Optimal project termination with an informed agent

Erik Madsen*

October 10, 2016

Abstract

I study optimal incentive contracting when a temporary project is overseen by a firm with access to two channels of information about the project's state - a noisy public news process, and the reports of a manager possessing superior private information. The firm faces the decision problem of when to wrap up the project, as well as the agency problem of a manager with borrowing constraints and incentives for delayed termination. I develop techniques for solving the resulting novel mechanism design problem, which abstracts from time-zero asymmetric information but allows the agent to extract dynamic information rents. The optimal contract features a project deadline, which induces occasional inefficient early termination but no late termination; and a golden parachute, paid upon termination, which declines as the deadline approaches and is exhausted when the deadline is reached. Further, both the deadline and parachute exhibit sensitivity to news which varies with proximity to the deadline. Initially incentive pay is low-powered and the contract imposes a hard deadline; while late in the project incentive pay is high-powered and the deadline becomes porus.

1 Introduction

Economic decision-makers commonly face the problem of determining when to wrap up temporary projects on the basis of noisy status reports. A classic example is a firm operating

*Department of Economics, New York University. Email: emadsen@nyu.edu. 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, Ben Golub, Brett Green, Johannes Hörner, Chad Jones, Nicolas Lambert, Edward Lazear, Michael Ostrovsky, Takuo Sugaya, Juuso Toikka, and Jeffrey Zwiebel for valuable discussions and feedback.

a depreciating plant which must eventually be retired or overhauled. The depreciation can take various forms, 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.

In these settings the decision-maker faces an optimal stopping problem of when to terminate the project given the history of the public signal. In practice, many such decisions are made with the assistance of agents who bring additional expertise to the problem. 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, he is likely to maintain superior knowledge of the plant's state. 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 her health status and day to day performance.

The firm would clearly benefit by incorporating the agent's superior information into the project termination decision. However, agents 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 agent must be compensated for revealing news which instigates early project termination, the firm may instead simply choose to set a project deadline to economize on payments. In this paper I study the optimal incorporation of the agent's information into the firm's termination decision when the project's state is binary (either good or bad) and the bad state is absorbing.[1]

I assume the project is commonly known to begin in the good state and that the agent observes only the project's current state. I therefore abstract from the screening of time-zero private information. In dynamic mechanism design settings with fully flexible transfers, this

[1]The single-person optimal stopping problem with a binary state and an absorbing bad state has been well-studied in operations research as the problem of "quickest detection" or "change-point detection"; see, e.g., Peskir and Shiryaev (2006) for a textbook treatment.

abstraction would yield an uninteresting contracting problem. For in such settings, agents extract no information rents for private information received after time zero. In fact, the form of an optimal contract is the same as in a setting where all (orthogonalized) information received by the agent after time zero is public. (See Eső and Szentes (2016) for a proof.) In the context of my model, an optimal contract under flexible transfers would simply sell the agent the project: the firm would charge the agent the NPV of the project's output plus rents and then rebate (or charge) him the realized flow of output.[2] By contrast, when the agent is capital-constrained he cannot buy into the project to compensate the firm for later payments that elicit his dynamic private information.

The contractual design problem in this paper is therefore quite different from the one faced in standard mechanism design models. In lieu of separating time-zero types, the firm's job is to dynamically screen low types who arrive over time. It does so by using payments tied to public news to distinguish good types from bad, and using termination deadlines to limit the scope of a bad type's information rents. An optimal contract sets deadlines to trade off efficiency versus information rents; and sets sensitivity to public news to trade off average project duration versus risk of early project termination. These tradeoffs produce rich project dynamics even in an environment with no risk-sharing or heterogeneous time-preference concerns.

I find that an optimal termination policy takes the simple form of a threshold rule in "virtual beliefs." At each moment the firm asks the agent to report on the current state of the project, but performs Bayesian updating about whether the state has transitioned as if no agent were present. It then terminates the project the first time either the agent reports that the state has transitioned or the firm's virtual beliefs drop below a (time- and history-independent) threshold. This threshold is strictly lower than the optimal threshold in the firm's problem without an agent. The firm is therefore able to partially utilize the agent's information, fully eliminating late terminations and partially mitigating early terminations that it would have incurred in the agent's absence. The optimal threshold also satisfies clean comparative statics: it is decreasing in the informativeness of news about the state, increasing in the rate of state transitions, and increasing in the severity of the agent's incentive misalignment.

Optimal payments take a similarly simple form: all payments are deferred until project termination, at which point the agent receives a lump sum "golden parachute." If the agent

[2]This contractual form is not unique - while selling the entire project is one way to achieve efficiency, an alternative is to merely buy the agent out of his rental rights. So another optimal contract charges the agent for the NPV of his stream of rents from perpetual operation of the project, and then rebates him the same amount at the time of project termination.

reports a state switch, the optimal golden parachute is exactly the expected discounted stream of project rents he would have received by never reporting the switch. On the other hand, if the project is terminated before a reported state switch he is paid nothing. The dynamics of the optimal golden parachute may be characterized by a one-dimensional HJB equation which determines how fast the parachute drifts down over time and how sensitively the parachute reacts to news about output.

Methodologically, my analysis reveals an unexpected connection between the approaches used to analyze dynamic problems of hidden information and hidden actions. I pursue two distinct paths to characterizing an optimal contract, one which derives the optimal dynamics of the golden parachute process and the other which determines the optimal incidence of early termination. As the two contracting variables control the cost of implementing the contract and the amount of output produced, respectively, I refer to these procedures as "price" and "quantity" approaches to optimization. The price approach yields a recursive problem with a continuation promise as the state variable, and is therefore the familiar route for practitioners of contracting under repeated moral hazard. A distinctive feature of my setting is a non-standard free boundary condition at a singularity of the HJB equation, a technically challenging feature not typically present in repeated moral hazard problems. Meanwhile, the quantity approach is conceptually analogous to the virtual valuation approach used in standard mechanism design settings, though requiring different mathematical tools to carry through. Each yields complementary insights into the form of an optimal contract - the price approach describes the tradeoff between hard and soft deadlines to incentivize truthful reporting, while the quantity approach clarifies the real option problem underpinning the optimal incidence of inefficient project termination.

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.[3] Papers in this tradition can be thought of as extending the canonical static model of Myerson (1981) to multi-period settings, though

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

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 et al. (2014) extend these results to an infinite horizon discrete-time setting 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.[4] By contrast, in my model the agent's type is payoff-relevant only to the principal and so tradeoffs between payments and allocations cannot separate different types. Instead, the principal observes public signals correlated with type and uses them to tie the agent'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 Eső and Szentes (2016), under fully flexible transfers the agent receives no information rents for any 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 no time-zero private information.

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 this paper, 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

[4]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).

[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.

by limited liability but can extract rents via time-zero private information. This distinction 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 studying 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 (2015) considers an environment with public experimentation by an agent who has private prior beliefs about the true state and a bias toward experimentation. She derives an optimal policy strikingly 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 et al. (2015) 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.

Although my paper abstracts from problems of moral hazard, it draws heavily on recursive techniques developed in the literature on incentive contracting under repeated moral hazard. I briefly discuss relevant entries in that literature and their relationship to the techniques used in this paper. Spear and Srivastava (1987) formulate the standard infinite-horizon recursive problem in discrete time, while Sannikov (2008) and Williams (2008) study the problem in continuous time when output follows a Brownian motion with drift, affording considerable tractability compared to the discrete-time case. The continuous-time model has been extended in various ways which preserve its basic recursive structure. Williams (2008, 2015) permits general forms of hidden savings which impact the agent's marginal utility of consumption. Sannikov (2014) allows actions to have long-run impacts on output. Prat and Jovanovic (2013) build in symmetric ex ante uncertainty about the quality of the agent, as in the career concerns literature. Cvitanic et al. (2013) suppose instead that the agent possesses ex ante private information about his marginal cost of effort, a setting which may lead the principal to offer a menu of contracts to separate or screen agents of different types as in mechanism design.

These papers solve the optimal contracting problem using recursive techniques with the agent's continuation utility as a state variable and eliminate effort from the problem using a first-order condition. Their basic technique is therefore in the same spirit as the price approach of this paper, although with important distinctions. In particular, in moral hazard models the agent's continuation utility process is automatically a martingale on-path, which is not true of the natural state variable in my model; so additional representation results are required to analyze my problem recursively. Also, while in moral hazard models the incentive constraint for effort has a natural local formulation, an analogous result for the binding incentive constraint in my model holds only for a restricted class of candidate contracts; verification of optimality is therefore more difficult in my setting.

Papers in this literature generally focus on risk-sharing and the wedge between induced and first-best effort rather than the incidence of early termination, which is the key inefficiency in my model. A notable exception is DeMarzo and Sannikov (2006), which abstracts from both risk-sharing (the agent is risk-neutral) and effort distortions (marginal cost of effort is constant and less than 1, so efficient effort is always implementable and optimal for the principal). They then characterize a simple implementation of the optimal contract in the presence of limited liability and hidden savings. They show that an optimal contract can be implemented via a constant-limit credit line which the agent may draw down at will, with the caveat that the project is terminated when the credit line is exhausted. This dynamic has close similarities to the optimal payment process in my paper, which features a golden parachute that can be taken by the agent at any time and which when exhausted due to poor project returns leads to termination. Further, their contracting problem boils down to the one-dimensional choice of an optimal credit limit, just as my problem can be reduced (via the quantity approach) to the choosing of a one-dimensional threshold in belief space.

2 The model

2.1 The environment

A firm possesses a project with a limited but uncertain lifespan and must decide when to scrap the project.[6] The project's lifespan Λ is unobserved by the firm, who has prior beliefs that Λ is exponentially distributed with failure rate α.[7] The state variable θ tracks the status of the project over time, and is defined as $\theta_t = G$ when $t < \Lambda$ and $\theta_t = B$ otherwise. Then from the perspective of the firm θ is a hidden Markov process with initial state $\theta_0 = G$ and transition rate α from G (the "good state") to B (the "bad state"). The project generates average profits r_θ per unit time when the current state is θ, with $r_G > 0 > r_B$. I will let $\Delta r \equiv r_G - r_B$ denote the spread in average profits between the two states.

Random variability in instantaneous flow profits prevents the firm from directly observing the current state of the project. For most of this paper, I study a project which generates profits in continuous time in Brownian increments. Letting Y_t be cumulative profits up to time t, I assume that the incremental contribution dY_t to profits over the time interval $[t, t + dt)$ is

$$dY_t = r_{\theta_t}\, dt + \sigma\, dZ_t,$$

where Z is a standard Brownian motion independent of Λ.[8]

Formally, I model the exogenous uncertainty of this setting with a canonical probability space $(\Omega, \mathcal{F}, \mathbb{P})$ sufficiently rich to admit Z and Λ, with $\mathbb{E}$ the expectation operator under $\mathbb{P}$. The cumulative profit process Y is then defined in terms of Z and Λ as

$$Y_t = r_G(t \wedge \Lambda) + r_B(t - t \wedge \Lambda) + \sigma Z_t.$$

I let $\mathbb{F}^Y = \{\mathcal{F}^Y_t\}_{t\geq 0}$ denote the natural filtration of $\mathcal{F}$ generated by Y; this filtration captures the information available to the firm from its observation of past profits. I write $\mathbb{E}^Y_t$ for the conditional expectation under $\mathbb{P}$ given $\mathcal{F}^Y_t$. I also let $\mathbb{F} = \{\mathcal{F}_t\}_{t\geq 0}$ denote the natural filtration of $\mathcal{F}$ generated by both Y and θ, with $\mathbb{E}_t$ the conditional expectation under $\mathbb{P}$ given $\mathcal{F}_t$. Finally, I define $\mathbb{P}^G$ and $\mathbb{P}^B$ to be the equivalent probability measures under which

[6] It does not matter whether the decision to scrap is reversible, as the firm learns nothing further about the project after it has ended.

[7] The analysis of this paper can be extended straightforwardly to accommodate arbitrary beliefs about Λ, at the cost of introducing explicit time dependence into the optimal contract. I abstract from such considerations to focus on dynamics arising endogenously from incentive frictions.

[8] I have also analyzed a variant of the model in which the state of the project impacts the arrival rate of Poisson profit shocks rather than the drift rate of profits. The mathematical analysis and qualitative conclusions are very similar to the Brownian case.

$Z_t^G \equiv \frac{1}{\sigma}(Y_t - r_G t)$ and $Z_t^B \equiv \frac{1}{\sigma}(Y_t - r_B t)$, respectively, are standard Brownian motions independent of Λ, with Λ's marginal distribution the same as under $\mathbb{P}$. I then write $\mathbb{E}_t^G$ and $\mathbb{E}_t^B$ for the conditional expectations under these measures given $\mathcal{F}_t^Y$.

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

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

where $\pi_t = \mathbb{P}_t^Y\{\theta_t = G\} = \mathbb{E}_t^Y[\mathbf{1}\{\theta_t = G\}]$ are the firm's posterior beliefs at time t about the current state of the project, based on its observation of past profits.

The firm may hire an agent to oversee the project and monitor its state. The agent is an expert who costlessly and privately observes the state process θ. He cannot, however, observe Λ in advance. That is, the filtration $\mathbb{F}$ captures the information available to the agent at any time. The agent enjoys intrinsic benefits from employment, in the form of flow benefits $b > 0$ per unit time while the project operates, regardless of its state.[9] I assume that $r_B + b < 0$, so that it is jointly unprofitable for the firm to operate the project in the bad state.[10] The agent possesses limited liability and has no initial wealth, so cannot be sold the project. He is risk-neutral and possesses the same discount rate ρ as the firm.

2.2 Contracts

The firm commits in advance to a contract eliciting reports from the agent over time and specifying a termination policy τ and a cumulative payment function Φ, which may condition on the public history of output and the agent's reports. By the revelation principle, I restrict attention to contracts which at each moment in time elicit a report $\widetilde{\theta}_t \in \{G, B\}$ of the current state. Equivalently, as the only information to communicate is the time of a state switch, the firm asks the agent to make a single report at the time of the switch. This formulation yields a contract space formally very similar to the space of mechanisms in a static mechanism design problem.[11]

[9]None of the results of this paper are impacted if the agent's flow benefits are 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 in the good state are slack at the optimum. Thus the optimal contract is independent of b_G.

[10]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 agent's reports. Essentially, when $r_B + b > 0$ there is scope for gains from trade by operating the project in the bad state, but because the agent has no wealth to pay the firm none of these gains are realized.

[11]In what follows, I do not allow contracts to condition on any exogenous randomization. As we shall see, this is without loss of generality, because the firm's objective function exhibits a natural concavity in an

Definition 1. *A* revelation contract $\mathcal{C} = (\Phi, \tau)$ *is a family of stochastic processes* $\Phi[t]$ *and random variables* $\tau[t]$ *for each* $t \in \mathbb{R}_+ \cup \{\infty\}$ *such that:*

- *Each* $\Phi[t]$ *is real-valued,* $\mathbb{F}^Y$*-adapted, right-continuous, non-decreasing, and satisfies* $\Phi[t]_0 = 0$*,*
- *Each* $\tau[t]$ *is an* $\mathbb{F}^Y$*-stopping time,*
- *For every* t, t', s *such that* $t' > t > s$*,* $\Phi[t']_s = \Phi[t]_s$ *and* $\mathbf{1}\{\tau[t'] \leq s\} = \mathbf{1}\{\tau[t] \leq s\}$*,*
- *For every* t, s *such that* $s > \tau[t]$*,* $\Phi[t]_s = \Phi[t]_{\tau[t]}$*.*

Intuitively, a revelation contract specifies a payment $\Phi[t]$ and an allocation $\tau[t]$ as a function of the reported type t, as in static mechanism design. However, in this setting payments and allocations are dynamic variables that can be conditioned on output history observed both before and after the agent's report. Definition 1 imposes an appropriate *no-foresight* property in this context, by requiring that the contract not "look ahead" and anticipate the timing of a future report when making current payment and termination decisions. Additionally, I require that no transfers are made subsequent to termination. As both parties have linear utility with the same discount rate and no information arrives after termination, this restriction is without loss of generality.

A revelation contract effectively specifies payments and allocations given any *deterministic* report by the agent. Of course, in general the agent may condition his report on the history of output and the state. Formally, the set of admissible reporting policies is the set of $\mathbb{F}$-stopping times. Clearly, given $\omega \in \Omega$, the payment stream and termination time induced by a contract (Φ, τ) and reporting policy Λ' are $\Phi[\Lambda'(\omega)](\omega)$ and $\tau[\Lambda'(\omega)](\omega)$. It is therefore natural to extend $\Phi[\cdot]$ and $\tau[\cdot]$ to arbitrary policies via $\Phi[\Lambda'](\omega) \equiv \Phi[\Lambda'(\omega)](\omega)$ and $\tau[\Lambda'](\omega) \equiv \tau[\Lambda'(\omega)](\omega)$. To ensure that this procedure yields processes which conform to the filtration $\mathbb{F}$, I impose a technical regularity requirement on the contract space.

Definition 2. *A revelation contract* $\mathcal{C} = (\Phi, \tau)$ *is* admissible *if for each* $T \in \mathbb{R}_+$*, the maps* $(t, \omega) \mapsto \Phi[t]_T(\omega)$ *and* $(t, \omega) \mapsto \mathbf{1}\{\tau[t](\omega) \leq T\}$ *are* $\mathcal{B}([0, T]) \otimes \mathcal{F}_T^Y$*-measurable.*

Admissibility ensures that a contract contains consistent data for constructing outcomes under any reporting policy chosen by the agent. The following proposition demonstrates this formally.

Proposition 1. *Suppose the revelation contract* $\mathcal{C} = (\Phi, \tau)$ *is admissible. Then for each* $\mathbb{F}$*-stopping time* Λ'*,* $\Phi[\Lambda']$ *is* $\mathbb{F}$*-progressively measurable and* $\tau[\Lambda']$ *is an* $\mathbb{F}$*-stopping time.*

---

appropriate state variable.

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

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

Incentive-compatibility is then defined in the natural way:

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

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

for all $\mathbb{F}$*-stopping times* Λ'.

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

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

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 agent to operate the project in addition to monitoring it. Therefore the firm cannot terminate the agent without also ceasing operation of the project. [12]

3 Preliminaries

I begin the analysis by performing several preliminary explorations of the contract space. I establish baseline implementable outcomes against which an optimal contract must improve, identify a key subset of IC constraints expected to bind at the optimum, and characterize how a candidate optimal contract conditions on the timing of the agent's report. By the end of this section, I will have reduced the contracting problem to a form amenable to more penetrating attacks.

[12] I have also considered alternative assumptions under which the firm can replace the agent or continue operating the project on his own. The major qualitative features of an optimal contract remain unchanged, though unsurprisingly the firm chooses a more aggressive termination policy given his improved outside option.

3.1 Baseline contracts

I begin the analysis by describing a pair of IC contracts which frame the basic tradeoff faced by the firm in its design problem.

Remark. *Let* $\Phi[t]_s = \frac{b}{\rho}\mathbf{1}\{s \geq t\}$ *and* $\tau[t] = t$. *Then* (Φ, τ) *is an IC contract which is profit-maximizing among all IC contracts implementing efficient project termination.*

This contract stops the project whenever the agent reports that the state has changed, and pays the agent a lump-sum transfer of b/ρ upon termination. Intuitively, such a contract is incentive compatible because at each moment the additional flow rents $b\,dt$ from waiting a moment longer to report a state switch are exactly offset by the interest expense $-\rho(b/\rho)\,dt$ of delaying receipt of the contract's terminal payment. Thus the agent is indifferent between all reporting policies regardless of the true state transition time, and in particular is willing to report truthfully. And it is profit-maximizing in the class of efficient contracts because any IC contract implementing Λ must commit to operating forever if the agent never reports a state switch. Therefore it must pay the agent at least as much to stop as he'd collect by letting the project operate forever, which is precisely b/ρ.

Thus at one extreme, the firm can operate the project efficiently provided it pays the agent enough. The following remark considers the opposite extreme, in which the firm disregards the agent's reports entirely.

