to existence of a solution to

$$\Gamma'(x) = \phi(x),$$

where $\phi(x) \equiv -\frac{1}{\pi^*-x}$ and $\pi^* \equiv \frac{b}{b+r_G}(1+\alpha/\rho)$. If a solution $\underline{\pi}$ to this equation exists, the corresponding constant of integration is $C = -\frac{1}{w'(\underline{\pi})}\frac{r_G+b}{\rho+\alpha}$. And if $\Gamma'(\underline{\pi}) < 0$, then given strict positivity of w it must be that $w'(\underline{\pi}) < 0$ and so $C > 0$.

Suppose first that $\pi^* \leq 1$. In this case $\lim_{x\uparrow\pi^*}\phi(x) = -\infty$ while $\phi(0) = -1/\pi^*$. Given the continuity of ϕ, it must be bounded on $[0, x_0]$ for any $x_0 < \pi^*$. Then given $\liminf_{x\downarrow 0}\Gamma'(x) = -\infty$, there exists an $x_1 \in (0, \pi^*)$ such that $\Gamma'(x_1) < \phi(x_1)$. And given that $\lim_{x\to 1}\Gamma'(x)$ is finite and Γ' is continuous on $(0,1)$, Γ' must be bounded on the closed interval $[x_1, \pi^*]$. So there exists an $x_2 \in (x_1, \pi^*)$ such that $\Gamma'(x_2) > \phi(x_2)$. Then as $\Gamma'(x) - \phi(x)$ is a continuous function on $[x_1, x_2]$, by the intermediate value theorem there exists an $\underline{\pi} \in (x_1, x_2)$ such that $\Gamma'(\underline{\pi}) = \phi(\underline{\pi})$. Further, $\Gamma'(\underline{\pi}) < 0$ given the negativity of ϕ on $[0, \pi^*)$.

Now suppose that $\pi^* > 1$. In this case ϕ is continuous and decreasing on $[0,1]$, with $\phi(1) = -1/(\pi^* - 1)$. The argument of the previous paragraph continues to establish existence of an $x_1 \in (0,1)$ such that $\Gamma'(x_1) < \phi(x_1)$. Meanwhile $\lim_{x\to 1}\Gamma'(x) = -\rho/\alpha$, while

$$\phi(1) = -\frac{1}{\frac{b}{b+r_G}(1+\alpha/\rho) - 1} < -\frac{\rho}{\alpha}.$$

Thus $\lim_{x\to 1}\Gamma'(x) > \phi(1)$, and so there exists an $x_2 \in (x_1, 1)$ such that $\Gamma'(x_2) > \phi(x_2)$. The intermediate value theorem then ensures existence of a solution $\underline{\pi} \in (x_1, x_2)$ to $\Gamma'(x) = \phi(x)$. Given the negativity of ϕ, it must be that $\Gamma'(\underline{\pi}) < 0$.

Finally, I establish that v is non-negative and $\underline{\pi} \leq b/(b + r_G)$. For $x \in [0,1]$, let $\mathbb{P}^x$ be the probability measure on $(\Omega, \mathcal{F})$ satisfying:

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

Under $\mathbb{P}^x$ the initial probability that the state is Good is x while the conditional state transition rate is α. Denote expectations wrt this measure by $\mathbb{E}^x$.

Let $\pi^x_t \equiv \mathbb{E}^x[\mathbf{1}\{\Lambda > t\} \mid \mathcal{F}^Y_t]$, and define $\tau^* \equiv \inf\{t \,:\, \pi^x_t \leq \underline{\pi}\}$. The proof of Lemma 9 establishes that

$$v(x) = \mathbb{E}^x\left[\int_0^{\tau^*} e^{-\rho t}(\pi^x_t r_G - (1-\pi^x_t)b)\,dt\right].$$

Suppose by way of contradiction that $\underline{\pi} > b/(b + r_G)$. Then $v(x) > 0$ for all $x > b/(b + r_G)$ given that the integrand in the previous expression for $v(x)$ is always strictly positive prior to τ^* and $\tau^* > 0$ a.s. In particular, $v(1) > 0$. Then inserting the boundary conditions into the HJB equation implies $v''(\underline{\pi}) < 0$, meaning that v and v' are both strictly negative for x close enough to $\underline{\pi}$. Let $\overline{\pi} = \inf\{x > \underline{\pi} \ : \ v'(x) \geq 0\} > \underline{\pi}$. Because $v(1) > 0$, it must be that $\overline{\pi} < 1$. By continuity, $v'(\overline{\pi}) = 0$. And as $v'(x) < 0$ on $(\underline{\pi}, \overline{\pi})$, it must be that $v(\overline{\pi}) < 0$ as well. But then from the HJB equation $v''(\overline{\pi}) < 0$, contradicting the negativity of v' for $x < \overline{\pi}$. So $\underline{\pi} \leq b/(b + r_G)$.