Remark. *Suppose* $\tau^\dagger$ *maximizes* $\mathbb{E}\left[\int_0^{\tau^Y} e^{-\rho t} dY_t\right]$ *among all* $\mathbb{F}^Y$*-stopping times* τ^Y. *Let* $\Phi[t] = 0$ *and* $\tau[t] = \tau^\dagger$. *Then* (Φ, τ) *is an IC contract which is profit-maximizing among all IC contracts utilizing only public information.*

Evidently, if the firm does not condition payments or project termination on the agent's reports, incentive compatibility is trivially satisfied. The firm's contracting problem is to improve upon $\tau^\dagger$ by extracting some of the agent's private information about the project's state, but to economize on transfers which compensate the agent for his lost benefit stream when the project is halted. Unsurprisingly, the firm can do better than either of the baseline contracts described above by striking a balance between efficient termination and the costs of information elicitation.

3.2 Relaxing the incentive constraints

Incentive compatibility boils down to the requirement that the agent not profit from falsely reporting a state change either too early or too late at any time following any history of output. In particular, when $\theta_t = G$ the agent's expected utility from waiting until Λ to report a state switch must exceed his expected utility from reporting a state switch immediately. And when $\theta_t = B$, the agent's expected utility from reporting the switch immediately must exceed the payoff of any delayed reporting strategy. It will be very helpful to characterize incentive compatibility as the combination of these two constraints.

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

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

for all $\mathbb{F}$*-stopping times* $\Lambda' \leq \Lambda$ *(respectively,* $\Lambda' \geq \Lambda$*).*

The IC-G and IC-B constraints collectively represent only a subset of the constraints required for full incentive-compatibility. Fortunately, the following lemma establishes that checking them is sufficient to establish that a contract is IC.

Lemma 1. *An admissible 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 lie slack under an optimal contract and can be dropped from the optimization problem. Intuitively, given the agent's preference to delay project termination absent contractual incentives, at least some of the IC-B constraints must bind. But as provisioning incentives is costly, the firm should impose them as lightly as possible; in particular, it would be surprising if the firm provisioned such strong incentives for reporting a switch on time that IC-G were violated and the agent could profit from reporting the switch prematurely. To validate this hypothesis, I solve the relaxed problem of characterizing an optimal IC-B contract, and then verify at the end that the resulting contract is indeed incentive-compatible and therefore an optimal contract.

3.3 Optimal usage of the agent's report

I now characterize how an optimal IC-B contract should respond to the timing of the agent's report. In general a contract could specify a complex series of payments to the agent over the course of his employment, both before and after his report; it could also utilize a termination policy which either stops before hearing the agent's report or runs long after the report is made. The following pair of lemmas prune away much of this complexity, establishing the optimality of a set of IC-B contracts which condition on the agent's report in a parsimonious way.

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:*

- $\tau'[t] = \tau[t] \wedge t = \tau[\infty] \wedge t$ *for all* $t \in \mathbb{R}_+ \cup \{\infty\}$*;*
- $\Pi[\mathcal{C}'] \geq \Pi[\mathcal{C}]$;
- $\Pi[\mathcal{C}'] > \Pi[\mathcal{C}]$ *if* $\mathbb{P}\{\tau[\Lambda] > \Lambda\} > 0$.

This lemma establishes that optimal IC-B contracts never terminate inefficiently late - that is, after the agent reports the state has switched. In principle late termination could be desirable as a way to compensate the agent with flow benefits for truthful reporting. However, the assumption that the project is jointly unprofitable in the bad state means that the firm can always compensate the agent more cheaply with a monetary transfer at the time of the state switch. The proof of the lemma exploits this observation, modifying an arbitrary contract by adding an additional transfer at the time of a report equal to the agent's expected flow benefits plus future transfers under the original contract. This change preserves IC-B while improving profitability whenever the original contract operated in the bad state.

Importantly, any contract terminating no later than the time of a report cannot condition on the report whenever the project is terminated early. Hence an optimal termination policy always consists of a rule of the form "Terminate the project as soon as the agent reports a state switch or a publicly observable threshold has been reached, whichever comes first."

Remark. *It is not true that full incentive-compatibility is always preserved by such a modification of the contract. In general IC-G may be enforced by monitoring output following a report to ensure that the state has really switched. If such monitoring is removed by truncating the project lifetime, an otherwise IC contract may create incentives to prematurely report a state switch. Lemma 2 therefore 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* $\tau[t] \leq t$ *for all* $t \in \mathbb{R}_+$*. Then there exists an* $\mathbb{F}^Y$*-progressively measurable stochastic process* $F \geq 0$*, with associated payment process* $\Phi'[t]_s = F_{\tau[t]}\mathbf{1}\{s \geq \tau[t]\}$*, such that* $\mathcal{C}' = (\Phi', \tau)$ *is an IC-B contract satisfying* $\Pi[\mathcal{C}'] = \Pi[\mathcal{C}]$.

This lemma demonstrates that all payments can be backloaded to a single *golden parachute*, denoted F, payable upon termination. Backloading is possible because both parties are risk-neutral and discount the future equally; thus expected profits and incentive-compatibility are undisturbed by deferring all promised payments until termination and accruing interest on them at rate ρ. And because termination occurs no later than the date of a reported state change, the size of the parachute need not be conditioned on the date of a past report, hence is $\mathbb{F}^Y$-adapted. It is then without loss of generality to study only transfer policies of the form "Pay nothing until project termination, at which point grant the agent a single lump sum golden parachute".

In light of Lemmas 2 and 3, I restrict attention to backloaded contracts with no late termination. Such contracts may be summarized by a pair (F, τ^Y), with F a golden parachute and τ^Y a public termination policy. Crucially, both F and τ^Y condition only on public information and not the reports of the agent. Lemma 3 establishes that an admissible contract with no later termination yields a $\mathbb{F}^Y$-progressively measurable golden parachute when backloaded. The following lemma establishes the converse result that every progressively measurable process can be utilized as the golden parachute of an admissible contract.

Lemma 4. *Suppose $F \geq 0$ is an $\mathbb{F}^Y$-progressively measurable process and τ^Y is an $\mathbb{F}^Y$-stopping time. Then letting $\tau[t] = \tau^Y \wedge t$ and $\Phi[t]_s = F_{\tau[t]}\mathbf{1}\{\tau[t] \leq s\}$ for every $t, s \in \mathbb{R}_+$ and $\tau[\infty] = \tau^Y, \Phi[\infty] = 0$, the contract $\mathcal{C} = (\Phi, \tau)$ is admissible.*

In light of this lemma, I equate the space of admissible contracts with the set of pairs (F, τ^Y) where F is $\mathbb{F}^Y$-progressively measurable and τ^Y is an $\mathbb{F}^Y$-stopping time.

3.4 Dynamic incentive-compatibility

Definition 4 characterizes IC-B from an ex ante point of view. It will be convenient for both intuition and proofs to recast the condition as a dynamic constraint on the size of the golden parachute. Recall that IC-B was motivated by the informal reasoning that under an IC contract, the agent should not want to delay the report of a state switch upon observing one. In other words, at each moment in time the size of the golden parachute should exceed the expected flow benefits plus deferred parachute earned under any delayed reporting strategy. The following lemma establishes that this condition is essentially necessary and sufficient for IC-B.

Lemma 5. *Let $\mathcal{C} = (F, \tau^Y)$ be an admissible contract. Suppose that for every $t \in \mathbb{R}_+$ and*

$\mathbb{F}^Y$-stopping time $\tau' \geq t$,

$$F_t \geq \mathbb{E}_t^B \left[\int_t^{\tau^Y \wedge \tau'} e^{-\rho(s-t)} b \, ds + e^{-\rho(\tau^Y \wedge \tau' - t)} F_{\tau^Y \wedge \tau'} \right] \tag{2}$$

on $\{\tau^Y > t\}$. Then $\mathcal{C}$ satisfies IC-B.

Conversely, suppose $\mathcal{C}$ is an IC-B contract. Then there exists another IC-B contract $\mathcal{C}' = (F', \tau^Y)$ such that $\mathbb{P}\{F'_{\tau^Y \wedge \Lambda} = F_{\tau^Y \wedge \Lambda}\} = 1$ and F' satisfies (2) for every $t \in \mathbb{R}_+$ and $\tau' \geq t$ on $\{\tau^Y > t\}$.

Remark. *Clearly the modified contract $\mathcal{C}'$ yields the same expected profits as $\mathcal{C}$.*

The inequality in the lemma statement is just a formal expression of the condition that the golden parachute at time t exceed the expected utility, conditional on the state having switched, of any delayed reporting strategy. The first half of the lemma shows that this condition holding at all times prior to exogenous termination is sufficient for IC-B. It is not quite necessary, as IC-B need hold only in ex ante expected terms, so the condition evaluated at $t = \Lambda$ could fail on a measure-zero subset of Ω. However, the second half of the lemma shows that every IC-B contract can be modified so that the condition holds for all t and Ω, and this modification preserves the termination policy and pays out the same golden parachute under truthful reporting a.s. Without loss of generality, I restrict attention to IC-B contracts satisfying (2) going forward and treat the condition as necessary as well as sufficient without comment.

4 Price and quantity elements of contracting

Over the next several sections I develop two alternative approaches to solving the firm's contracting problem. In this section I show how to eliminate either the payment scheme or termination policy from the problem, reducing the analysis to design of a single contractual element of one's choosing. The freedom to choose between optimizing payments versus termination times is closely analogous to the equivalence between choosing prices or quantities in classical monopoly theory, a perspective I return to later in the section. I go on to develop the "price approach" of optimizing the payment scheme in Section 5, and then pursue the alternative "quantity approach" of designing the termination policy in Section 6. The two approaches provide distinct and complementary insights into properties of the optimal contract. In particular, the price approach characterizes how the firm should trade off hard and soft deadlines when incentivizing truthful reporting. It is the better approach for understanding how incentives evolve over the lifetime of the contract. Meanwhile the quantity approach pinpoints how the presence of an agent distorts the firm's inference problem of deciding the state has switched. It is the preferred method for describing how and when the presence of an agent improves operational efficiency.

4.1 Designing the golden parachute

One approach to the optimal contracting problem is to design the golden parachute F, treating the termination policy τ^Y as an auxiliary variable and eliminating it. To do so, one must determine the optimal termination policy which incentivizes truthful reporting under a given golden parachute F.

It turns out that early termination serves a very limited role in an optimal contract. To incentivize truthful reporting, the firm must commit to limit the agent's continuation utility (flow rents plus golden parachute) from delaying his report of a state switch. In particular, if the golden parachute ever falls to zero, then any delayed report must also yield at most zero expected utility. Given the agent's limited liability, immediate project termination is of course the only way to enforce such a promise. Is early termination optimal in any other circumstance? The following lemma answers this question in the negative.

Lemma 6. *Let $\mathcal{C} = (F, \tau^Y)$ be an IC-B contract. Then $\tau^Y \leq \inf\{t \,:\, F_t = 0\}$, and there exists another IC-B contract $\widetilde{\mathcal{C}} = (\widetilde{F}, \widetilde{\tau}^Y)$ such that $\widetilde{F}_t = \min\{F_t, b/\rho\}$ whenever $\tau^Y \geq t$ and $\widetilde{\tau}^Y = \inf\{t \,:\, \widetilde{F}_t = 0\}$. Further, $\widetilde{\tau}^Y \geq \tau^Y$ and $\widetilde{\mathcal{C}}$ yields expected profits at least as high as $\mathcal{C}$.*

The first claim of the contract is just a restatement of the conclusion that no contract can operate past the point where the golden parachute hits zero. The lemma then constructs a new contract illustrating two design principles: first, the golden parachute can be capped at b/ρ; and second, if a contract (F, τ^Y) is ever terminated before F hits zero, it can be extended in an incentive-compatible way without decreasing profits. The first principle is closely related to the fact that $(b/\rho, \infty)$ is an IC contract, as I demonstrated in Section 3.1. Essentially, whenever F_t hits b/ρ, delayed reporting can be deterred indefinitely simply by holding the fee at that level; in particular, it can be delayed until the next time F_t falls below b/ρ. Thus $\min\{F_t, b/\rho\}$ is also incentive compatible. The second principle can be proven by explicitly constructing an IC continuation contract. In particular, set

$$\widetilde{F}_t = \max\left\{b/\rho - (b/\rho - \min\{F_{\tau^Y}, b/\rho\})e^{\rho(t-\tau^Y)}, 0\right\}$$

for $t \geq \tau^Y$. Note that so long as $\widetilde{F}_t > 0$, $\frac{d}{dt}\widetilde{F}_t = -b + \rho\widetilde{F}_t$, which exactly offsets the flow rents and interest expense of delaying termination by a moment. Hence incentive compatibility is ensured for $t \geq \tau^Y$. As for times before τ^Y, $\widetilde{F}_t$ provides no more expected utility to the agent than a contract that terminates at τ^Y and pays $\widetilde{F}_{\tau^Y}$. So it provides no additional incentives to delay reporting prior to τ^Y. Finally, because this extension of the contract pays no more than $\widetilde{F}_{\tau^Y}$ to the agent while generating positive revenue as long as the state is good, it must improve profitability relative to termination at τ^Y.

Lemma 6 solves the design problem for τ^Y given an incentive-compatible golden parachute F. The firm's problem therefore reduces to the design of the dynamics of the golden parachute, subject to the requirement that the project terminate as soon as the parachute reaches zero. I derive the optimal golden parachute in Section 5.

4.2 Designing the termination policy

The second approach to deriving an optimal contract is to design the termination policy τ^Y while treating the golden parachute F as an auxiliary variable. To proceed this way, one must determine the optimal golden parachute which incentivizes truthful reporting under a given termination policy τ^Y.

Consider that one reporting strategy an agent may follow is to never report a state switch, collecting rents until the firm halts the project according to τ^Y. The value of any IC-B golden parachute implementing τ^Y must therefore exceed this amount, implying the inequality

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

after all histories. Setting $F_{\tau^Y} = 0$ and F_t to saturate this inequality minimizes the golden parachute subject to deterring a deviation to never reporting a state change. The following lemma establishes that such a payment scheme deters all other delayed reporting strategies as well and is therefore optimal within the class of IC-B golden parachutes implementing τ^Y.

Lemma 7. *Every $\mathbb{F}^Y$-stopping time τ^Y is supportable as part of an IC-B contract, and the profit-maximizing IC-B golden parachute implementing τ^Y satisfies*

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

for all t and all histories such that $\tau^Y \geq t$.

The surprising conclusion of Lemma 7 is that the firm faces no real double deviation problem when designing an optimal golden parachute. The parachute in the lemma is derived by pretending that the agent can do just one of two things at the time of a state change - report the change now, or stay silent forever. In general, the agent's ability to deviate again by reporting a state change at some later date implies additional constraints that might bind at the optimum. In my model, however, delaying a report of a state change for any interval of time leaves the agent in an "on-path" setting identical to a counterfactual agent with the same output history but a later state switch. And under the contract of Lemma 7, that agent's golden parachute is exactly his expected rents from never reporting. Therefore the true agent's utility from delaying a report by that interval is identical to his utility from never reporting. This argument establishes the lemma's claim that no strategy of delayed reporting yields a profitable deviation.

Lemma 7 allows F to be eliminated from the firm's objective function, reducing the contracting problem to the design of τ^Y alone. I derive the optimal termination policy in Section 6.

4.3 The firm as a uniform-price monopsonist

Lemmas 6 and 7 reveal that the firm's choice of how long to operate the project entails tradeoffs familiar from standard monopoly theory. In particular, consider a uniform-price monopsonist with a concave value function facing an upward-sloping supply function for procurement of an input good. The basic tension of that setting is that purchasing an additional unit on the margin entails paying a higher price for all inframarginal units. The same tradeoff appears in the contracting problem of this paper. In particular, extending the expected lifetime of the project (increasing τ^Y) requires offering a larger golden parachute at all prior times. As the total cost of administering the contract is $\mathbb{E}[e^{-\rho\Lambda}F_\Lambda]$, an additional purchase on the margin therefore increases the cost of all inframarginal units. Meanwhile the firm's marginal valuation curve is downward-sloping, as an extra increment of project lifetime is discounted both by time preference and the probability of a state switch. So the firm can be helpfully viewed as a uniform-price monopsonist trading off output and payments.

In addition to the conceptual analogy, there is a close analogy in solution technique between the contracting and monopoly settings. The uniform-price monopsonist problem can be solved using one of two equivalent approaches. In the first, the supply curve is used to eliminate quantity from the profit function, and the resulting objective function is maximized for price. In the second, the supply curve is inverted and used to eliminate price, resulting in a maximization problem for quantity. I pursue a closely analogous pair of approaches for finding an optimal contract in this setting.

While the analogy to a static monopoly problem is illuminating, it is limited. A distinctive feature of this paper's contracting problem not present in a static setting is that the firm must decide not only how long to run the project on average, but also how variable to make the lifetime. This has significant implications for the value of pursuing multiple solution approaches. In the static problem the economic insights of both avenues are identical, as the optimal price (quantity) may be readily inserted into the supply (inverse supply) function to obtain the remaining variable. By contrast, in the present setting the techniques directly characterize distinct aspects of the optimal contract which are difficult to extract directly from the generalized supply and inverse supply relationships. It is therefore necessary to analyze both to fully understand the behavior of an optimal contract.

5 The price approach

In this section I pursue the approach suggested by Lemma 6 and eliminate τ^Y from the optimal contracting problem, reducing the contracting problem to the design of F alone. I then use techniques of recursive dynamic contracting to maximize the firm's profits wrt incentive-compatible F and derive an optimal golden parachute. The main theorem of this section derives a differential equation, the HJB equation, whose solution characterizes the optimal incentive-compatible golden parachute. The associated optimal termination policy is left implicit, but could in principle be calculated by Lemma 6. Throughout this section, I will often summarize a contract by the golden parachute process F alone; whenever τ^Y is not explicitly stated as an element of the contract, it is understood to be the first time the golden parachute reaches zero.

5.1 The golden parachute as state variable

In this subsection I show that the current value of the golden parachute constitutes a sufficient statistic for the state of the contract. In the context of recursive contractual design, this assertion amounts to two statements. First, the firm's profit function can be written as a discounted sum of flow payments involving only the current value of F at each moment of time. Second, the agent's IC constraints can be reduced to a set of local restrictions on the dynamics of F. With these results in hand, standard techniques of continuous-time dynamic contracting (see, e.g., Sannikov (2008)) can be used to derive the firm's value function using

F as a state variable.

Recall that the firm's ex post profits from an IC-B contract F consist of flow payments r_G up to the stopping time $\tau = \Lambda \wedge \tau^Y$, followed by a lump-sum payment of F_τ once τ is reached. From an ex ante perspective, this payment flow can be analyzed as follows: at each time t, there is a probability $e^{-\alpha t}$ that the project has not yet terminated. In this case the firm receives a discounted flow payment of $e^{-\rho t} r_G \, dt$, and with probability $\alpha \, dt$ the state changes in the next instant and the firm must pay out $e^{-\rho t} F_t$. The net expected contribution of time t's flows to time-zero profits are therefore $e^{-(\rho+\alpha)t}(r_G - \alpha F_t)$. These flows are then summed up to the public stopping time τ^Y and appropriately averaged to obtain the firm's total expected profits from F. This result is formalized in the following lemma.

Lemma 8. *The firm's expected profits under any contract F when the agent reports truthfully are*

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

Lemma 8 establishes the first prong of the claim that F_t appropriately summarizes the state of the contract at time t. Note that the expectation in the profit function of Lemma 8 is wrt to the measure $\mathbb{P}^G$ under which Z^G is a standard Brownian motion. In other words, it assumes a world in which the project is always good. In this formulation, the distortion caused by state switching is entirely captured by the higher discount rate $\rho + \alpha$ applied to flow profits.

As for the second prong, recall that IC-B means stopping immediately following a state switch must provide higher utility to the agent than any delayed reporting policy. In particular, it must pay more, on average, than waiting a moment and then reporting. This latter constraint is purely local, restricting only the instantaneous dynamics of the golden parachute given its current value. In fact, in this model these local constraints collectively imply global IC-B. Intuitively, an agent who deviates from immediate reporting of a state switch to delay momentarily finds himself in a situation identical to an agent with the same public output history and a slightly later state switch who has not (yet) deviated. Therefore if agents have incentives to report a state switch at time $t + dt$, an appropriate local constraint will ensure that agents at time t do as well.

This intuition is difficult to formalize for general golden parachute processes, as I do not rule out the possibility that F evolves very irregularly. However, when F is sufficiently well-behaved incentive compatibility can be characterized quite sharply.

Lemma 9. *Fix a contract* $\mathcal{C} = (F, \tau^Y)$. *Suppose there exist* $\mathbb{F}^Y$*-adapted, right-continuous processes* γ *and* β *such that for each* t, F *satisfies*

$$F_t = F_0 + \int_0^t \gamma_s \, ds + \int_0^t \beta_s \, dZ_s^G$$

a.s. Then (F, τ^Y) *is an IC-B contract iff*

$$b - \rho F_t + \gamma_t - \frac{\Delta r}{\sigma} \beta_t \le 0$$

whenever $t < \tau^Y$.

Lemma 9 decomposes the increments dF_t of the golden parachute into two components: an instantaneous average change $\gamma_t \, dt$, and a "surprise" term $\beta_t \, dZ_t^G$ linked to the arrival of information on project output. In a world where the project's state is known to be good, the surprise term has mean zero and $\gamma_t \, dt$ captures the expected change in the golden parachute in the next instant. The derivation of the incentive constraint deterring late reporting arises from the observation that after the state has switched, Z^B rather than Z^G has zero-mean increments. Further, the increments of the two processes are related via $dZ_t^G = dZ_t^B - \frac{\Delta r}{\sigma} \beta_t \, dt$.

It is then straightforward to compute the expected change in utility of an agent who momentarily delays reporting a state change. After the delay, the agent collects additional flow rents $b \, dt$; pays an interest cost $-\rho F_t \, dt$ from collecting the golden parachute a bit later; and receives the expected change in the value of the golden parachute, which from his perspective is $\left(\gamma_t - \frac{\Delta r}{\sigma} \beta_t\right) dt$. If the sum of these three factors is non-positive everywhere, then the agent never benefits from delaying a report of a state change.

Unfortunately Lemma 9 does not fully characterize the set of IC-B contracts, as it applies only to a particularly well-behaved class of stochastic processes. Indeed, the most involved technical step in the proof of Theorem 1 is establishing the suboptimality of more general golden parachutes.[13] However, Lemma 9 captures the key intuition for the incentive constraint appearing in Theorem 1.

Lemmas 8 and 9 together indicate that the current value of the golden parachute is an appropriate state variable for a recursive formulation of the firm's contracting problem. The next subsection carries out the recursive derivation of an optimal contract.

Remark. *The choice of* F_t *as a state variable for this problem contrasts with standard con-*

[13] Examples of such processes include Ito processes with drift or diffusion terms that are not right-continuous; semimartingales whose BV components are not absolutely continuous or which contain jumps; as well as even more irregular processes which are not semimartingales, such as fractional Brownian motions.

tracting environments, which generally use the agent's continuation utility as a state variable. One distinctive feature of my model is the presence of asymmetric information, which in general requires tracking a pair of continuation utilities for the agent in each state. Letting (U_t^G, U_t^B) *be the agent's continuation utilities in each state at time* t*, these variables can be computed explicitly for any IC contract satisfying backloading and no late termination:*

$$U_t^G = \mathbb{E}_t^G \left[\int_t^{\tau^Y} e^{-(\rho+\alpha)(s-t)}(b + \alpha F_s)\, ds + e^{-(\rho+\alpha)(\tau^Y - t)} F_{\tau^Y} \right], \quad U_t^B = F_t.$$

The choice of F_t *as a state variable is therefore equivalent, after appropriate reduction of the contract space, to choosing* U_t^B *as a state variable. The key property of my model which further reduces the state space is the fact that IC constraints in the good state don't bind at the optimum. The analysis therefore amounts to designing the dynamics of* U_t^B*, i.e. the terminal utility of the contract.*

5.2 A recursive formulation of the optimal contract

Consider the class of all IC-B contracts promising a golden parachute of $F_0 \geq 0$ if the agent reports a state switch at time zero, and let $V(F_0)$ be the maximal profit achievable by the firm among all such contracts. V is then the firm's *value function.* To calculate V, I utilize the accounting identity that $V(F_0)$ should be equal to the firm's instantaneous flow profits at time 0 plus continuation profits from following an optimal contract a moment later. Of course, what state is reached a moment later depends on the contract increment dF chosen at time 0; as V is the value function dF must of course be chosen to maximize lifetime profits over all incentive-compatible contract increments. Heuristically, V must satisfy the "Bellman equation"

$$V(F_0) = \max_{dF \in \mathrm{IC}(F_0)} \mathbb{E}^G \left[(r_G - \alpha F_0)\, dt + e^{-(\rho+\alpha)dt} V(F_0 + dF) \right].$$

Supposing the firm chooses $dF = \gamma\, dt + \beta\, dZ_t^G$, its control variables are (γ, β), which by Lemma 9 must satisfy $b - \rho F_0 + \gamma - \frac{\Delta r}{\sigma}\beta \leq 0$. Derivation of a proper HJB equation (the continuous-time analog of a Bellman equation) is then accomplished by taking $dt \to 0$ appropriately.

Ito's lemma provides the proper machinery to take this limit. It says that, to first order,

$$V(F_0 + dF) \simeq V(F_0) + \left(\gamma V'(F_0) + \frac{1}{2}\beta^2 V''(F_0) \right) dt + \beta V'(F_0)\, dZ_t^G.$$

This expression differs from a standard Taylor expansion due to the presence of the second-order term $\frac{1}{2}\beta^2 V''(F_0)$. Intuitively, the extra term captures the impact of the curvature of V on $\mathbb{E}^G[V(F_0 + dF)]$, which by Jensen's inequality differs from $V(F_0 + \mathbb{E}^G[dF])$ in the same direction as the sign of V''.

Inserting the Ito expansion of $V(F_0 + dF)$ into the "Bellman equation", the final term $\beta V'(F_0)\, dZ_t^G$ has mean zero and vanishes.[14] Expanding $e^{-(\rho+\alpha)dt}$ to first order as $1-(\rho+\alpha)dt$, eliminating second-order terms, and re-arranging yields the HJB equation

$$(\rho + \alpha)V(F_0) = \max_{(\gamma,\beta)\in \mathrm{IC}(F_0)} \left\{ r_G - \alpha F_0 + \gamma V'(F_0) + \frac{1}{2}\beta^2 V''(F_0) \right\}.$$

This is a second-order differential equation which, combined with appropriate boundary conditions, pins down the firm's value function. One boundary condition is easy: when $F_0 = 0$ the unique IC contract entails immediate termination, so $V(0) = 0$. The second condition is more subtle, and turns out to be a so-called "free boundary condition". Note that an upper bound on the value of any contract to the firm is $r_G/(\rho + \alpha)$, as this is what the firm would earn if it could directly observe the state of the project. Conversely, as discussed earlier there is no benefit to offering a golden parachute larger than b/ρ. Therefore V is bounded above, and it should reach a maximum for some finite value of the golden parachute. The second boundary condition is then that $V'(F^*) = 0$ for some undetermined value F^* of the golden parachute.

The following theorem provides a formal statement of the conditions under which the heuristic approach just outlined produces a solution to the firm's contracting problem. It constitutes the main technical result of this section.

Theorem 1. *Suppose there exists a constant $F^* > 0$ and a C^2 function $V : [0, F^*] \to \mathbb{R}$ satisfying $V(0) = 0$, $V'(F^*) = 0$, and the* HJB equation

$$(\rho + \alpha)V(F_0) = \sup_{(\gamma,\beta)\in IC(F_0)} \left\{ r_G - \alpha F_0 + \gamma V'(F_0) + \frac{1}{2}\beta^2 V''(F_0) \right\}$$

for every $F_0 \in [0, F^*]$, where $IC(F_0) \equiv \{(\gamma, \beta) \ : \ b - \rho F_0 + \gamma - \beta \frac{\Delta r}{\sigma} \leq 0\}$. Then:

- *$V(F^*)$ is an upper bound on the expected profits of any IC-B contract;*
- *$F^* < b/\rho$, V is strictly increasing and strictly concave on $[0, F^*]$, and $V(F^*) = \frac{r_G - \alpha F^*}{\rho+\alpha}$;*

[14] Strictly speaking, stochastic integrals wrt Z^G are martingales only if they satisfy appropriate regularity conditions. Checking that these conditions are met is a standard step in a rigorous verification proof. See the proof of Theorem 1 for details.

- *There exist unique continuous functions* $\gamma^*, \beta^* : \mathbb{R} \to \mathbb{R}$ *such that* $(\gamma^*(F_0), \beta^*(F_0))$ *maximize the HJB equation on* $[0, F^*]$ *and* γ^*, β^* *are constant on* $\mathbb{R} \setminus [0, F^*]$*;*
- $\beta^* \geq 0$, $\beta^*(F^*) = 0$, *and* $\gamma^*(F^*) < 0$*;*
- *There exists a weak solution* F *to the stochastic differential equation (SDE)*

$$F_t = F^* + \int_0^t \gamma^*(F_s)\, ds + \int_0^t \beta^*(F_s)\, dZ_s^G \tag{3}$$

  for all time;
- *If* $\lim_{F_0 \to F} V'(F_0)V'''(F_0)$ *exists and is finite, then there exists a unique strong solution* F *to* (3) *for all time;*
- *If* F *is a strong solution to* (3)*, then* $\max\{F, 0\}$ *is an optimal contract whose expected profits are* $V(F^*)$*.*

Theorem 1 states that solving a particular boundary value problem yields the firm's value function as well as an optimal contract. Importantly, the value function is actually *increasing* over a range of golden parachutes. This is because when the parachute reaches zero, the firm is forced to terminate the project and loses all future output from its operation. Thus for sufficiently small values of F_0, the revenue from a longer project lifespan outweighs the cost of increasing expected payments to the agent, leading to an upward-sloping value function. The critical point F^* determines when these two factors exactly balance. Because $\gamma^*(F^*) < 0$ and $\beta^*(F^*) = 0$, an optimal contract never promises a golden parachute larger than F^* no matter how good the history of output. Instead, as F_t approaches F^* the volatility of the parachute vanishes and the parachute drifts downward deterministically, holding its value below F^*.

The HJB equation also allows one to explore the economics of the firm's contractual design problem. Given concavity of V as assured by the theorem, the optimal controls $(\gamma^*(F_0), \beta^*(F_0))$ are readily computed as the first-order conditions of the HJB equation given a binding IC constraint:

$$\beta^*(F_0) = -\frac{\Delta r}{\sigma} \frac{V'(F_0)}{V''(F_0)}, \quad \gamma^*(F_0) = -(b - \rho F_0) + \frac{\Delta r}{\sigma} \beta^*(F_0).$$

Intuitively, the tradeoff the firm faces when designing a golden parachute is whether to provision incentives through downward drift or volatility of the parachute. Both are costly for the firm, because each brings the contract closer to premature termination on average.

An optimal contract therefore provisions incentives as lightly as possible, leading to a binding IC-B constraint at all points along the contract. In other words, $(\gamma^*(F_0), \beta^*(F_0))$ satisfy

$$b - \rho F_0 + \gamma^*(F_0) - \frac{\Delta r}{\sigma}\beta^*(F_0) = 0.$$

The optimal mix of incentives is then determined by the slope and curvature of the value function: each unit of downward drift added to the contract incurs a constant marginal cost $V'(F_0)$, while each unit of volatility added incurs a linearly increasing marginal cost $-\beta V''(F_0)$. Equating the ratio of marginal costs to the marginal rate of substitution $\frac{\Delta r}{\sigma}$ between volatility and drift in the incentive constraint determines the optimal mix of incentives.

The fact that an optimal contract provisions incentives as lightly as possible and increases the golden parachute in response to good news verifies an earlier conjecture that the IC constraint for the agent in the good state is slack under an optimal contract. To see this, consider any contract evolving as

$$dF_t = \gamma_t\, dt + \beta_t\, dZ_t^G,$$

where $b - \rho F_t + \gamma_t - \frac{\Delta r}{\sigma}\beta_t = 0$ and $\beta \geq 0$. Reasoning as in the paragraphs following Lemma 9, one may compute the expected change in the agent's utility when the state is good and he considers momentarily delaying a (false) report of a state change, which turns out to be $b - \rho F_t + \gamma_t \geq 0$. Thus it is always at least weakly optimal for an agent in the good state to delay (falsely) reporting a state change a moment. Since delay continues to be optimal as long as the state as good, global incentive compatibility obtains.

Finally, note that Theorem 1 falls short of completely solving the optimal contracting problem, as it does not ensure that a solution to the HJB equation actually exists. The problem is essentially technical, as in principle the firm's value function may not have the required degree of smoothness to satisfy a standard HJB equation. I defer a thorough examination of the technical issues to Section 5.4. In the next subsection, I take a different approach and exhibit an explicit solution to the HJB equation under a particular parameter restriction. For this class of models, Theorem 1 ensures that such a solution is in fact the firm's value function.

5.3 A complete solution to a special case

For a particular subspace of model parameters, the boundary value problem posed in Theorem 1 has an elegant closed-form solution. The following proposition states the parameter restriction and describes the form of the value function and the associated optimal controls.

Proposition 2. *Suppose $\rho = \alpha + \left(\frac{\Delta r}{\sigma}\right)^2$. Then there exist unique constants a_1 and a_2 such that $a_1 > 0 > a_2$ and (a_1, a_2) satisfy*

$$\begin{cases} \rho a_1 + 2ba_2 + \alpha = 0, \\ \frac{1}{4}\left(\frac{\Delta r}{\sigma}\right)^2 a_1^2 + ba_1a_2 - r_G a_2 = 0. \end{cases}$$

Further, $V(F_0) = a_2F_0^2 + a_1F_0$ satisfies the HJB equation of Theorem 1 on $[0, F^*]$ and the boundary conditions $V(0) = 0$, $V'(F^*) = 0$, and $V''(F^*) < 0$, where

$$F^* \equiv \left(\frac{\Delta r}{\sigma}\right)^{-2}\left(b + r_G\frac{\rho}{\alpha} - \sqrt{(b + r_G)^2 + r_G^2\left(\left(\frac{\rho}{\alpha}\right)^2 - 1\right)}\right) \in (0, b/\rho).$$

The corresponding continuous maximizers of the HJB equation are

$$\beta^*(F_0) = \frac{\Delta r}{\sigma}(F^* - F_0)$$

and

$$\gamma^*(F_0) = -\alpha(F^* - F_0) - (b - \rho F^*).$$

Further, $V''' = 0$, so there exists a unique strong solution to (3).

Proposition 2 solves the firm's optimal contracting problem when $\rho = \alpha + \left(\frac{\Delta r}{\sigma}\right)^2$. In this case the firm's value function is quadratic and the associated controls are linear in the current value of the golden parachute. One implication of this linearity is that the optimal contract's sensitivity to news increases the closer the contract moves to termination, and conversely vanishes as $F_0 \to F^*$. This dynamic conforms to a sensible intuition that news should "matter more" for incentives as the contract progresses and it becomes ex ante more likely that the state has transitioned. In Section 6 I explore the issue of optimal sensitivity to news from another angle, when considering how news affects progress toward project termination.

Proposition 2 characterizes all the essential features of the firm's optimal contract but stops short of deriving the contract explicitly. Given the linearity of the optimal controls, the optimal contract satisfies the linear SDE

$$dF_t = -(\alpha(F^* - F) + (b - \rho F^*))\, dt + \frac{\Delta r}{\sigma}(F^* - F)\, dZ_t^G.$$

Such SDEs are solvable in closed form, yielding an explicit solution for the optimal contract:

Corollary. *Suppose* $\rho = \alpha + \left(\frac{\Delta r}{\sigma}\right)^2$. *Then letting* F^* *be as in Proposition 2, an optimal contract is given by*

$$F_t = F^* - (b - \rho F^*) \exp\left(\left(\frac{3}{2}\alpha - \frac{1}{2}\rho\right) t - \frac{\Delta r}{\sigma} Z_t^G\right) \int_0^t \exp\left(-\left(\frac{3}{2}\alpha - \frac{1}{2}\rho\right) s + \frac{\Delta r}{\sigma} Z_s^G\right) ds$$

for $t \leq \inf\{t \ : \ F_t = 0\}$.

An immediate implication of this corollary is the following: fix a time t and two states of the world $\omega_1, \omega_2 \in \Omega$ satisfying $Y_s(\omega_1) > Y_s(\omega_2)$ for all $s \in (0, t)$ and $Y_t(\omega_1) = Y_t(\omega_2)$. Then $F_t(\omega_1) < F_t(\omega_2)$. In words, ω_1 and ω_2 encode two different paths of output that reach the same cumulative output by time t. However, output under ω_1 outperforms early and lags late compared to ω_2; consequently the contract is brought closer to termination by time t under ω_1. Thus in a very strong sense the optimal contract weights news more heavily later in the contract.

5.4 Existence of solutions to the HJB equation

This subsection is devoted to discussing technical existence issues of solutions to the HJB equation derived in Theorem 1. It can be safely omitted by readers uninterested in this detail.

Theorem 1 establishes a program for characterizing an optimal contract by solving a particular boundary value problem. However, it does not ensure that this program will succeed, as it could be that no solution to the boundary value problem exists. A standard solution to this problem is to apply ODE existence theorems directly to the HJB equation to ensure existence of solutions to a family of initial value problems, and then to argue that there exist initial conditions whose solutions satisfy the right boundary conditions. The following proposition demonstrates the difficulties of that approach in this context.

Proposition 3. *Let* $V : [0, F^*] \to \mathbb{R}$ *be a* C^2 *function satisfying* $V(0) = 0$ *and* $V'(F^*) = 0$. *Then* V *satisfies the HJB equation on* $[0, F^*]$ *iff it is strictly increasing and strictly concave* *on* $[0, F^*)$ *and satisfies*

$$V''(F_0) = -\frac{1}{2}\left(\frac{\Delta r}{\sigma}\right)^2 \frac{V'(F_0)^2}{(\rho+\alpha)V(F_0) - (r_G - \alpha F_0) + (b - \rho F_0)V'(F_0)}$$

for each $F_0 \in [0, F^*)$.

Further, suppose V *is such a function. Then* $V(F^*) = \frac{r_G - \alpha F^*}{\rho+\alpha}$, $F^* < b/\rho$, *and* $V''(F^*) < 0$.

Proposition 3 demonstrates that finding solutions to the HJB equation boils down to solving a particular nonlinear second-order ODE of the form $V'' = G(F_0, V, V')$. Morever, the solutions of interest to the optimal contracting problem are *singular* solutions to this ODE. To see this, consider the numerator and denominator of G as $F_0 \to F^*$. The numerator becomes $-\frac{1}{2}\left(\frac{\Delta r}{\sigma}\right)^2 V'(F^*)^2$, which vanishes by definition of the free boundary condition. As for the denominator, $(\rho+\alpha)V(F^*) - (r_G - \alpha F^*) + (b - \rho F_0)V'(F^*) = 0$ given the proposition's final result that $V(F^*) = \frac{r_G - \alpha F^*}{\rho+\alpha}$. This condition follows from the observation that the constant function $V(F_0) = V(F^*)$ is a local solution to the ODE whenever $V(F^*) \neq \frac{r_G - \alpha F^*}{\rho+\alpha}$. Given that G is continuously differentiable and thus uniformly Lipschitz on sufficiently small neighborhoods around $(F^*, V(F^*), 0)$, standard local ODE uniqueness results then imply that there can be no strictly increasing solutions to the ODE near F^*. Therefore candidate solutions of the ODE must satisfy $V(F^*) = \frac{r_G - \alpha F^*}{\rho+\alpha}$.

For this approach to succeed, then, it is necessary at minimum to show local existence of a solution to the ODE $V'' = G(F_0, V, V')$ around the singularity $(F_0, V(F_0), V'(F_0)) = (F^*, (r_G - \alpha F^*)/(\rho + \alpha), 0)$ for some range of $F^* > 0$. As far as I am aware no general existence results apply in this setting. I return to this question in Section 6, assisted by the full characterization of an optimal contract obtainable by the quantity approach.

6 The quantity approach

In this section I pursue the dual approach to Section 5, eliminating F from the firm's problem and optimizing the resulting objective function to derive an optimal termination policy. The main result of this section shows that the optimal incentive-compatible termination policy solves a particular quickest detection problem, closely analogous to the firm's problem without an agent. I then develop an optimal stopping approach to explicitly solve the quickest detection problem.