Now, suppose by way of contradiction that v is not non-negative on $[\underline{\pi}, b/(b+r_G)]$. Define $x^* \equiv \inf\{x \geq \underline{\pi} \ : \ v(x) < 0\} < b/(b + r_G)$. By continuity $v(x^*) = 0$, and since $v(x) < 0$ for x sufficiently close to x^*, it must be that $v'(x^*) \leq 0$. As $x^* < b/(b + r_G)$, the HJB equation implies that $v''(x^*) > 0$, implying $v'(x) > 0$ for x close to x^*. But then $v(x) > 0$ for x close to x^*, a contradiction. So it must be that $v(x) \leq 0$ for $x \leq b/(b + r_G)$. So let $\tau^\dagger = \inf\{t \ : \ \pi_t \leq b/(b + r_G)\}$. Given $\tau^\dagger \leq \tau^*$, the proof of Lemma 9 establishes that

$$v(x) = \mathbb{E}^x \left[\int_0^{\tau^\dagger} e^{-\rho t}(\pi_t r_G - (1 - \pi_t)b)\, dt + e^{-\rho \tau^\dagger} v(\pi_{\tau^\dagger}) \right].$$

Since both terms in the expectation are non-negative, $v(x) \geq 0$ everywhere. $\square$

C Proofs of results from the text

C.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.

C.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 $\mathbb{P}\{\tau[\Lambda] > \Lambda\} = \mathbb{P}^\Lambda\{\tau > \Lambda\} > 0$.

C.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 termination fee 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.

C.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 $\{\tau^Y \le \Lambda\}$ 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 $F_{\Lambda\wedge\tau^Y} \ge F^*_{\Lambda\wedge\tau^Y}$ 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

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

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.

C.5 Proof of Lemma 5

Fix any reporting policy $\Lambda' \ge \Lambda$. 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 $\{\Lambda' \le \tau^Y\}$ this identity follows from the definition of $\mathbb{E}^B_t$ and the fact that $\Lambda' \ge \Lambda$. 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.

C.6 Proof of Proposition 1

Fix an $\mathbb{F}^Y$-stopping time τ^Y, and let F^* be the fee 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

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

C.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 $\tau^{(x,t)} \leq \widetilde{\tau} \leq \tau^{(x,t)} + \varepsilon$ 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),

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

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

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

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, $\tau' \leq \tau^Y$, and $\tau' \leq \tau^{(x,t)} + \varepsilon$. In particular, note that if $\tau^Y \geq \tau^{(x,t)}$, then also $\tau' \geq \tau^{(x,t)}$.

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 $\tau^{(x,t)} \le \widetilde{\tau}^n \le \tau^{(x,t)} + 1/n$ 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 $v(x,t) \le v^{\tau^{(x,t)}}$. 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, $\tau^* \ge \tau^{(x,t)}$ 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 $\widetilde{\tau}^{(x,t)} \equiv \inf\{t \ : \ \pi^{(x,t)} \le b/(b+r_G)\}$. By right-continuity of $\pi^{(x,t)}$, $\widetilde{\tau}^{(x,t)} > 0$. Hence $v^{\widetilde{\tau}^{(x,t)}}(x,t) > 0$, a contradiction. So $\underline{\pi}(t) \le b/(b+r_G)$ for all time.

C.8 Proof of Lemma 6

Written in integral form, the filtering equation for beliefs is

$$\pi_{t'} = \pi_t - \int_t^{t'} \frac{\pi_{s-}}{1 - H(s-)}\, dH(s) + \int_t^{t'} \frac{(\lambda^B - \lambda^G)\pi_{s-}(1-\pi_{s-})}{\lambda^G \pi_{s-} + \lambda^B(1-\pi_{s-})}\, d\overline{Z}_s.$$

As H is a monotone function, it has an a.e. derivative h satisfying $\int_a^b h(s)\, ds \le \int_a^b dH(s)$ for every $a < b$. Therefore

$$\pi_{t'} \le \pi_t - \int_t^{t'} \frac{\pi_{s-} h(s)}{1 - H(s-)}\, ds + \int_t^{t'} \frac{(\lambda^B - \lambda^G)\pi_{s-}(1-\pi_{s-})}{\lambda^G \pi_{s-} + \lambda^B(1-\pi_{s-})}\, d\overline{Z}_s.$$