6.1 The firm's objective function

Recall that by Lemma 7, any $\mathbb{F}^Y$-stopping time τ^Y is implementable by an IC-B contract, with associated profit-maximizing golden parachute

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

The following proposition proves 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 4. *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. Then*

$$\Pi(\tau^Y) = \mathbb{E} \left[\int_0^{\tau^Y} e^{-\rho t} (\pi_t r_G - (1 - \pi_t) b) \, dt \right]. \tag{4}$$

Though the formal derivation of (4) is somewhat involved, the intuition behind it is very simple. Consider any contract with termination policy τ^Y and golden parachute as specified in Lemma 7. Upon a reported state change the firm pays the agent his expected discounted flow of rents, conditional on the state being bad, from allowing the project to operate until τ^Y. The firm's expected profits under such a contract are therefore identical to a counterfactual setting in which its project pays expected returns $-b$ rather than r_B when the state is bad. And conditional on the public information available at time t, the probability of being in the good state at that time is π_t, so instantaneous expected returns at any time are $\pi_t r_G - (1 - \pi_t) b$.

Proposition 4 reduces the contracting problem with an agent to solving a particular virtual quickest detection problem, closely related to the firm's problem without an agent (cf. equation (1)). In this virtual problem the firm learns about the state of the project as if no agent were available but incurs a flow cost of operating the project in the bad state of $-b$ rather than r_B. Importantly, this reduction does *not* imply that the optimal contract implements the same termination policy as in a quickest detection problem with no agent and r_B replaced by $-b$. First, the firm's termination policy is $\tau = \Lambda \wedge \tau^Y$, so the presence of an agent prevents operation of the project in the bad state; this is in stark contrast to the problem without an agent, where both early and late terminations occur. Second, the firm learns about the current state more quickly in the virtual problem, as its signal-to-noise ratio (SNR) is $\frac{\Delta r}{\sigma}$ with an agent, versus $\frac{r_G + b}{\sigma}$ in the quickest detection problem with r_B replaced by b.

6.2 Threshold termination policies

Given the form of $\Pi(\tau^Y)$ in Proposition 4 and the fact that posterior beliefs drift downward over time, a natural conjecture for an optimal termination policy is a threshold rule in the firm's posterior beliefs, i.e. $\tau^{\underline{\pi}} \equiv \inf\{t : \pi_t \leq \underline{\pi}\}$ is an optimal policy for some $\underline{\pi} \in [0,1]$. To verify this conjecture, it is necessary to establish that π_t is a sufficient statistic for the history of the project when projecting the future path of beliefs; in other words, one must rule out the possibility that the time-t conditional distribution of future beliefs depends on $\{Y_s\}_{s\leq t}$ in a more complicated way than through π_t. The following lemma accomplishes this task by fully characterizing the dynamics of π :

Lemma 10. π *satisfies the SDE*

$$d\pi_t = -\alpha\pi_t\,dt + \frac{\Delta r}{\sigma}\pi_t(1-\pi_t)d\overline{Z}_t$$

with initial condition $\pi_0 = 1$, *where* $\overline{Z}$ *is a* $\mathbb{P}$*-standard Brownian motion adapted to* $\mathbb{F}^Y$ *with increments*

$$d\overline{Z}_t = \frac{1}{\sigma}(dY_t - (\pi_t r_G + (1-\pi_t)r_B)\,dt).$$

The SDE describing the evolution of π_t is derivable from Bayes' rule, and consists of two contributions. First, the known transition rate α from the good to the bad state leads to a deterministic drop in beliefs over time at rate $\alpha\pi_t$. Second, the firm learns about the state via unexpected deviations in observed incremental output. The quadratic dependence of this learning on π is related to the fact that beliefs are most dispersed at $\pi = 0.5$, and hence most subject to revision as new information is received; more extreme beliefs require stronger evidence to change.

The process $\overline{Z}$ is known as the *innovation process*, and tracks deviations of the project's output from time-t expectations. The fact that $\overline{Z}$ is a standard Brownian motion implies that the distribution of the future path of beliefs conditional on $\mathcal{F}_t^Y$ depends on the public history only through π_t. Thus π_t is indeed a sufficient statistic for the history of the project.

This insight reduces the firm's contractual design problem to the choice of a single number $\underline{\pi} \in [0,1]$. A natural baseline contract is the break-even termination policy, under which $\underline{\pi}$ satisfies $\underline{\pi} r_G - (1-\underline{\pi})b = 0$. This termination policy would be optimal if, at the breakeven point, the firm had to either end the project immediately or commit to continuing it forever. However, the firm retains a real option to terminate the project in the future, which increases the value of continuing past the break-even point in the hopes of receiving good news about the project's state which raises flow returns. An optimal threshold policy should therefore satisfy $\underline{\pi} < b/(b + r_G)$. The next subsection determines the optimal threshold by calculating the value of the firm's real option over time.

6.3 Solving the quickest detection problem

Let

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

be the value of the firm's real option from continuing the project at time t. Given that π_t is a sufficient statistic for the project history, this option value can be written $R_t = \widetilde{V}(\pi_t)$ for some function $\widetilde{V} : [0,1] \to \mathbb{R}_+$, which I refer to as the firm's *virtual value function.* It is natural to expect that $\widetilde{V}$ is a weakly increasing function which vanishes for sufficiently small posterior beliefs, at which point the firm's real option is worthless and an optimal contract is immediately terminated. An optimal threshold is then the largest $\underline{\pi}$ such that $\widetilde{V}(\underline{\pi}) = 0$.

The following theorem derives a differential equation characterizing $\widetilde{V}$ and establishes that the procedure just outlined maximizes $\Pi(\tau^Y)$.

Theorem 2. *Suppose there exists a $\underline{\pi} \in (0,1)$ and a C^2 function $\widetilde{V} : [\underline{\pi}, 1] \to \mathbb{R}$ satisfying*

$$\rho\widetilde{V}(\pi_0) = \pi_0 r_G - (1-\pi_0)b - \alpha\pi_0\widetilde{V}'(\pi_0) + \frac{1}{2}\left(\frac{\Delta r}{\sigma}\right)^2 \pi_0^2(1-\pi_0)^2\widetilde{V}''(\pi_0)$$

on $[\underline{\pi}, 1]$ and boundary conditions $\widetilde{V}(\underline{\pi}) = \widetilde{V}'(\underline{\pi}) = 0$. Then $\underline{\pi} \leq b/(b + r_G)$, $\widetilde{V}(\pi_0) > 0$ for $\pi_0 \in (\underline{\pi}, 1]$, and $\widetilde{V}$ may be extended to a C^1, piecewise C^2 function on $[0,1]$ by setting $\widetilde{V}(\pi_0) = 0$ for $\pi_0 \in [0, \underline{\pi})$. On this extended domain, $\widetilde{V}$ is the firm's virtual value function.

Further, let $\tau^* \equiv \inf\{t \ : \ \pi_t \leq \underline{\pi}\}$. Then τ^* maximizes $\Pi(\tau^Y)$ among all $\mathbb{F}^Y$-stopping times τ^Y, and $\Pi(\tau^*) = \widetilde{V}(1)$.

The differential equation stated in Theorem 2 may be derived heuristically in a fashion analogous to the method outlined in Section 5.2 for obtaining the HJB equation of Theorem 1. The boundary condition $\widetilde{V}(\underline{\pi}) = 0$ reflects the fact that at $\underline{\pi}$ the project is terminated. It turns out that, given any choice of $\underline{\pi} \in (0,1)$, the boundary condition $\widetilde{V}(\underline{\pi}) = 0$ along with the requirement that $\widetilde{V}$ be well-behaved (in particular, finite) at $\pi_0 = 1$ select a unique solution $\widetilde{V}$ to the ODE in Theorem 2.[15] Further, $\widetilde{V}(\pi_t)$ gives the time$-t$ continuation profits for the termination policy $\tau^{\underline{\pi}} = \inf\{t \ : \ \pi_t \leq \underline{\pi}\}$.

[15]Technically speaking, the requirement that $\widetilde{V}$ be C^2 at the right boundary is a transversality condition

The final boundary condition, $\widetilde{V}'(\underline{\pi}) = 0$, is a so-called smooth pasting condition pinning down the optimal $\underline{\pi}$. If $\widetilde{V}'(\underline{\pi}) < 0$, then $\underline{\pi}$ is too low, because continuation profits eventually become negative prior to termination. The firm could increase profits by terminating sooner in the region of beliefs where continuation profits are strictly negative. On the other hand, if $\widetilde{V}'(\underline{\pi}) > 0$, then the threshold is too high. The reasoning here is more subtle due to the kink in the value function at $\underline{\pi}$, but it can still be understood heuristically. Consider a firm at the threshold belief who instead of terminating immediately continues an instant dt further and then reverts to terminating the next time beliefs fall below $\underline{\pi}$. The firm's expected gains from such a deviation are, to first order,

$$\Delta\Pi \simeq (\underline{\pi} r_G - (1-\underline{\pi})b)\ dt + \mathbb{E}\left[\widetilde{V}\left(\underline{\pi} - \alpha\underline{\pi}\, dt + \frac{\Delta r}{\sigma}\underline{\pi}(1-\underline{\pi})d\overline{Z}_t\right)\right].$$

Now, assume as a simplification that $d\overline{Z}_t$ is a binary random variable taking values $\pm\sqrt{dt}$ with equal probability (i.e. having the same mean and variance as a Brownian increment over the same time interval). Then for sufficiently small dt

$$\mathbb{E}\left[\widetilde{V}\left(\underline{\pi} - \alpha\underline{\pi}\, dt + \frac{\Delta r}{\sigma}\underline{\pi}(1-\underline{\pi})d\overline{Z}_t\right)\right] = \frac{1}{2}\widetilde{V}\left(\underline{\pi} - \alpha\underline{\pi}\, dt + \frac{\Delta r}{\sigma}\underline{\pi}(1-\underline{\pi})\sqrt{dt}\right),$$

or to first order in $\sqrt{dt}$,

$$\mathbb{E}\left[\widetilde{V}\left(\underline{\pi} - \alpha\underline{\pi}\, dt + \frac{\Delta r}{\sigma}\underline{\pi}(1-\underline{\pi})d\overline{Z}_t\right)\right] \simeq \frac{1}{2}\widetilde{V}'(\underline{\pi}^+)\frac{\Delta r}{\sigma}\underline{\pi}(1-\underline{\pi})\sqrt{dt}.$$

Thus any first-order gains or losses from the flow of profits over the (sufficiently short) interval dt are dwarfed by the lower-order expected gain in continuation profits. $\widetilde{V}'(\underline{\pi}) = 0$ is therefore a necessary condition for optimality of the threshold $\underline{\pi}$. Theorem 2 ensures that smooth pasting is also sufficient for optimality.

The next lemma establishes that the function $\widetilde{V}$ stipulated in Theorem 2 actually exists by explicitly exhibiting one. The lemma makes use of Tricomi's confluent hypergeometric function $U(m, n, z)$, which is a solution to Kummer's differential equation defined and may be written in closed form as

$$U(m,n,z) = \frac{1}{\Gamma(m)}\int_0^\infty e^{-zt}t^{m-1}(1+t)^{n-m-1}\, dt$$

closely analogous to the growth conditions common in macroeconomics and growth theory. It is imposed for the same reason that the growth condition is required to rule out spurious explosive solutions and select a unique growth path for the economy in those models.

when $m, z > 0$.

Lemma 11. *Let $k \equiv \frac{\Delta r}{\sigma}$ 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 $\underline{\pi} \in (0,1)$ such that*

$$\widetilde{V}(\pi_0) = \frac{r_G + b}{\rho + \alpha}\pi_0 - \frac{b}{\rho} + C\pi_0^{\beta}(1-\pi_0)^{1-\beta} U\left(\beta - 1, 2\beta - \frac{2\alpha}{k^2}, \frac{2\alpha}{k^2}\frac{\pi_0}{1-\pi_0}\right)$$

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

$$\rho\widetilde{V}(\pi_0) = \pi_0 r_G - (1-\pi_0)b - \alpha\pi_0\widetilde{V}'(\pi_0) + \frac{1}{2}\left(\frac{\Delta r}{\sigma}\right)^2 \pi_0^2(1-\pi_0)^2\widetilde{V}''(\pi_0)$$

and $\widetilde{V}(\underline{\pi}) = \widetilde{V}'(\underline{\pi}) = 0$. Further, $\widetilde{V}$ may be extended to a C^2 function on $[\underline{\pi}, 1]$ satisfying the ODE on the extended domain.

This lemma relies on the fortuitous fact that the ODE of Theorem 2 is solvable in closed form in terms of known special functions. As the ODE is second-order, the general solution contains two linearly independent components; however, one of them diverges as $\pi_0 \to 1$, and so must be discarded. Most of the work of the lemma then goes toward establishing existence of constants C and $\underline{\pi}$ satisfying the remaining boundary conditions.

The final theorem of this section ties the previous results together and verifies that the stopping time maximizing $\Pi(\tau^Y)$ is actually incentive-compatible.

Theorem 3. *There exists a $\underline{\pi} \in (0, b/(b + r_G)]$ such that $\tau^* \equiv \inf\{t \ : \ \pi_t \leq \underline{\pi}\}$ maximizes*

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

among all $\mathbb{F}^Y$-stopping times τ^Y. Letting F^* be the associated golden parachute defined in Lemma 7, (F^*, τ^*) is an optimal contract.

To see why (F^*, τ^*) is incentive compatible, note first that by the martingale representation theorem $F_t^* = \mathbb{E}_t^B\left[\int_t^{\tau^Y} e^{-\rho(s-t)} b\, ds\right]$ evolves as

$$dF_t^* = (\rho F_t^* - b)\, dt + \beta_t\, dZ_t^B = \left(\rho F_t^* - b + \frac{\Delta r}{\sigma}\beta_t\right) dt + \beta_t\, dZ_t^G$$

for some $\mathbb{F}^Y$-adapted, progressively measurable process β. Further, given that τ^* is a threshold policy in π_t, time to termination and the size of the golden parachute are increasing in π_t.

The arrival of good news about the state must therefore increase F^*, in which case $\beta \geq 0$. This contract is then a member of a class shown to be incentive compatible in the discussion following Theorem 1.

6.4 More on existence of solutions to the HJB equation

The results of Section 6.3 prove affirmatively that the firm's virtual value function $\widetilde{V}$ is a C^2 function. It is natural to wonder whether this result can be leveraged to establish that the firm's value function V in the dual problem satisfies the HJB equation derived in Theorem 1. This subsection explores that possibility. As with Section 5.4, this subsection can be safely skipped by a reader uninterested in technical existence details.

Let (F^*, τ^*) be defined as in Theorem 3. Because π_t is a sufficient statistic for the state of the project, there exists a function $f : [\underline{\pi}, 1] \to \mathbb{R}_+$ such that $F_t^* = f(\pi_t)$ for all time a.s. Clearly for all $x \in [\underline{\pi}, 1]$, $f(x)$ satisfies

$$f(x) = \mathbb{E}^B \left[\int_0^{\tau_x^*} e^{-\rho t} b \, dt \right],$$

where τ_x^* is an $\mathbb{F}^Y$-stopping time defined by $\tau^*(x) \equiv \inf\{t \ : \ p_t^x \leq \underline{\pi}\}$ for the stochastic process p^x satisfying $p_0^x = x$ and

$$dp_t^x = -\alpha p_t^x \, dt + \frac{\Delta r}{\sigma} p_t^x (1 - p_t^x) \frac{1}{\sigma} (dY_t - (p_t^x r_G + (1 - p_t^x) r_B) \, dt).$$

Informally, τ_x^* gives the total remaining project lifetime when the current posterior belief is $x \in [0, 1]$.

Lemma 12. *f is a continuous, strictly increasing function on $[\underline{\pi}, 1]$ satisfying $f(\underline{\pi}) = 0$ and $f(1) = F_0^*$.*

Similarly, there exists a function $V : [0, F_0^*] \to \mathbb{R}$ such that

$$V(F_t^*) = \mathbb{E}_t^G \left[\int_t^{\tau^*} e^{-(\rho+\alpha)s} (r_G - \alpha F_t^*) \, ds \right]$$

for all time a.s. Using f it is straightforward to relate V, the firm's value function when designing F, and $\widetilde{V}$, the firm's virtual value function when designing τ^Y:

Lemma 13.

$$\widetilde{V}(x) = xV(f(x)) - (1 - x) f(x) \tag{5}$$

for all $x \in [\underline{\pi}, 1]$.

Essentially, when the firm's posterior belief is x, then either the state is good; in which case the continuation value of the contract is $V(f(x))$; or the state is bad, in which case the continuation value of the contract is the termination fee paid to the agent at that point.

Given the tight linkage between V and $\widetilde{V}$ provided by Lemma 5 and the smoothness of $\widetilde{V}$, it seems reasonable that V should also be smooth enough to satisfy the HJB equation. The main stumbling block to this argument is establishing that the function f mapping beliefs onto golden parachutes is sufficiently smooth. The following proposition represents a closest approach to showing existence of solutions to the HJB equation of Theorem 1 via this method.

Proposition 5. *Suppose there exists a* C^2 *function* $g : [\underline{\pi}, 1] \to \mathbb{R}$ *satisfying*

$$\rho g(x) = b - \left(\alpha x + \left(\frac{\Delta r}{\sigma}\right)^2 x^2(1-x)\right) g'(x) + \frac{1}{2}\left(\frac{\Delta r}{\sigma}\right)^2 x^2(1-x)^2 g''(x)$$

on $[\underline{\pi}, 1]$ *and the boundary condition* $g(\underline{\pi}) = 0$. *Then* $g = f$, *and* V *is a* C^2 *function satisfying the HJB equation on* $[0, F_0^*]$ *and the boundary conditions* $V(0) = 0$ *and* $V'(F_0^*) = 0$.

Standard ODE existence results guarantee a unique C^2 solution to the ODE of Proposition 5 on $[\underline{\pi}, 1)$ for any choice of $g'(\underline{\pi})$. However, the ODE is singular at $x = 1$ due to the vanishing coefficient on g'' there, and so it is not guaranteed that any solution is well-behaved, in particular finite, as $x \to 1$. Thus despite the explicit characterization of an optimal contract afforded by the quantity approach, some additional regularity assumptions are needed to ensure that V satisfies the HJB equation.

7 The relationship to static mechanism design

The technical developments in this paper have so far been framed in terms of dual approaches to solving a monopsonist's procurement problem. A helpful alternative framing directly connects the techniques here to the canonical solution of a static mechanism design problem. Consider a standard procurement contracting setting: an agent with private cost parameter θ supplying allocation x and receiving transfer T obtains utility $U(x, T; \theta) = T - C(x, \theta)$. The canonical solution, the "first-order approach", derives an inverse supply curve by applying an envelope theorem argument:[16] any IC revelation contract $(x(\cdot), T(\cdot))$ must satisfy the

[16]To amplify the connection with the current paper this argument is developed for the special case of differentiable contracts, though the result holds much more generally.

FOCs $T'(\theta) - C_x(x(\theta),\theta)x'(\theta) = 0$ reflecting the local optimality of truthtelling. Then by the chain rule $U'(\theta) = -C_\theta(x(\theta),\theta)$, where $U(\theta) = U(x(\theta), t(\theta); \theta)$. An inverse supply curve is obtained by integrating the FOCs beginning at the extremal type $\overline{\theta}$ of highest cost:

$$U(\theta) = U(\overline{\theta}) + \int_\theta^{\overline{\theta}} C_\theta(x(\nu),\nu)\, d\nu.$$

This expression pins down $T(\cdot)$ given $x(\cdot)$, up to a surface term $U(\overline{\theta})$ which is determined by the participation constraint $U(\overline{\theta}) \geq 0$. The optimal procurement contract is then derived by eliminating T from the principal's objective function and optimizing wrt x.

A closely analogous heuristic argument can be traced in the setting of my model. My goal is to build up an inverse supply curve via a "quasi-first order approach", integrating local FOCs as in the static case. Given an IC contract (F, τ^Y), suppose F decomposes as $dF_t = \gamma_t\, dt + \beta_t\, dZ_t^G$. Integrating backward from τ^Y and taking time-t expectations wrt $\mathbb{P}^B$, F_t can be seen to satisfy the integral equation

$$F_t = \mathbb{E}_t^B\left[e^{-\rho(\tau^Y - t)} F_{\tau^Y}\right] - \mathbb{E}_t^B\left[\int_t^{\tau^Y} e^{-\rho(s-t)}\left(\gamma_s - \frac{\Delta r}{\sigma}\beta_s - \rho F_s\right) ds\right].$$

Now recall the IC constraint derived in Lemma 9, which is a FOC for local optimality of truthtelling:

$$b - \rho F_t + \gamma_t - \frac{\Delta r}{\sigma}\beta_t \leq 0.$$

Inserting this inequality into the integral equation yields a quasi-inverse supply curve

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

A true inverse supply curve is obtained by the additional assumption that the FOC is satisfied with equality everywhere, and the surface term is pinned down by the limited liability constraint $F_{\tau^Y} \geq 0$.

Framed in this way, the techniques of the current paper can be understood as separately developing two insights which are typically intertwined in static mechanism design. The price approach directly analyzes the control problem posed by the FOCs for local optimality; while in parallel the quantity approach builds up an inverse supply curve from first principles without resorting to FOCs.

8 Conclusion

I have studied the problem of a firm seeking to dynamically elicit the information of a rent-seeking agent about when to terminate a project with a limited lifespan. I show that the firm's problem can be analyzed through the lens of classic monopoly theory, as the firm faces a tight price-quantity tradeoff between payments to the agent and output from the project. I find that the firm's optimal contract can be derived by eliminating either price or quantity from the objective function and optimizing with respect to the remaining variable alone. While the price approach is equivalent to a standard recursive analysis using continuation utility as a state variable, the quantity approach provides a much sharper characterization of the firm's optimal contract.

In an optimal contract, the firm tracks its posterior beliefs about the project's state *as if no agent were present* even though the agent truthfully reports whether the state has switched. It optimally terminates the project the first time either the agent reports the state has changed or its "virtual beliefs" fall below a threshold. This threshold solves a modified version of the unassisted optimal stopping problem, from which I obtain simple and intuitive comparative statics. The threshold is strictly lower than the optimal threshold in the firm's unassisted optimal stopping problem, and is decreasing in the informativeness of news about the state, increasing in the rate of state transitions, and increasing in the severity of the agent's incentive misalignment. The benefits of hiring an agent can therefore be succinctly summarized: the agent's information completely eliminates late termination and partially mitigates early termination universally across all projects and output histories. The optimal payment rule can also be stated briefly: when the agent is terminated for reporting a state change, he receives a lump sum golden parachute equal to his expected lifetime rents from withholding his report forever.

While I have assumed that the firm's project produces a Brownian output flow, the major conclusions of the paper all hold under alternative Poisson technologies involving either steady output flows and occasional costly "breakdowns" or steady expenditures and occasional profitable "breakthroughs." Another natural generalization of the model is to agents with imperfect knowledge of the state. This change creates an additional incentive for the agent to delay reporting bad news inducing project termination, as there is a chance his beliefs about the state will improve in the future and he will not have to report the bad news after all. In this setting the agent must be additionally compensated for the value of his real option to wait and see.

Another generalization is to projects with many states. In this setting the agent expects the project to continue deteriorating over time after a termination threshold is reached, so that delaying a report produces a growing belief asymmetry which can be profitably exploited by the agent. (By contrast, with only two states the belief asymmetry is constant over time, and equal to its initial value at the time of the state switch.) The agent must therefore be additionally compensated for the wedge between the marginal belief asymmetry at the time of a state switch and his average belief asymmetry over the lifetime of the contract.

References

Baron, D.P. and D. Besanko, "Regulation, Asymmetric Information, and Auditing," *The RAND Journal of Economics*, 1984, *15* (4), 447–470.

Battaglini, M., "Long-Term Contracting with Markovian Consumers," *American Economic Review*, 2005, *95* (3), 637–658.

Besanko, D., "Multi-period contracts between principal and agent with adverse selection," *Economics Letters*, 1985, *17* (1), 33 – 37.

Courty, P. and H. Li, "Sequential Screening," *The Review of Economic Studies*, 2000, *67* (4), 697–717.

Cvitanic, J., X. Wan, and H. Yang, "Dynamics of Contract Design with Screening," *Management Science*, 2013, *59* (5), 1229–1244.

DeMarzo, P.M. and Y. Sannikov, "Optimal Security Design and Dynamic Capital Structure in a Continuous-Time Agency Model," *The Journal of Finance*, 2006, *61* (6), 2681–2724.

Eső, P. and B. Szentes, "Dynamic Contracting: An Irrelevance Result," *Theoretical Economics*, 2016. Forthcoming.

Garrett, D.F. and A. Pavan, "Managerial Turnover in a Changing World," *Journal of Political Economy*, 2012, *120* (5), 879–925.

Grenadier, S.R., A. Malenko, and N. Malenko, "Timing Decisions in Organizations: Communication and Authority in a Dynamic Environment," 2015. Working paper.

Guo, Y., "Dynamic Delegation of Experimentation," 2015. Working paper.

Harrison, J.M., *Brownian Models of Performance and Control*, Cambridge University Press, 2013.

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

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

Myerson, R.B., "Optimal auction design," *Mathematics of Operations Research*, 1981, *6* (1), 58–73.

Pavan, A., I. Segal, and J. Toikka, "Dynamic Mechanism Design: A Myersonian Approach," *Econometrica*, 2014, *82* (2), 601–653.

Peskir, G. and A. Shiryaev, *Optimal Stopping and Free-Boundary Problems*, Birkhäuser Basel, 2006.

Prat, J. and B. Jovanovic, "Dynamic contracts when the agent's quality is unknown," *Theoretical Economics*, 2013, *9* (3), 865–914.

Sannikov, Y., "A Continuous- Time Version of the Principal-Agent Problem," *The Review of Economic Studies*, 2008, *75* (3), 957–984.

_ , "Moral Hazard and Long-Run Incentives," 2014. Working paper.

Spear, S.E. and S. Srivastava, "On Repeated Moral Hazard with Discounting," *The Review of Economic Studies*, 1987, *54* (4), 599–617.

Williams, N., "On Dynamic Principal-Agent Problems in Continuous Time," 2008. Working paper.

_ , "Persistent Private Information," *Econometrica*, 2011, *79* (4), 1233–1275.

_ , "A solvable continuous time dynamic principal-agent model," *Journal of Economic Theory*, 2015, *159, Part B*, 989 – 1015.

Appendices

A Proofs of theorems

A.1 Proof of Theorem 1

Fix F^* and V as described in the theorem statement. Lemma 24 in the technical appendix establishes that V is a positive, strictly increasing, strictly concave function on $[0, F^*)$, and that maximizers γ^*, β^* of the HJB equation exist on $[0, F^*]$ and are continuous except possibly at F^*. And by Proposition 3, $V(F^*) = \frac{r_G - \alpha F^*}{\rho+\alpha}$ and $F^* < b/\rho$. Finally, as this proposition also implies $V''(F^*) < 0$, the maximizers (γ^*, β^*) are continuous on $[0, F^*]$.

Extend V to $\mathbb{R}_+$ by setting $V(F_0) = V(F^*)$ for $F_0 > F^*$. Then V is a C^1, piecewise C^2 function on this extended domain and satisfies

$$(\rho + \alpha)V(F_0) \geq \sup_{(\gamma,\beta)\in \text{IC}(F)} \left\{ r_G - \alpha F + \gamma V'(F_0) + \frac{1}{2}\beta^2 V''(F_0) \right\}$$

everywhere, with equality for $F_0 \in [0, F^*]$.

I now carry out a verification argument for IC-B contracts $(\widetilde{F}, \widetilde{\tau})$ such that $\widetilde{F}$ is cadlag and bounded above by b/ρ and $\widetilde{\tau} = \inf\{t \ : \ \widetilde{F}_t = 0\}$. Lemma 6 justifies the boundedness restriction and form of the termination policy, while Lemma 23 in the technical appendix justifies the focus on cadlag golden parachutes. Verification within this class therefore suffices to prove that V is an upper bound on the lifetime value of any IC-B contract.

Fix an IC-B contract $(\widetilde{F}, \widetilde{\tau})$ satisfying the conditions just described. Wlog assume that $\widetilde{F} = F^{\widetilde{\tau}}$, as the stopped version of any contract remains IC-B and yields the same profits for the firm in all cases. Define the process M by

$$M_t = \frac{b}{\rho}(1 - e^{-\rho t}) + e^{-\rho t} F_t.$$

By Lemma 22 in the technical appendix, M is a $\mathbb{P}^B$-supermartingale, and as F is cadlag and bounded above by b/ρ, so is M. M then satisfies the conditions of the Doob-Meyer decomposition theorem (see Theorem 1.4.10 of Karatzas and Shreve (1991) for a statement), hence

$$M = \widetilde{F}_0 + X - A$$

for some $\mathbb{F}^Y$-adapted process A which is increasing, right-continuous, and starts at zero; and some $\mathbb{F}^Y$-adapted process X which is a cadlag $\mathbb{P}^B$-martingale starting at zero.

Further, by the martingale representation theorem there exists an $\mathbb{F}^Y$-adapted, progressively measurable process β^M, satisfying $\mathbb{P}^B\left\{\int_0^t (\beta^M)_s^2\, ds < \infty\right\} = 1$ for all t, such that

$$X_t = \int_0^t \beta_s^M\, dZ_s^B$$

for all t. And as $Z_t^B = Z_t^G + \frac{\Delta r}{\sigma}t$, X may be equivalently decomposed as

$$X_t = \int_0^t \frac{\Delta r}{\sigma}\beta_s^M\, ds + \int_0^t \beta_s^M\, dZ_s^G.$$

Now write $\widetilde{F}_t$ in terms of M_t as

$$\widetilde{F}_t = e^{\rho t}M_t - \frac{b}{\rho}(e^{\rho t} - 1),$$

and use Ito's lemma to obtain

$$\widetilde{F}_t = \widetilde{F}_0 + \int_0^t e^{\rho s}\left(\rho M_s - b + \frac{\Delta r}{\sigma}\beta_s^M\right) ds + \int_0^t e^{\rho s} dA_s + \int_0^t e^{\rho s}\beta_s^M\, dZ_s^G,$$

or equivalently,

$$\widetilde{F}_t = \widetilde{F}_0 + \int_0^t \left(\rho\widetilde{F}_t - b + \frac{\Delta r}{\sigma}\widetilde{\beta}_s\right) dt - \int_0^t e^{\rho s}\, dA_s + \int_0^t \widetilde{\beta}_s\, dZ_s^G,$$

where $\widetilde{\beta}_t \equiv e^{\rho t}\beta_t^M$. Let $\widetilde{A}_t \equiv \int_0^t e^{\rho s}\, dA_s$. Then $\widetilde{A}$ is a monotone increasing right-continuous process started at zero. Thus $\widetilde{F}$ is a semimartingale to which Ito's lemma, generalized to processes with jumps, can be applied. Let $\Delta\widetilde{A}_t \equiv \widetilde{A}_t - \lim_{s\uparrow t}\widetilde{A}_t$ track the jumps of $\widetilde{A}$, of which there are at most countably many pathwise. I will let $D_t \equiv \widetilde{A}_t - \sum_{0\le s\le t}\Delta\widetilde{A}_t$ denote the continuous part of A.

In general it is not assured that $\widetilde{\beta}$ is sufficiently regular to guarantee that $\int \widetilde{\beta}\, dZ^G$ is a $\mathbb{P}^G$-martingale. I therefore perform the verification procedure using a localized version of the process. For each $N = 1, 2, ...$, define the $\mathbb{F}^Y$-stopping time $\tau^{l,N} \equiv \inf\left\{t\ :\ \left|\int_0^t \widetilde{\beta}_s\, dZ_s^G\right| \ge N\right\}$. For each N, the Ito isometry implies that $\mathbb{E}^G\left[\int_0^{t\wedge\tau^{l,N}} \widetilde{\beta}_s^2\, ds\right] \le N^2$, and so $\int_0^{t\wedge\tau^{l,N}} \widetilde{\beta}_s\, dZ_s^G$ is a $\mathbb{P}^G$-martingale.

Now fix a time t. Ito's lemma says that

$$
\begin{aligned}
e^{-(\rho+\alpha)(t\wedge\widetilde{\tau}\wedge\tau^{l,N})}V(\widetilde{F}_{t\wedge\widetilde{\tau}\wedge\tau^{l,N}}) = V(\widetilde{F}_0) &+ \sum_{0\le s\le t\wedge\widetilde{\tau}\wedge\tau^{l,N}} e^{-(\rho+\alpha)s}\Delta V_s \\
&+ \int_0^{t\wedge\widetilde{\tau}\wedge\tau^{l,N}} e^{-(\rho+\alpha)s}\left(-(\rho+\alpha)V(\widetilde{F}_s)+\frac{1}{2}\widetilde{\beta}_s^2V''(\widetilde{F}_s)\right)ds \\
&+ \int_0^{t\wedge\widetilde{\tau}\wedge\tau^{l,N}} e^{-(\rho+\alpha)s}V'(\widetilde{F}_s)\left(\left(\rho\widetilde{F}_t-b+\frac{\Delta r}{\sigma}\widetilde{\beta}_s\right)ds-dD_s\right) \\
&+ \int_0^{t\wedge\widetilde{\tau}\wedge\tau^{l,N}} e^{-(\rho+\alpha)s}\widetilde{\beta}_sV'(\widetilde{F}_s)\,dZ_s^G,
\end{aligned}
$$

where $\Delta V_t \equiv V(\widetilde{F}_t) - \lim_{s\uparrow t} V(\widetilde{F}_s)$. Note that $\Delta V \le 0$ given that V is monotone increasing and all jumps in $\widetilde{F}$ are downward. As V' is bounded, the final term is a martingale. Taking expectations wrt $\mathbb{P}^G$ leaves

$$
\begin{aligned}
V(\widetilde{F}_0) = \mathbb{E}^G&\left[e^{-(\rho+\alpha)(t\wedge\widetilde{\tau}\wedge\tau^{l,N})}V(\widetilde{F}_{t\wedge\widetilde{\tau}\wedge\tau^{l,N}})\right] - \mathbb{E}^G\left[\sum_{0<s\le t\wedge\widetilde{\tau}\wedge\tau^{l,N}} e^{-(\rho+\alpha)s}\Delta V_s\right] \\
&- \mathbb{E}^G\left[\int_0^{t\wedge\widetilde{\tau}\wedge\tau^{l,N}} e^{-(\rho+\alpha)s}\left(-(\rho+\alpha)V(\widetilde{F}_s)+\frac{1}{2}\widetilde{\beta}_s^2V''(\widetilde{F}_s)\right)ds\right] \\
&- \mathbb{E}^G\left[\int_0^{t\wedge\widetilde{\tau}\wedge\tau^{l,N}} e^{-(\rho+\alpha)s}V'(\widetilde{F}_s)\left(\left(\rho\widetilde{F}_t-b+\frac{\Delta r}{\sigma}\widetilde{\beta}_s\right)ds-dD_s\right)\right].
\end{aligned}
$$

As $\Delta V \le 0$, D is monotone increasing, and $V' \ge 0$, this expression implies the inequality

$$
\begin{aligned}
V(\widetilde{F}_0) \ge \mathbb{E}^G&\left[e^{-(\rho+\alpha)(t\wedge\widetilde{\tau}\wedge\tau^{l,N})}V(\widetilde{F}_{t\wedge\widetilde{\tau}\wedge\tau^{l,N}})\right] \\
&- \mathbb{E}^G\left[\int_0^{t\wedge\widetilde{\tau}\wedge\tau^{l,N}} e^{-(\rho+\alpha)s}\left(-(\rho+\alpha)V(\widetilde{F}_s)+\frac{1}{2}\widetilde{\beta}_s^2V''(\widetilde{F}_s)\right)ds\right] \\
&- \mathbb{E}^G\left[\int_0^{t\wedge\widetilde{\tau}\wedge\tau^{l,N}} e^{-(\rho+\alpha)s}V'(\widetilde{F}_s)\left(\rho\widetilde{F}_t-b+\frac{\Delta r}{\sigma}\widetilde{\beta}_s\right)ds\right].
\end{aligned}
$$

Letting $\widetilde{\gamma}$ on $[0,t]$ via $\widetilde{\gamma}_s = \frac{\Delta r}{\sigma}\widetilde{\beta}_s - (b-\rho\widetilde{F}_s)$, the previous inequality may be written

$$
\begin{aligned}
V(\widetilde{F}_0) \ge \mathbb{E}^G&\left[e^{-(\rho+\alpha)(t\wedge\widetilde{\tau}\wedge\tau^{l,N})}V(\widetilde{F}_{t\wedge\widetilde{\tau}\wedge\tau^{l,N}})\right] \\
&- \mathbb{E}^G\left[\int_0^{t\wedge\widetilde{\tau}\wedge\tau^{l,N}} e^{-(\rho+\alpha)s}\left(-(\rho+\alpha)V(\widetilde{F}_s)+\widetilde{\gamma}_sV'(\widetilde{F}_s)+\frac{1}{2}\widetilde{\beta}_s^2V''(\widetilde{F}_s)\right)ds\right].
\end{aligned}
$$

Note that $(\widetilde{\gamma}_s, \widetilde{\beta}_s) \in \text{IC}(\widetilde{F}_s)$ for all $s \in [0, t]$, so the extension of the HJB equation stated at the beginning of the proof implies

$$V(\widetilde{F}_0) \geq \mathbb{E}^G \left[e^{-(\rho+\alpha)(t \wedge \widetilde{\tau} \wedge \tau^{l,N})} V(\widetilde{F}_{t \wedge \widetilde{\tau} \wedge \tau^{l,N}}) + \int_0^{t \wedge \widetilde{\tau} \wedge \tau^{l,N}} e^{-(\rho+\alpha)s} \left(r_G - \alpha \widetilde{F}_s \right) ds \right].$$

Finally, take $N \to \infty$ and $t \to \infty$. The interior of the expectation is uniformly bounded over all t and N, so the bounded converge theorem yields

$$V(\widetilde{F}_0) \geq \mathbb{E}^G \left[e^{-(\rho+\alpha)\widetilde{\tau}} V(\widetilde{F}_{\widetilde{\tau}}) + \int_0^{\widetilde{\tau}} e^{-(\rho+\alpha)s} \left(r_G - \alpha \widetilde{F}_s \right) ds \right].$$

As $\widetilde{F}_{\widetilde{\tau}} = 0$ and $V(0) = 0$, this reduces to

$$V(\widetilde{F}_0) \geq \mathbb{E}^G \left[\int_0^{\widetilde{\tau}} e^{-(\rho+\alpha)s} \left(r_G - \alpha \widetilde{F}_s \right) ds \right].$$

By Lemma 8, the rhs is the expected profit of the contract $\widetilde{F}$. Therefore $V(\widetilde{F}_0)$ is an upper bound on the expected profits of $\widetilde{F}$. As F^* is a global maximum of V, it follows that $V(F^*)$ is as well.

The final step of verification is establishing that the optimal controls to the HJB equation trace out a contract whose profits are equal to $V(F^*)$. Extend γ^* and β^* continuously to all F_0 by setting $\gamma^*(F_0) = \gamma^*(0)$ and $\beta^*(F_0) = \beta^*(0)$ for all $F_0 < 0$, and similarly $\gamma^*(F_0) = \gamma^*(F^*)$ and $\beta^*(F_0) = 0$ for all $F_0 > F^*$.

Lemma 14. *There exists a weak solution to* (3) *for all time. If* $\lim_{F_0 \to F^*} V'(F_0)V'''(F_0)$ *exists and is finite, then there exists a unique strong solution* F *for all time. Any strong solution* F *satisfies* $F \leq F^*$ *a.s.*

Proof. Because γ^* and β^* are continuous and bounded, Theorem 5.4.22 of Karatzas and Shreve (1991) ensures existence of a weak solution. To show existence of a unique strong solution, I establish under the additional condition in the lemma statement that γ^* and β^* are both globally Lipschitz continuous, from which Theorem 5.2.9 and 5.2.5 of Karatzas and Shreve (1991) ensure strong existence and uniqueness, respectively.

It is sufficient to show that β^* is differentiable on $[0, F^*]$ with bounded derivative. For then γ^* is differentiable with bounded derivative as well, and the mean value theorem implies that both are Lipschitz continuous on $[0, F^*]$. Lipschitz continuity on the extended domain follows trivially given that γ^* and β^* are constant outside of $[0, F^*]$.