Now, the increments of the innovation process may be written

$$d\overline{Z}_s = (\lambda^G \pi_{s-} + \lambda^B(1-\pi_{s-}))\, ds - dN_s,$$

where N is an inhomogeneous Poisson counting process with rate $\lambda^G \mathbf{1}\{\Lambda > t\} + \lambda^B \mathbf{1}\{\Lambda \le t\}$. Thus

$$\pi_{t'} \le \pi_t + \int_t^{t'} \pi_{s-}\left(-\frac{h(s)}{1 - H(s-)} + (\lambda^B - \lambda^G)(1-\pi_{s-}) \right) ds.$$

Now, the value of the remaining integral is unchanged by modifying the integrand on a measure-zero set of points. As H has at most countably many discontinuities, I may replace

$H(s-)$ by $H(s)$, and use the essential infimum α of $h/(1-H)$ to obtain the bound

$$\pi_{t'} \leq \pi_t + \int_t^{t'} \pi_{s-} \left(-\alpha + (\lambda^B - \lambda^G)(1 - \pi_{s-})\right) ds.$$

Now, suppose the inequality in the lemma statement holds. Define $z_t \equiv \pi_t/(1-\pi_t)$, and consider any time t such that $z_t \leq b/r_G$. I claim that $z_{t'} \leq b/r_G$ for all $t' > t$. For suppose $z_{t'} > b/r_G$ for some $t' > t$. Let $t'' \equiv \sup\{s < t' \ : \ z_s \leq b/r_G\} \geq t$. Because z has left limits everywhere and possesses only downward jump discontinuities, $z_{t''} \leq z_{t''-} \leq b/r_G$ and so $t'' < t'$. Now, $z_s > b/r_G$ on $(t'', t']$ by construction, hence by the inequality in the lemma statement $-\alpha + (\lambda^B - \lambda^G)(1 - \pi_{s-}) \leq 0$ on (t'', t'). But then using the upper bound on future beliefs derived above, it must be that $z_{t'} \leq z_{t''}$, contradicting $z_{t'} > r_G/b \geq z_{t''}$.

Hence once $\tau^\dagger = \inf\{t \ : \ z_t \leq b/r_G\}$ is reached, virtual profits are non-positive forever after. Thus $\tau^\dagger$ must be an optimal stopping rule.

C.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.

C.10 Proof of Proposition 3

Let $\mathcal{D} = \{d_1, ..., d_n\}$. By assumption, each Y^θ is a Lévy process. Then under the hypothesis in the proposition statement, each Y^θ may be written

$$Y_t^\theta = r_\theta t + \sigma_\theta Z_t^\theta + \sum_{i=1}^n d_i (N_i^\theta(t) - \lambda_i^\theta t),$$

where Z^θ is a standard Brownian motion and N_i^θ are Poisson processes with rates λ_i^θ, where $Z^\theta, N_1^\theta, ..., N_n^\theta$ are mutually independent. If $\sigma_G \neq \sigma_B$, then $\pi_t = 0$ for all $t > 0$ $\mathbb{P}^B$-a.s., as differences in quadratic variation of Brownian motions are immediately detectable. In this case τ^Y is deterministic under $\mathbb{P}^B$, and so U is trivially a $\mathbb{P}^G$-martingale. So assume $\sigma_G = \sigma_B = \sigma$. I will also maintain $\sigma > 0$, with the $\sigma = 0$ case being an easy modification of what follows.

Let $Y_t^c \equiv Y_t - \sum_{s \le t} \Delta Y_s$ be the continuous part of Y, 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 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

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

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 $\{t \leq \tau^Y\}$.

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 $(Y_s)_{s \geq t}$. Thus τ^Y is increasing in π_t conditional on $(Y_s)_{s\geq t}$. And under $\mathbb{P}^B$, $(Y_s)_{s \geq t}$ is independent of π_t, since $(Y_s)_{s\geq t}$ is identical in law to $(Y_s^B)_{s \geq t}$ 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 $(\widetilde{r}_G - \widetilde{r}_B)\phi_0 \geq 0$ and $(\lambda_i^G - \lambda_i^B)\phi_i \geq 0$ for all $i \geq 1$ $\mathbb{P}^B$-a.e. On $\{(t,\omega) : t > \tau^Y(\omega)\}$ this is trivial, as U is a constant process for $t \geq \tau^Y$ and so every ϕ_i must be zero a.e. on this set. So consider the claim on $\{(t,\omega) : t \leq \tau^Y(\omega)\}$. 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

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

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.

C.11 Proof of Lemma 8

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 $\underline{f}$.