Recall from the proof of Proposition 3 that for $F_0 \in [0, F^*)$, V satisfies

$$V''(F_0) = -\left(\frac{\Delta r}{\sigma}\right)^2 \frac{V'(F_0)^2}{\Gamma(F_0)},$$

where $\Gamma(F_0) \equiv (\rho+\alpha)V(F_0) - (r_G - \alpha F_0) + (b - \rho F_0)V'(F_0) > 0$. As both V' and Γ are continuously differentiable, $V'''(F_0)$ exists and is a continuous function on $[0, F^*)$. Then the derivative of β^* exists, is continuous, and is given by

$$\frac{d\beta^*}{dF_0}(F_0) = \frac{\Delta r}{\sigma}\left(-1 + \frac{V'(F_0)}{V''(F_0)^2}V'''(F_0)\right)$$

for $F_0 < F^*$. Recall that by Proposition $V''(F^*) < 0$. Then under the assumption of the lemma statement, the limit of this expression exists and is finite as $F_0 \to F^*$. Finally, the mean value theorem implies that $\frac{d\beta^*}{dF_0}(F^*) = \lim_{F_0 \to F^*} \frac{d\beta^*}{dF_0}(F_0)$. Thus β^* is continuously differentiable on the compact interval $[0, F^*]$, meaning it has a bounded derivative on that domain.

Now, fix any strong solution F. Suppose that on some $\omega \in \Omega$, F satisfies $F_t(\omega) > F^*$ for some $t > 0$. Let $t' = \sup\{s < t \ : F_s(\omega) = F^*\}$. Then a.s. $F_t(\omega)$ satisfies the ODE

$$\frac{df}{ds} = \gamma^*(F^*)$$

on $[t', t]$ with boundary conditions $f(t') = F^*$ and $f(t) = F_t(\omega) > F^*$. This is a contradiction of the fact that this ODE has a unique global solution f satisfying $f(s) > f(t)$ for all $s < t$ given $\gamma^*(F^*) < 0$. Thus such paths can occur with at most probability zero. In other words, $F \leq F^*$ a.s. □

Assume a strong solution F to (3) exists for all time, and let $\tau^Y = \inf\{t \ : \ F_t = 0\}$. I show that the expected profits of F are precisely $V(F^*)$. Because β^* is bounded, $\int_0^t \beta^*(F_s)V'(F_s)\,dZ_s^G$ is a martingale for all time, so no regularization is needed for this step of verification. Fix t and use Ito's lemma to write

$$\begin{aligned} &e^{-(\rho+\alpha)(t\wedge\tau^Y)}V(F_{t\wedge\tau^Y}) \\ &= V(F^*) + \int_0^{t\wedge\tau^Y} e^{-(\rho+\alpha)s}\left(-(\rho+\alpha)V(F_s) + \gamma^*(F_s)V'(F_s) + \frac{1}{2}\beta^*(F_s)^2 V''(F_s)\right) ds \\ &\quad + \int_0^{t\wedge\tau^Y} e^{-(\rho+\alpha)s}\beta^*(F_s)V'(F_s)\,dZ_s^G. \end{aligned}$$

Take expectations wrt $\mathbb{P}^G$ to eliminate the martingale term, leaving

$$\mathbb{E}^G[e^{-(\rho+\alpha)(t\wedge\tau^Y)}V(F_{t\wedge\tau^Y})]$$
$$= V(F^*) + \mathbb{E}^G\left[\int_0^{t\wedge\tau^Y} e^{-(\rho+\alpha)s}\left(-(\rho+\alpha)V(F_s) + \gamma^*(F_s)V'(F_s) + \frac{1}{2}\beta^*(F_s)^2V''(F_s)\right)ds\right].$$

Because $F_t \in [0, F^*]$ for $t \leq \tau^Y$ and $(\gamma^*(F_0), \beta^*(F_0))$ maximize the rhs of the HJB equation for $F_0 \in [0, F^*]$, the final term is equal to $-\mathbb{E}^G\left[\int_0^{t\wedge\tau^Y} e^{-(\rho+\alpha)s}(r_G - \alpha F_s)\,ds\right]$, or after rearrangement

$$V(F^*) = \mathbb{E}^G\left[\int_0^{t\wedge\tau^Y} e^{-(\rho+\alpha)s}(r_G - \alpha F_s)\,ds\right] + \mathbb{E}^G[e^{-(\rho+\alpha)(t\wedge\tau^Y)}V(F_{t\wedge\tau^Y})].$$

Now take $t \to \infty$. The interior of each expectation is uniformly bounded for all time, so the bounded convergence theorem allows limits and expectations to be swapped. As $V(F_{\tau^Y}) = 0$, this leaves

$$V(F^*) = \mathbb{E}^G\left[\int_0^{\tau^Y} e^{-(\rho+\alpha)s}(r_G - \alpha F_s)\,ds\right].$$

By Lemma 8, the rhs are the expected profits of F, which was the desired result.

Finally, note that $\max\{F, 0\}$ is an Ito process for $t \leq \tau^Y$ with progressively measurable increments γ, β, where $\gamma_t = \gamma^*(F_t)$ and $\beta_t = \beta^*(F_t) \geq 0$ satisfy $b - \rho F_t + \gamma_t - \frac{\Delta r}{\sigma}\beta_t = 0$. Then Lemma 21 in the technical appendix implies that $(\max\{F, 0\}, \tau^Y)$ is an IC contract.

A.2 Proof of Theorem 2

Let $\widetilde{V}$ be as stated in the theorem, with $\tau^* \equiv \inf\{t : \pi_t \leq \underline{\pi}\}$ and $\tau^x \equiv \inf\{t : \pi_t \leq x\}$ for any $x \geq \underline{\pi}$. Then by Ito's lemma, for all t

$$e^{-\rho(\tau^*\wedge t-\tau^x)}\widetilde{V}(\pi_{\tau^*\wedge t}) = \widetilde{V}(x) + \int_{\tau^x}^{\tau^*\wedge t} e^{-\rho(s-\tau^x)}\left(-\rho\widetilde{V}(\pi_s) - \alpha\pi_s\widetilde{V}'(\pi_s) + \frac{1}{2}\left(\frac{\Delta r}{\sigma}\right)^2\widetilde{V}''(\pi_s)\right)ds$$
$$+ \int_{\tau^x}^{\tau^*\wedge t} e^{-\rho(s-\tau^x)}\widetilde{V}'(\pi_s)\,d\overline{Z}_s.$$

Now, using the fact that $\widetilde{V}$ satisfies the ODE in the theorem statement on $s \leq \tau^*$, the previous expression reduces to

$$e^{-\rho(\tau^*\wedge t-\tau^x)}\widetilde{V}(\pi_{\tau^*\wedge t}) = \widetilde{V}(x) - \int_{\tau^x}^{\tau^*\wedge t} e^{-\rho(s-\tau^x)}\left(\pi_s r_G - (1-\pi_s)b\right)\, ds + \int_{\tau^x}^{\tau^*\wedge t} e^{-\rho(s-\tau^x)}\widetilde{V}'(\pi_s)\, d\overline{Z}_s.$$

Because $\widetilde{V}'$ is continuous on a compact domain and therefore bounded, the final term is a martingale. Further, by the strong Markov property of stochastic integrals wrt Brownian motions, its conditional expectation wrt $\mathcal{F}_{\tau^x}$ vanishes. Taking conditional expectations then yields

$$\mathbb{E}_{\tau^x}[e^{-\rho(\tau^*\wedge t-\tau^x)}\widetilde{V}(\pi_{\tau^*\wedge t})] = \widetilde{V}(x) - \mathbb{E}_{\tau^x}\left[\int_{\tau^x}^{\tau^*\wedge t} e^{-\rho(s-\tau^x)}\left(\pi_s r_G - (1-\pi_s)b\right)\, ds\right].$$

Now take $t \to \infty$. As the terms in both expectations are bounded, limits and expectations may be swapped, yielding

$$\widetilde{V}(x) = \mathbb{E}_{\tau^x}\left[\int_{\tau^x}^{\tau^*} e^{-\rho(s-\tau^x)}\left(\pi_s r_G - (1-\pi_s)b\right)\, ds\right]$$

given the boundary condition $\widetilde{V}(\underline{\pi}) = 0$. In particular, $\widetilde{V}(1) = \Pi(\tau^*)$.

Now, suppose $\underline{\pi} > b/(b+r_G)$. Then by the expression just derived $\widetilde{V}(x) > 0$ for all $x > \underline{\pi}$, as the flow benefits of operation are always strictly positive on $s \leq \tau^x$ and $\tau^x < \tau^*$. But on the other hand, by the ODE in the theorem statement it follows that

$$\frac{1}{2}\left(\frac{\Delta r}{\sigma}\right)^2 \underline{\pi}^2(1-\underline{\pi})^2\widetilde{V}''(\underline{\pi}) = -(\underline{\pi} r_G - (1-\underline{\pi})b) < 0,$$

meaning $\widetilde{V}''(\underline{\pi}) < 0$ and implying that for x sufficiently close to $\underline{\pi}$, $\widetilde{V}(x) < 0$. Thus $\underline{\pi} \leq b/(b+r_G)$.

Next I claim that $\widetilde{V}$ is strictly increasing on $[\underline{\pi}, b/(b+r_G)]$. This claim is trivial if $\underline{\pi} = b/(b+r_G)$, so assume $\underline{\pi} < b/(b+r_G)$ for the following argument. By a similar argument to the previous paragraph, $\widetilde{V}''(\underline{\pi}) > 0$; then given the boundary conditions at $\underline{\pi}$, $\widetilde{V}'(x) > 0$ for x sufficiently close to $\underline{\pi}$. Suppose $\widetilde{V}'$ vanishes somewhere on $[\underline{\pi}, b/(b+r_G))$, and let $x^* = \inf\{x > \underline{\pi} \,:\, \widetilde{V}'(x) = 0\}$. Continuity ensures $\widetilde{V}'(x^*) = 0$, and by assumption $x^* < b/(b+r_G)$. As $\widetilde{V}'(x) > 0$ for all $x \in (\underline{\pi}, x^*)$, and as $\widetilde{V}(\underline{\pi}) = 0$, it must be that $\widetilde{V}(x^*) > 0$. But then by the ODE $\widetilde{V}''(x^*) > 0$, implying that $\widetilde{V}'(x) < 0$ for x below and sufficiently close to x^*. This contradicts the definition of x^* as the *minimal* point at which $\widetilde{V}'$ vanishes. Thus no such point can exist, i.e. $\widetilde{V}' > 0$ on $[\underline{\pi}, b/(b + r_G))$, proving the claim.

Now fix $x > b/(b + r_G)$. Using the expression for $\widetilde{V}(x)$ derived earlier and the law of iterated expectations, one can show that

$$\widetilde{V}(x) = \mathbb{E}_{\tau^x}\left[\int_{\tau^x}^{\tau^{b/(b+r_G)}} e^{-\rho(s-\tau^x)}\left(\pi_s r_G - (1-\pi_s)b\right)\, ds\right] + \mathbb{E}_{\tau^x}\left[e^{-\rho(\tau^{b/(b+r_G)}-\tau^x)}\widetilde{V}(b/(b+r_G))\right].$$

Then as the first term is strictly positive while the second is non-negative, I conclude that $\widetilde{V}(x) > 0$ for all $x > b/(b + r_G)$, and hence $\widetilde{V}(x) > 0$ for all $x > \underline{\pi}$.

I have now established enough properties of $\widetilde{V}$ to verify the optimality of τ^*. Extend $\widetilde{V}$ to $[0, 1]$ by setting $\widetilde{V}(x) = 0$ for $x < \underline{\pi}$. Note that on this extended domain $\widetilde{V}$ is C^1 and piecewise C^2, so Ito's lemma still applies. Further, given $\underline{\pi} \le b/(b + r_G)$, $\widetilde{V}$ satisfies

$$\rho\widetilde{V}(\pi_0) \ge \pi_0 r_G - (1-\pi_0)b - \alpha\pi_0\widetilde{V}'(\pi_0) + \frac{1}{2}\left(\frac{\Delta r}{\sigma}\right)^2 \pi_0^2(1-\pi_0)^2\widetilde{V}''(\pi_0)$$

on $[0, 1]\setminus\{b/(b+r_G)\}$, with equality iff $\pi_0 > \underline{\pi}$. ($\widetilde{V}''$ is potentially discontinuous at $\pi_0 = \underline{\pi}$, so I do not specify an inequality holding at that point. This will not matter for what follows.) Fix any $\mathbb{F}^Y$-stopping time τ^Y (not necessarily a threshold policy), and use Ito's lemma as before to write

$$e^{-\rho(\tau^Y\wedge t)}\widetilde{V}(\pi_{\tau^Y\wedge t}) = \widetilde{V}(1) + \int_0^{\tau^Y\wedge t} e^{-\rho s}\left(-\rho\widetilde{V}(\pi_s) - \alpha\pi_s V'(\pi_s) + \frac{1}{2}\left(\frac{\Delta r}{\sigma}\right)^2 \pi_s^2(1-\pi_s)^2\widetilde{V}''(\pi_s)\right) ds$$
$$+ \int_0^{\tau^Y\wedge t} e^{-\rho s}\widetilde{V}'(\pi_s)\, d\overline{Z}_s.$$

Given the inequality just derived, I obtain

$$e^{-\rho(\tau^Y\wedge t)}\widetilde{V}(\pi_{\tau^Y\wedge t}) \le \widetilde{V}(1) - \int_0^{\tau^Y\wedge t} e^{-\rho s}(\pi_s r_G - (1-\pi_s)b)\, ds + \int_0^{\tau^Y\wedge t} e^{-\rho s}\widetilde{V}'(\pi_s)\, d\overline{Z}_s.$$

Take expectations to eliminate the final term, and then take $t \to \infty$ and swap limits and expectations. I am left with

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

Given that $\widetilde{V} \geq 0$, the final term is non-negative and we're left with $\widetilde{V}(1) \geq \Pi(\tau^Y)$. Thus given that $\Pi(\tau^*) = \widetilde{V}(1)$, τ^* is an optimal termination policy.

Finally, I verify that $\widetilde{V}$ is the firm's virtual value function, in the sense that $\widetilde{V}(\pi_t) = R_t$ for all time. To verify this, fix any $\tau^Y \geq t$ and use a variant of the argument above to conclude that

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

with equality when $\tau^Y = \inf\{s \geq t \ : \pi_t \leq \underline{\pi}\}$ given that the two inequalities in the verification argument hold with equality for such a policy. Thus

$$\widetilde{V}(\pi_t) = \sup_{\tau^Y \geq t} \mathbb{E}_t\left[\int_t^{\tau^Y} e^{-\rho(s-t)}(\pi_s r_G - (1-\pi_s)b)\,ds\right]$$

for all time a.s., as desired.

A.3 Proof of Theorem 3

Theorem 2 in conjunction with Lemma 11 imply the existence of a threshold $\underline{\pi} \in (0, b/(r_G+b)]$ such that $\tau^* = \inf\{t \ : \ \pi_t \leq \underline{\pi}\}$ optimizes $\Pi(\tau^Y)$ among all $\mathbb{F}^Y$-stopping times τ^Y. Therefore by Proposition 4 (F^*, τ^*) is an optimal IC-B contract. It remains only to establish that (F^*, τ^*) is an IC contract.

Define a process X by

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

for $t \leq \tau^*$, with $X_t = X_{\tau^*}$ for $t > \tau^*$. Then X is a $\mathbb{F}^Y$-adapted square-integrable (indeed bounded) $\mathbb{P}^B$-martingale, so by the martingale representation theorem there exists an $\mathbb{F}^Y$-adapted, progressively measurable process β^X satisfying $\mathbb{E}^B\left[\int_0^t (\beta_s^X)^2\,ds\right] < \infty$ for all t such that

$$X_t = \int_0^t \beta_s^X\,dZ_s^B$$

for all time. Next note that

$$X_t = \frac{b}{\rho}(1 - e^{-\rho t}) + e^{-\rho t} F_t^*$$

for $t \leq \tau^Y$, hence by Ito's lemma

$$F_t^* = F^* + \int_0^t e^{\rho s}(\rho X_s - b)\,ds + \int_0^t e^{\rho s}\beta_s^X\,dZ_s^B = \int_0^t (\rho F_s - b)\,ds + \int_0^t \beta_s\,dZ_s^B,$$

where $\beta_t \equiv e^{\rho t}\beta_t^X$. Written in terms of Z^G, this becomes

$$F_t = F^* + \int_0^t \gamma_s \, ds + \int_0^t \beta_s \, dZ_s^G,$$

where $\gamma_t \equiv \rho F_t - b + \frac{\Delta r}{\sigma}\beta_t$. So F is an Ito process satisfying $b - \rho F_t + \gamma_t - \frac{\Delta r}{\sigma}\beta_t = 0$ for all $t \leq \tau^Y$. Additionally, a change of measure combined with the Cauchy-Schwartz inequality and the Ito isometry implies that $\mathbb{E}^G\left[\int_0^t \beta_s^2 \, ds\right] < \infty$ for all t.

If in addition $\beta \geq 0$, then Lemma 21 in the technical appendix implies that (F^*, τ^*) is an IC contract. This must be true, because F_t^* is increasing in the current value of π_t (see Lemma 12), and π_t increases in response to positive news shocks. If the loading on dZ_t^G were ever negative, it would be possible to find paths of output over which π_t grows but F_t^* shrinks, a contradiction.

B Proofs of propositions

B.1 Proof of Proposition 1

Observe that $\Phi[\Lambda']$ is pathwise right-continuous, so to establish progressive measurability it suffices to show adaptedness. Fix $T \in \mathbb{R}_+$. The no-foresight property implies $\Phi[\Lambda']_T$ and $\mathbf{1}\{\tau[\Lambda'] \leq T\}$ are decomposable as

$$\Phi[\Lambda']_T = \Phi[\Lambda' \wedge T]_T \mathbf{1}\{\Lambda' \leq T\} + \Phi[\infty]_T \mathbf{1}\{\Lambda' > T\}$$

and

$$\mathbf{1}\{\tau[\Lambda'] \leq T\} = \mathbf{1}\{\tau[\Lambda' \wedge T] \leq T\}\mathbf{1}\{\Lambda' \leq T\} + \mathbf{1}\{\tau[\infty] \leq T\}\mathbf{1}\{\Lambda' > T\}.$$

Thus it is sufficient to prove that $\Phi[\Lambda' \wedge T]_T$ and $\mathbf{1}\{\tau[\Lambda' \wedge T] \leq T\}$ are $\mathcal{F}_T$-measurable, as all remaining terms are $\mathcal{F}_T$-measurable by definition. In what follows, I will let $\Sigma = \mathcal{B}([0,T]) \otimes \mathcal{F}^Y_T$.

Define a stochastic process X by $X_t(\omega) = \Phi[t \wedge T]_T(\omega)$ for all $(t,\omega) \in \mathbb{R}_+ \times \Omega$. I first establish that X is $\mathbb{F}^Y$-progressively measurable. For $T' \leq T$, Σ-measurability of the map $(t,\omega) \mapsto X_t(\omega)$ restricted to $[0,T'] \times \Omega$ follows from admissibility. For $T' > T$, fix any $E \in \mathcal{B}(\mathbb{R}_+)$ and note that

$$\begin{aligned} &\{(t,\omega) \in [0,T'] \times \Omega \ : \ X_t(\omega) \in E\} \\ &= \{(t,\omega) \in [0,T] \times \Omega \ : \ \Phi[t]_T(\omega) \in E\} \cup ((T,T'] \times (\Phi[T]_T)^{-1}(E)). \end{aligned}$$

The first set on the rhs is in Σ by assumption of admissibility; meanwhile the second is the direct product of an element of $\mathcal{B}([0,T'])$ and an element of $\mathcal{F}^Y_T$ given adaptedness of $\Phi[T]$, hence is in Σ as well. Then so is their union given that Σ is a sigma-algebra, meaning X is $\mathbb{F}^Y$-progressively measurable. In this case X is trivially $\mathbb{F}$-progressively measurable as well. Thus $X_{\Lambda' \wedge T} = \Phi[\Lambda' \wedge T]_T$ is $\mathcal{F}_{\Lambda' \wedge T}$-measurable, hence $\mathcal{F}_T$-measurable as well, which is the desired result.

$\mathcal{F}_T$-measurability of $\mathbf{1}\{\tau[\Lambda' \wedge T] \leq T\}$ follows by applying the argument of the previous paragraph to the process U defined by $U_t(\omega) = \mathbf{1}\{\tau[t \wedge T] \leq T\}$ for all $(t,\omega) \in \mathbb{R}_+ \times \Omega$, invoking the fact that $\tau[T]$ is an $\mathbb{F}^Y$-stopping time in place of adaptedness of $\Phi[T]_T$ at the appropriate step.

B.2 Proof of Proposition 2

Divide the first equation of the system in the proposition statement through by a_1 and solve for $1/a_1$ to obtain

$$\frac{1}{a_1} = -\frac{1}{\alpha}\left(\rho + 2b\frac{a_2}{a_1}\right).$$

Divide the second equation through by $a_1 a_2$ and insert the equality above to obtain

$$\frac{1}{4}\left(\frac{\Delta r}{\sigma}\right)^2 \frac{a_1}{a_2} + \frac{2br_G}{\alpha}\frac{a_2}{a_1} + b + \frac{r_G \rho}{\alpha} = 0.$$

Multiply through by a_1/a_2 to obtain a quadratic equation satisfied by a_1/a_2 :

$$\frac{1}{4}\left(\frac{\Delta r}{\sigma}\right)^2 \left(\frac{a_1}{a_2}\right)^2 + \left(b + \frac{\rho r_G}{\alpha}\right)\frac{a_1}{a_2} + \frac{2br_G}{\alpha} = 0.$$

Letting $F^* = -a_1/(2a_2)$, F^* is a root of the quadratic polynomial

$$\phi(F) \equiv \frac{\alpha}{2}\left(\frac{\Delta r}{\sigma}\right)^2 F^2 - (\alpha b + \rho r_G)F + br_G.$$

The two roots of ϕ are

$$F^*_\pm = \frac{(\alpha b + \rho r_G) \pm \sqrt{(\alpha b + \rho r_G)^2 - 2\alpha b r_G \left(\frac{\Delta r}{\sigma}\right)^2}}{\alpha \left(\frac{\Delta r}{\sigma}\right)^2}.$$

Now, by assumption $\left(\frac{\Delta r}{\sigma}\right)^2 = \rho - \alpha$. Therefore

$$(\alpha b + \rho r_G)^2 - 2\alpha b r_G \left(\frac{\Delta r}{\sigma}\right)^2 > (\alpha b + \rho r_G)^2 - 2\alpha b \rho r_G = (\alpha b)^2 + (\rho r_G)^2 > 0.$$

So the discriminant of the quadratic is strictly positive, and there exist two distinct positive roots of the equation. Note that

$$\phi(b/\rho) = \frac{\alpha}{2}\left(\frac{\Delta r}{\sigma}\right)^2 \frac{b^2}{\rho^2} - (\alpha b + \rho r_G)\frac{b}{\rho} + br_G = \frac{\alpha}{2}(\rho - \alpha)\frac{b^2}{\rho^2} - \frac{\alpha b^2}{\rho} = -\frac{\alpha^2 b^2}{2\rho^2} < 0,$$

hence $F^*_- < b/\rho < F^*_+$. Recall also that

$$a_1 = -\frac{\alpha}{\rho - b/F^*}.$$

Hence there exists one solution to the system of equations, corresponding to F^*_-, with $a_1 > 0$, and another, corresponding to F^*_+, with $a_1 < 0$. As F^* is positive for both solutions, there is one solution satisfying $a_1 > 0 > a_2$, and another satisfying $a_2 > 0 > a_1$. The first solution is the one claimed in the theorem statement, and I will restrict attention to it going forward.

Now, let $V(F) = a_2 F^2 + a_1 F$. Note that V is strictly increasing and strictly concave on $[0, F^*]$, and that $V(0) = 0$ while $V'(F^*) = 0$ and $V''(F^*) = 2a_2 < 0$. Then inserting V into the rhs of the HJB equation, there are unique maximizers

$$\beta^*(F) = -\frac{\Delta r}{\sigma}\frac{V'(F)}{V''(F)} = \frac{\Delta r}{\sigma}(F^* - F)$$

and

$$\begin{aligned}\gamma^*(F) &= \frac{\Delta r}{\sigma}\beta^*(F) - (b - \rho F) \\ &= \left(\left(\frac{\Delta r}{\sigma}\right)^2 - \rho\right)(F^* - F) - (b - \rho F^*) \\ &= -\alpha(F^* - F) - (b - \rho F^*).\end{aligned}$$

for all $F \in [0, F^*)$. The continuous extension of γ^* and β^* to $F = F^*$ constitute the unique continuous maximizers of the HJB equation on the extended domain. The HJB equation then becomes

$$(\alpha+\rho)(a_2F^2+a_1F) = r_G - \alpha F - (\alpha(F^* - F) + (b - \rho F^*))\,(2a_2F + a_1) + \left(\frac{\Delta r}{\sigma}\right)^2 (F^* - F)^2 a_2.$$

As this equation must hold for all F, the two sides of the equation must match term by term. The zeroeth order term is

$$0 = r_G - a_1((\alpha - \rho)F^* + b) + \left(\frac{\Delta r}{\sigma}\right)^2 F^{*2} a_2.$$

Using the fact that $F^* = -a_1/(2a_2)$ and $\rho - \alpha = \left(\frac{\Delta r}{\sigma}\right)^2$, this may equivalently be written

$$0 = r_G - ba_1 - \frac{1}{4}\left(\frac{\Delta r}{\sigma}\right)^2 \frac{a_1^2}{a_2},$$

which is just a rearrangement of the second equation defining a_1 and a_2. So the zeroeth order terms match.

Meanwhile the first-order terms match if

$$(\alpha+\rho)a_1 = -\alpha - 2a_2((\alpha-\rho)F^* + b) + a_1\alpha - 2a_2F^*\left(\frac{\Delta r}{\sigma}\right)^2.$$

Again using $\rho - \alpha = \left(\frac{\Delta r}{\sigma}\right)^2$, this reduces to

$$\rho a_1 = -\alpha - 2a_2 b,$$

which is just a rearrangement of the first equation defining a_1 and a_2. So the first-order terms match.

Finally, the second-order terms match if

$$(\alpha+\rho)a_2 = 2\alpha a_2 + \left(\frac{\Delta r}{\sigma}\right)^2 a_2.$$

This is implied by $\rho = \alpha + \left(\frac{\Delta r}{\sigma}\right)^2$, so second-order terms match. Thus V satisfies the HJB equation on $[0, F^*]$.

B.3 Proof of Proposition 3

Suppose first that V satisfies the HJB equation. Then Lemma 24 in the technical appendix establishes that V is strictly increasing and strictly concave on $[0, F^*)$, and that

$$\beta^*(F_0) = -\frac{\Delta r}{\sigma}\frac{V'(F_0)}{V''(F_0)}, \quad \gamma^*(F_0) = \frac{\Delta r}{\sigma}\beta^*(F_0) + \rho F - b$$

are maximizers of the HJB equation on $[0, F^*)$. Substituting these maximizers into the HJB equation yields

$$(\rho+\alpha)V(F_0) = r_G - \alpha F_0 + (b - \rho F_0)V'(F_0) - \frac{1}{2}\left(\frac{\Delta r}{\sigma}\right)^2 \frac{V'(F_0)^2}{V''(F_0)}$$

for $F_0 \in [0, F^*)$. And because $V'(F_0) > 0$ and $V''(F_0) < 0$, it must be that $(\rho+\alpha)V(F_0) - (r_G - \alpha F_0) - (b - \rho F_0)V'(F_0) > 0$, so this equation may be solved for $V''(F_0)$ to obtain the ODE in the proposition statement.

Conversely, suppose V is strictly increasing and strictly concave and satisfies the ODE in the proposition statement. Rearranging it yields

$$(\rho+\alpha)V(F_0) = r_G - \alpha F_0 + \left(b - \rho F_0 - \left(\frac{\Delta r}{\sigma}\right)^2 \frac{V'(F_0)}{V''(F_0)}\right) V'(F_0) + \frac{1}{2}\left(\frac{\Delta r}{\sigma}\frac{V'(F_0)^2}{V''(F_0)}\right)^2 V''(F_0),$$

or

$$(\rho+\alpha)V(F_0) = r_G - \alpha F_0 + \gamma^*(F_0)V'(F_0) + \frac{1}{2}\beta^*(F_0)^2 V''(F_0).$$

And the proof of Lemma 24 shows that $(\gamma^*(F_0), \beta^*(F_0))$ maximizes the rhs subject to IC-B given any strictly increasing, strictly concave V. So V satisfies the HJB equation.

Now, for V satisfying the conditions of the proposition, write the ODE satisfied by V on $[0, F^*)$ as $V'' = G(f, V, V')$, where

$$G(f, V, V') = -\frac{1}{2}\left(\frac{\Delta r}{\sigma}\right)^2 \frac{(V')^2}{(\rho+\alpha)V - (r_G - \alpha f) + (b - \rho f)V'}$$

is defined on the domain $U = \mathbb{R}^3 \setminus \{(f, V, V') \ : \ (\rho+\alpha)V - (r_G - \alpha f) + (b - \rho f)V' = 0\}$, with U an open set. Suppose $V(F^*) \neq \frac{r_G - \alpha F^*}{\rho+\alpha}$. Then $x_0 = (F^*, V(F^*), 0) \in U$, and as V is C^2 by assumption, V satisfies $V'' = G(f, V, V')$ on the entire domain $[0, F^*]$. Let . For sufficiently small open sets U' containing x_0, $U' \subset U$ and G is continuous in f and has bounded first derivatives wrt V and V' on U', hence is uniformly Lipschitz continuous wrt (V, V') for fixed f in U. Then by the Picard-Lindelof theorem, there exists a unique solution of $V'' = G(f, V, V')$ locally around x_0. Now observe that $\widetilde{V}(f) = V(F^*)$ is a solution of the ODE sufficiently close to x_0, thus $V(F_0) = V(F^*)$ for F_0 sufficiently close to F^*. This contradicts the fact that V is strictly increasing on $[0, F^*)$. So $V(F^*) = \frac{r_G - \alpha F^*}{\rho+\alpha}$.

Finally, define $\Gamma : [0, F^*] \to \mathbb{R}$ by

$$\Gamma(F_0) \equiv (b - \rho F_0)V'(F_0) + (\rho+\alpha)V(F_0) - (r_G - \alpha F_0).$$

As V is a C^2 function, Γ is a C^1 function. Using Γ, the ODE satisfied by V on $[0, F^*)$ may be written

$$\Gamma(F_0) = -\frac{1}{2}\left(\frac{\Delta r}{\sigma}\right)^2 \frac{V'(F_0)^2}{V''(F_0)}.$$

Given $V'(F_0) > 0$ and $V''(F_0) < 0$ for $F < F^*$, it must be that $\Gamma(F_0) > 0$ for $F < F^*$. And as $V(F^*) = \frac{r_G - \alpha F^*}{\rho+\alpha}$ and $V'(F^*) = 0$, $\Gamma(F^*) = 0$. It must therefore be that $\Gamma'(F^*) \leq 0$. As

$$\Gamma'(F_0) = (b - \rho F_0)V''(F_0) + \alpha(1 + V'(F_0))$$

and $V'(F^*) = 0$, this implies

$$(b - \rho F^*)V''(F^*) \leq -\alpha.$$

As $V''(F^*) \leq 0$, if $F^* \geq b/\rho$ this inequality would be violated. So $F^* < b/\rho$. The inequality would also be violated if $V''(F^*) = 0$, so $V''(F^*) < 0$.

B.4 Proof of Proposition 4

For each $\mathbb{F}^Y$-stopping time τ^Y, let F be the golden parachute defined by

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

Lemma 7 says that

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

Lemma 15. $e^{-\rho(\Lambda\wedge\tau^Y)} F_{\Lambda\wedge\tau^Y} = \mathbb{E}_\Lambda \left[\int_{\tau^Y \wedge \Lambda}^{\tau^Y} e^{-\rho t} b \, dt \right]$ *a.s.*

Proof. Note that $\mathbf{1}\{\tau^Y \geq \Lambda\}$ is measurable wrt $\mathcal{F}_\Lambda$. Then as

$$\int_{\tau^Y\wedge\Lambda}^{\tau^Y} e^{-\rho t} b \, dt = \mathbf{1}\{\tau^Y \geq \Lambda\} \int_\Lambda^{\tau^Y} e^{-\rho t} b \, dt$$

it follows that

$$\mathbb{E}_\Lambda \left[\int_{\tau^Y\wedge\Lambda}^{\tau^Y} e^{-\rho t} b \, dt \right] = \mathbf{1}\{\tau^Y \geq \Lambda\} \mathbb{E}_\Lambda \left[\int_\Lambda^{\tau^Y} e^{-\rho t} b \, dt \right] = 0 = e^{-\rho(\Lambda\wedge\tau^Y)} F_{\Lambda\wedge\tau^Y}$$

on $\{\tau^Y < \Lambda\}$. On the other hand, on $\{\tau^Y \geq \Lambda\}$

$$\mathbb{E}_\Lambda \left[\int_{\tau^Y\wedge\Lambda}^{\tau^Y} e^{-\rho t} b \, dt \right] = \mathbb{E}_\Lambda \left[\int_\Lambda^{\tau^Y} e^{-\rho t} b \, dt \right].$$

Further, on $\{\Lambda = s\}$, a standard property of stopping time sigma algebras yields

$$\mathbb{E}_\Lambda \left[\int_\Lambda^{\tau^Y} e^{-\rho t} b \, dt \right] = \mathbb{E}_s \left[\int_\Lambda^{\tau^Y} e^{-\rho t} b \, dt \right] = \mathbb{E}_s \left[\int_s^{\tau^Y} e^{-\rho t} b \, dt \right].$$

And on $\{s \geq \Lambda\}$
$$\mathbb{E}_s\left[\int_s^{\tau^Y} e^{-\rho t} b\, dt\right] = \mathbb{E}_s^B\left[\int_s^{\tau^Y} e^{-\rho t} b\, dt\right].$$
Thus on $\{\tau^Y \geq \Lambda\} \cap \{\Lambda = s\}$,
$$\mathbb{E}_\Lambda\left[\int_{\tau^Y \wedge \Lambda}^{\tau^Y} e^{-\rho t} b\, dt\right] = \mathbb{E}_s^B\left[\int_s^{\tau^Y} e^{-\rho t} b\, dt\right] = e^{-\rho s} F_s = e^{-\rho(\Lambda \wedge \tau^Y)} F_{\Lambda \wedge \tau^Y}.$$
This relationship holds for all s, proving the result. □

Therefore by the law of iterated expectations,
$$\Pi(\tau^Y) = \mathbb{E}\left[\int_0^{\Lambda \wedge \tau^Y} e^{-\rho t} r_G\, dt - \int_{\Lambda \wedge \tau^Y}^{\tau^Y} e^{-\rho t} b\, dt\right] = \mathbb{E}\left[\int_0^{\Lambda \wedge \tau^Y} e^{-\rho t}(r_G + b)\, dt - \int_0^{\tau^Y} e^{-\rho t} b\, dt\right].$$
Another application of the law of iterated expectations yields
$$\begin{aligned}
\mathbb{E}\left[\int_0^{\Lambda \wedge \tau^Y} e^{-\rho t}(r_G + b)\, dt\right] &= \mathbb{E}\left[\int_0^{\Lambda \wedge \tau^Y} e^{-\rho t}(r_G + b)\, dt\right] \\
&= \mathbb{E}\left[\int_0^{\infty} \mathbf{1}\{t \leq \Lambda\}\mathbf{1}\{t \leq \tau^Y\} e^{-\rho t}(r_G + b)\, dt\right] \\
&= \mathbb{E}\left[\int_0^{\infty} \mathbb{E}_t^Y[\mathbf{1}\{t \leq \Lambda\}]\mathbf{1}\{t \leq \tau^Y\} e^{-\rho t}(r_G + b)\, dt\right] \\
&= \mathbb{E}\left[\int_0^{\tau^Y} \pi_t e^{-\rho t}(r_G + b)\, dt\right].
\end{aligned}$$
Thus
$$\Pi(\tau^Y) = \mathbb{E}\left[\int_0^{\tau^Y} e^{-\rho t}(\pi_t r_G - (1 - \pi_t) b)\, dt\right].$$

B.5 Proof of Proposition 5

Let g be as hypothesized. Fix $x \in [\underline{\pi}, 1]$ and let p^x be the unique strong solution to the SDE
$$p_t = p_0 + \int_0^t \left(-\alpha p_s - \left(\frac{\Delta r}{\sigma}\right)^2 p_s^2(1 - p_s)\right) ds + \int_0^t \frac{\Delta r}{\sigma} p_s(1 - p_s)\, dZ_s^B$$

with initial condition $p_0 = x$ satisfying $0 < p^x_t < 1$ for all $t > 0$ a.s. (See the proof of Lemma 12 for a proof of existence and uniqueness.) Define $\tau^Y_x \equiv \inf\{t \ : \ p^x_t \leq \underline{\pi}\}$. By Ito's lemma

$$\begin{aligned} g(x) &= e^{-\rho\tau^Y_x} g(p^x_{\tau^Y_x}) \\ &\quad - \int_0^{\tau^Y_x} e^{-\rho t} \left(-\rho g(p^x_t) - \left(\alpha p^x_t + \left(\frac{\Delta r}{\sigma}\right)^2 (p^x_t)^2 (1-p^x_t) \right) g'(p^x_t) \right. \\ &\qquad \left. + \frac{1}{2}\left(\frac{\Delta r}{\sigma}\right)^2 (p^x_t)^2(1-p^x_t)^2 g''(p^x_t) \right) dt \\ &\quad - \int_0^{\tau^Y_x} \frac{\Delta r}{\sigma} e^{-\rho t} p^x_t (1-p^x_t)\, dZ^B_t. \end{aligned}$$

As p^x is pathwise continuous a.s., $g(p^x_{\tau^Y_x}) = g(\underline{\pi}) = 0$. Hence

$$\begin{aligned} g(x) = &- \int_0^{\tau^Y_x} e^{-\rho t} \left(-\rho g(p^x_t) - \left(\alpha p^x_t + \left(\frac{\Delta r}{\sigma}\right)^2 (p^x_t)^2 (1-p^x_t) \right) g'(p^x_t) \right. \\ &\qquad \left. + \frac{1}{2}\left(\frac{\Delta r}{\sigma}\right)^2 (p^x_t)^2(1-p^x_t)^2 g''(p^x_t) \right) dt \\ &- \int_0^{\tau^Y_x} e^{-\rho t} \frac{\Delta r}{\sigma} p^x_t (1-p^x_t)\, dZ^B_t. \end{aligned}$$

Given the ODE satisfied by g, this reduces to

$$g(x) = \int_0^{\tau^Y_x} e^{-\rho t} b\, dt - \int_0^{\tau^Y_x} e^{-\rho t} \frac{\Delta r}{\sigma} p^x_t (1-p^x_t)\, dZ^B_t.$$

Now take expectations wrt $\mathbb{P}^B$ to eliminate the martingale term, yielding

$$g(x) = \mathbb{E}^B \left[\int_0^{\tau^Y_x} e^{-\rho t} b\, dt \right] = f(x).$$

For the remainder of this proof, assume f satisfies the ODE in the theorem statement. I first establish that $f' > 0$. Given that f is strictly increasing by Lemma 12 and $f(1) < b/\rho$, $f < b/\rho$. Then for any $x < 1$ such that $f'(x) = 0$, the ODE implies that $f''(x) < 0$. This contradicts the fact that f is strictly increasing, hence $f'(x) > 0$ for all $x < 1$. And when $x = 1$ the ODE implies $f'(x) = (b - \rho f(1))/\alpha > 0$. Thus $f' > 0$.

Then using the implicit function theorem, f^{-1} exists and is a C^2 function, implying V is also a C^2 function given Lemma 12. Differentiating the expression in that lemma twice yields

$$\widetilde{V}'(x) = xV'(f(x))f'(x) + V(x) + f(x) - (1-x)f'(x)$$

and

$$\widetilde{V}''(x) = 2V'(f(x))f'(x) + xf''(x)V'(f(x)) + f'(x)^2V''(f(x)) + 2f'(x) - (1-x)f''(x).$$

Substituting these expressions into the ODE of Theorem 2 and rearranging yields

$$\begin{aligned}
&(\rho+\alpha)V(f(x)) \\
&= r_G - \alpha f(x) + \left(-\alpha x f'(x) + k^2x(1-x)^2f'(x) + \frac{1}{2}k^2x^2(1-x)^2f''(x)\right)V'(f(x)) \\
&\quad + \frac{1}{2}k^2x^2(1-x)^2f'(x)^2V''(f(x)) \\
&\quad + \frac{1-x}{x}\left(\rho f(x) - b + (\alpha x + k^2x^2(1-x))f'(x) - \frac{1}{2}k^2x^2(1-x)^2f''(x)\right),
\end{aligned}$$

where $k \equiv \frac{\Delta r}{\sigma}$. Now, given the ODE satisfied by f, the final term vanishes and the coefficient in front of V' may be simplified, yielding

$$\begin{aligned}
(\rho+\alpha)V(f(x)) = r_G - \alpha f(x) &+ (\rho f(x) - b + k^2x(1-x)f'(x))V'(f(x)) \\
&+ \frac{1}{2}k^2x^2(1-x)^2f'(x)^2V''(f(x)).
\end{aligned}$$

Defining $\beta(x) \equiv kx(1-x)f'(x)$ and $\gamma(x) \equiv \rho f(x) - b + k\beta(x)$, this equation may be written

$$(\rho+\alpha)V(f(x)) = r_G - \alpha f(x) + \gamma(x)V'(f(x)) + \frac{1}{2}\beta(x)^2V''(f(x)).$$

Note that $(\gamma(x), \beta(x)) \in \mathrm{IC}(f(x))$ for all x. Thus

$$(\rho+\alpha)V(F_0) \le \sup_{(\gamma,\beta)\in\mathrm{IC}(F_0)} \left\{ r_G - \alpha F_0 + \gamma V'(F_0) + \frac{1}{2}\beta^2V''(F_0) \right\}$$

for all $F_0 \in [0, F_0^*]$.

To establish the opposite inequality, fix any $F_0 \in (0, F_0^*)$ and $(\gamma, \beta) \in \mathrm{IC}(F_0)$. Without loss suppose that (γ, β) lies in the interior of $\mathrm{IC}(F_0)$, i.e. $b - \rho F_0 + \gamma - k\beta < 0$. For if

$$(\rho+\alpha)V(F_0) \ge r_G - \alpha F_0 + \gamma V'(F_0) + \frac{1}{2}\beta^2V''(F_0)$$

for all such (γ,β), then continuity of the rhs in (γ,β) shows that it holds everywhere on $\mathrm{IC}(F_0)$, and thus that

$$(\rho+\alpha)V(F_0) \geq \sup_{(\gamma,\beta)\in \mathrm{IC}(F_0)} \left\{ r_G - \alpha F_0 + \gamma V'(F_0) + \frac{1}{2}\beta^2 V''(F_0) \right\}.$$

Define a fee process $\widetilde{F}$ by

$$\widetilde{F}_t = F_0 + \gamma t + \beta Z_t^G.$$

Also let $\phi : [0, F^*] \to \mathbb{R}$ be the auxiliary function

$$\phi(f) \equiv -(\rho+\alpha)V(f) + r_G - \alpha f + \gamma V'(f) + \frac{1}{2}\beta^2 V''(f).$$

As V is a C^2 function, ϕ is continuous. It is sufficient to show that $\phi(F_0) \leq 0$.

Suppose by way of contradiction that $\phi(F_0) > 0$. Define a stopping time $\widetilde{\tau}$ by

$$\widetilde{\tau} \equiv \inf\{t \;:\; b - \rho \widetilde{F}_t + \gamma - k\beta = 0 \text{ or } \phi(\widetilde{F}_t) = 0 \text{ or } \widetilde{F}_t \in \{0, F_0^*\}\}.$$

Note that $\widetilde{\tau} > 0$ a.s. given that ϕ is continuous and $\widetilde{F}$ is pathwise continuous a.s. Also, by construction $\phi(\widetilde{F}_t) > 0$ for all $t < \widetilde{\tau}$. By Ito's lemma

$$\begin{aligned} e^{-(\rho+\alpha)\widetilde{\tau}} V(\widetilde{F}_{\widetilde{\tau}}) = V(F_0) &+ \int_0^{\widetilde{\tau}} e^{-(\rho+\alpha)s} \left(-(\rho+\alpha)V(\widetilde{F}_s) + \gamma V'(\widetilde{F}_s) + \frac{1}{2}\beta^2 V''(\widetilde{F}_s) \right) ds \\ &+ \int_0^{\widetilde{\tau}} \beta V'(\widetilde{F}_s)\, dZ_s^G. \end{aligned}$$

The final term is a martingale whose expectation wrt $\mathbb{P}^G$ vanishes. Hence

$$\mathbb{E}^G\left[e^{-(\rho+\alpha)(\widetilde{\tau})} V(\widetilde{F}_{\widetilde{\tau}})\right] = V(F_0) + \mathbb{E}^G\left[\int_0^{\widetilde{\tau}} e^{-(\rho+\alpha)s} \left(-(\rho+\alpha)V(\widetilde{F}_s) + \gamma V'(\widetilde{F}_s) + \frac{1}{2}\beta^2 V''(\widetilde{F}_s) \right) ds\right].$$

By construction

$$\mathbb{E}^G\left[\int_0^{\widetilde{\tau}} e^{-(\rho+\alpha)s} \left(-(\rho+\alpha)V(\widetilde{F}_s) + \gamma V'(\widetilde{F}_s) + \frac{1}{2}\beta^2 V''(\widetilde{F}_s) \right) ds\right] > -\mathbb{E}^G\left[\int_0^{\widetilde{\tau}} e^{-(\rho+\alpha)s}(r_G - \alpha \widetilde{F}_s)\, ds\right],$$

so

$$V(F_0) < \mathbb{E}^G\left[\int_0^{\widetilde{\tau}} e^{-(\rho+\alpha)s}(r_G - \alpha \widetilde{F}_s)\, ds + e^{-(\rho+\alpha)(\widetilde{\tau})} V(\widetilde{F}_{\widetilde{\tau}})\right].$$

But then the contract employing golden parachute $\widetilde{F}$ until $\widetilde{\tau}$ and then following an optimal parachute started at $\widetilde{F}_{\widetilde{\tau}}$ afterward satisfies IC-B and provides strictly higher expected profits than using the optimal contract starting at F_0 for all time. This is a contradiction, hence $\phi(F_0) \leq 0$. As discussed earlier, it follows that

$$(\rho + \alpha)V(F_0) \geq \sup_{(\gamma,\beta)\in \mathrm{IC}(F_0)} \left\{ r_G - \alpha F_0 + \gamma V'(F_0) + \frac{1}{2}\beta^2 V''(F_0) \right\}$$

for all $F_0 \in (0, F_0^*)$. The fact that V is C^2 and $\mathrm{IC}(\cdot)$ is a continuous correspondence implies that the inequality holds also at the boundaries. Hence V satisfies the HJB equation.

As for the boundary conditions, the fact that $V(0) = 0$ is automatic from the fact that any contract with initial golden parachute of 0 must terminate immediately. It remains only to show that $V'(F_0^*) = 0$. Surely $V'(F_0^*) \geq 0$, as otherwise at F^* the rhs of the HJB equation would be infinite. So suppose by way of contradiction that $V'(F_0^*) > 0$. Surely $V''(F^*) < 0$, or else again the rhs of the HJB equation would be infinite. Thus at F_0^* there exist unique maximizers of the HJB equation

$$\beta^*(F_0^*) = -k\frac{V'(F_0^*)}{V''(F_0^*)}, \quad \gamma^*(F_0^*) = -b + \rho F_0^* + k\beta^*(F_0^*).$$

In particular, $\beta^*(F_0^*) > 0$. But we saw earlier that $\beta(x) = kx(1-x)f'(x)$ and $\gamma(x) = \rho f(x) - b + k\beta(x)$ saturate the rhs of the HJB equation for all x, hence given the uniqueness of the maximizers $\beta^*(F_0^*) = 0$. This is a contradiction, so $V'(F_0^*) = 0$.

C Proofs of lemmas

C.1 Proof of Lemma 1

Clearly IC-G and IC-B are 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,

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

or after subtracting the final term from both sides,

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

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

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

Sum these two inequalities to obtain

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

As this inequality holds for arbitrary Λ', $\mathcal{C}$ is incentive-compatible.

C.2 Proof of Lemma 2

Fix an IC-B contract $\mathcal{C} = (\Phi, \tau)$. Define a new contract $\widetilde{\mathcal{C}} = (\widetilde{\Phi}, \widetilde{\tau})$ by setting $\widetilde{\tau}[t] = \tau[t] \wedge t$ for all t, $\widetilde{\Phi}[t]_s = \Phi[t]_s$ for all t and $s < \widetilde{\tau}[t]$, and

$$\widetilde{\Phi}[t]_s = \Phi[t]_{\widetilde{\tau}[t]} + \mathbb{E}^B_{\widetilde{\tau}[t]}\left[\int_{\widetilde{\tau}[t]}^{\tau[t]} e^{-\rho(u-\widetilde{\tau}[t])}(b\,du + d\Phi[t]_u)\right]$$

for all t and $s \geq \widetilde{\tau}[t]$. To compute the agent's expected utility from a particular reporting strategy under $\widetilde{\mathcal{C}}$, I employ the following technical result. Let $\widetilde{\mathbb{E}}^B_t$ be the conditional expectation under $\mathbb{P}^B$ given $\mathcal{F}_t$. (Recall that $\mathbb{E}^B_t$ conditions only on $\mathcal{F}^Y_t$, which does not capture the full information available to the agent at time t.)

Lemma 16. *For any $\mathbb{F}$-stopping time Λ',*

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

Proof. Fix $t \in \mathbb{R}_+$. As $\Phi[t]$ and $\tau[t]$ depend only on the history of output and not on Λ, and as Y is independent of Λ under $\mathbb{P}^B$,

$$\mathbb{E}^B_{\widetilde{\tau}[t]}\left[\int_{\widetilde{\tau}[t]}^{\tau[t]} e^{-\rho(u-\widetilde{\tau}[t])}(b\,du + d\Phi[t]_u)\right] = \widetilde{\mathbb{E}}^B_{\widetilde{\tau}[t]}\left[\int_{\widetilde{\tau}[t]}^{\tau[t]} e^{-\rho(u-\widetilde{\tau}[t])}(b\,du + d\Phi[t]_u)\right].$$

So

$$\widetilde{\Phi}[t]_{\widetilde{\tau}[t]} = \Phi[t]_{\widetilde{\tau}[t]} + \widetilde{\mathbb{E}}^B_{\widetilde{\tau}[t]}\left[\int_{\widetilde{\tau}[t]}^{\tau[t]} e^{-\rho(u-\widetilde{\tau}[t])}(b\,du + d\Phi[t]_u)\right].$$

Further,

$$\int_{\widetilde{\tau}[t]}^{\tau[t]} e^{-\rho(u-\widetilde{\tau}[t])}(b\,du + d\Phi[t]_u) = \mathbf{1}\{\widetilde{\tau}[t] = t\}\int_{\widetilde{\tau}[t]}^{\tau[t]} e^{-\rho(u-\widetilde{\tau}[t])}(b\,du + d\Phi[t]_u),$$

since $\widetilde{\tau}[t] = \tau[t]$ whenever $\widetilde{\tau}[t] < t$, in which case the integral vanishes. Then as $\mathbf{1}\{\widetilde{\tau}[t] = t\}$ is $\mathcal{F}_{\widetilde{\tau}[t]}$-measurable,

$$\widetilde{\mathbb{E}}^B_{\widetilde{\tau}[t]}\left[\int_{\widetilde{\tau}[t]}^{\tau[t]} e^{-\rho(u-\widetilde{\tau}[t])}(b\,du + d\Phi[t]_u)\right] = \mathbf{1}\{\widetilde{\tau}[t] = t\}\widetilde{\mathbb{E}}^B_{\widetilde{\tau}[t]}\left[\int_{\widetilde{\tau}[t]}^{\tau[t]} e^{-\rho(u-\widetilde{\tau}[t])}(b\,du + d\Phi[t]_u)\right].$$

A basic property of of conditional expectations implies that for every $\omega \in \{\widetilde{\tau}[t] = t\}$,

$$\widetilde{\mathbb{E}}^B_{\widetilde{\tau}[t]}\left[\int_{\widetilde{\tau}[t]}^{\tau[t]} e^{-\rho(u-\widetilde{\tau}[t])}(b\,du + d\Phi[t]_u)\right](\omega) = \widetilde{\mathbb{E}}^B_t\left[\int_{\widetilde{\tau}[t]}^{\tau[t]} e^{-\rho(u-\widetilde{\tau}[t])}(b\,du + d\Phi[t]_u)\right](\omega).$$

Hence

$$\widetilde{\mathbb{E}}^B_{\widetilde{\tau}[t]}\left[\int_{\widetilde{\tau}[t]}^{\tau[t]} e^{-\rho(u-\widetilde{\tau}[t])}(b\,du + d\Phi[t]_u)\right] = \mathbf{1}\{\widetilde{\tau}[t] = t\}\widetilde{\mathbb{E}}^B_t\left[\int_{\widetilde{\tau}[t]}^{\tau[t]} e^{-\rho(u-\widetilde{\tau}[t])}(b\,du + d\Phi[t]_u)\right].$$

As the indicator variable is $\mathcal{F}_t$-measurable, it may be moved inside the expectation and then eliminated by the previous identity. Then we may rewrite our expression for $\widetilde{\Phi}[t]_{\widetilde{\tau}[t]}$ as

$$\widetilde{\Phi}[t]_{\widetilde{\tau}[t]} = \Phi[t]_{\widetilde{\tau}[t]} + \widetilde{\mathbb{E}}^B_t\left[\int_{\widetilde{\tau}[t]}^{\tau[t]} e^{-\rho(u-\widetilde{\tau}[t])}(b\,du + d\Phi[t]_u)\right]$$

for all t.

Now fix $\omega \in \Omega$ and let $t' = \Lambda'(\omega)$. (Note that t' is a constant, not a random variable.) Then

$$\widetilde{\mathbb{E}}^B_{t'}\left[\int_{\widetilde{\tau}[t']}^{\tau[t']} e^{-\rho(u-\widetilde{\tau}[t'])}(b\,du + d\Phi[t']_u)\right](\omega) = \widetilde{\mathbb{E}}^B_{\Lambda'}\left[\int_{\widetilde{\tau}[t']}^{\tau[t']} e^{-\rho(u-\widetilde{\tau}[t'])}(b\,du + d\Phi[t']_u)\right](\omega).$$

Decompose the integral inside the expectation as

$$\begin{aligned}\int_{\widetilde{\tau}[t']}^{\tau[t']} e^{-\rho(u-\widetilde{\tau}[t'])}(b\,du + d\Phi[t']_u) &= \mathbf{1}\{\Lambda' = t'\}\int_{\widetilde{\tau}[\Lambda']}^{\tau[\Lambda']} e^{-\rho(u-\widetilde{\tau}[\Lambda'])}(b\,du + d\Phi[\Lambda']_u) \\ &\quad + \mathbf{1}\{\Lambda' \neq t'\}\int_{\widetilde{\tau}[t']}^{\tau[t']} e^{-\rho(u-\widetilde{\tau}[t'])}(b\,du + d\Phi[t']_u).\end{aligned}$$

As $\mathbf{1}\{\Lambda' = t'\}$ and $\mathbf{1}\{\Lambda' \neq t'\}$ are both $\mathcal{F}_{\Lambda'}$-measurable,

$$\begin{aligned}\widetilde{\mathbb{E}}^B_{\Lambda'}\left[\int_{\widetilde{\tau}[t']}^{\tau[t']} e^{-\rho(u-\widetilde{\tau}[t'])}(b\,du + d\Phi[t']_u)\right] &= \mathbf{1}\{\Lambda' = t'\}\widetilde{\mathbb{E}}^B_{\Lambda'}\left[\int_{\widetilde{\tau}[\Lambda']}^{\tau[\Lambda']} e^{-\rho(u-\widetilde{\tau}[\Lambda'])}(b\,du + d\Phi[\Lambda']_u)\right] \\ &\quad + \mathbf{1}\{\Lambda' \neq t'\}\widetilde{\mathbb{E}}^B_{\Lambda'}\left[\int_{\widetilde{\tau}[t']}^{\tau[t']} e^{-\rho(u-\widetilde{\tau}[t'])}(b\,du + d\Phi[t']_u)\right].\end{aligned}$$

In particular,

$$\widetilde{\mathbb{E}}^B_{\Lambda'}\left[\int_{\widetilde{\tau}[t']}^{\tau[t']} e^{-\rho(u-\widetilde{\tau}[t'])}(b\,du + d\Phi[t']_u)\right](\omega) = \widetilde{\mathbb{E}}^B_{\Lambda'}\left[\int_{\widetilde{\tau}[\Lambda']}^{\tau[\Lambda']} e^{-\rho(u-\widetilde{\tau}[\Lambda'])}(b\,du + d\Phi[\Lambda']_u)\right](\omega).$$

Hence

$$\widetilde{\mathbb{E}}^B_{\Lambda'(\omega)}\left[\int_{\widetilde{\tau}[\Lambda'(\omega)]}^{\tau[\Lambda'(\omega)]} e^{-\rho(u-\widetilde{\tau}[\Lambda'(\omega)])}(b\,du + d\Phi[\Lambda'(\omega)]_u)\right](\omega) = \widetilde{\mathbb{E}}^B_{\Lambda'}\left[\int_{\widetilde{\tau}[\Lambda']}^{\tau[\Lambda']} e^{-\rho(u-\widetilde{\tau}[\Lambda'])}(b\,du + d\Phi[\Lambda']_u)\right](\omega)$$

for all $\omega \in \Omega$. Combining this result with the expression for $\widetilde{\Phi}[t]_{\widetilde{\tau}[t]}$ derived earlier yields

$$\widetilde{\Phi}[\Lambda'(\omega)]_{\widetilde{\tau}[\Lambda'(\omega)]}(\omega) = \Phi[\Lambda'(\omega)]_{\widetilde{\tau}[\Lambda'(\omega)]}(\omega) + \widetilde{\mathbb{E}}^B_{\Lambda'}\left[\int_{\widetilde{\tau}[\Lambda']}^{\tau[\Lambda']} e^{-\rho(u-\widetilde{\tau}[\Lambda'])}(b\,du + d\Phi[\Lambda']_u)\right](\omega)$$

for all $\omega \in \Omega$. Finally, by definition of $\widetilde{\Phi}[\Lambda']$ and $\widetilde{\tau}[\Lambda']$,

$$\begin{aligned}\widetilde{\Phi}[\Lambda'(\omega)]_{\widetilde{\tau}[\Lambda'(\omega)]}(\omega) &= \widetilde{\Phi}[\Lambda']_{\widetilde{\tau}[\Lambda'(\omega)]}(\omega)\\ &= \widetilde{\Phi}[\Lambda']_{\widetilde{\tau}[\Lambda'(\omega)](\omega)}(\omega)\\ &= \widetilde{\Phi}[\Lambda']_{\widetilde{\tau}[\Lambda'](\omega)}(\omega)\\ &= \widetilde{\Phi}[\Lambda']_{\widetilde{\tau}[\Lambda']}(\omega).\end{aligned}$$

And analogously, $\Phi[\Lambda'(\omega)]_{\widetilde{\tau}[\Lambda'(\omega)]}(\omega) = \Phi[\Lambda']_{\widetilde{\tau}[\Lambda']}(\omega)$. This establishes the expression in the lemma statement. □

Now fix an arbitrary reporting policy $\Lambda' \geq \Lambda$. In this case $\mathbb{P}|_{\mathcal{F}_{\Lambda'}} = \mathbb{P}^B\big|_{\mathcal{F}_{\Lambda'}}$, hence by the previous lemma

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

The agent's expected utility under policy $\Lambda' \geq \Lambda$ may therefore be written

$$U[\Lambda'] = \mathbb{E}\left[\int_0^{\widetilde{\tau}[\Lambda']} e^{-\rho t}(b\,dt + d\Phi[\Lambda']_t) + \widetilde{\mathbb{E}}_{\Lambda'}\left[\int_{\widetilde{\tau}[\Lambda']}^{\tau[\Lambda']} e^{-\rho s}(b\,ds + d\widetilde{\Phi}[\Lambda']_s)\right]\right].$$

By the law of iterated expectations, this reduces to

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

The agent's expected utility under any reporting strategy $\Lambda' \geq \Lambda$ is therefore the same under $\mathcal{C}$ and $\widetilde{\mathcal{C}}$, so the fact that the former contract is IC-B implies the latter is as well.

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

$$\begin{aligned}\Pi[\widetilde{\mathcal{C}}] &= \mathbb{E}\left[\int_0^{\widetilde{\tau}[\Lambda]} e^{-\rho t}(r_G\,dt - d\Phi[\Lambda]_t) - \widetilde{\mathbb{E}}_\Lambda\left[\int_{\widetilde{\tau}[\Lambda]}^{\tau[\Lambda]} e^{-\rho t}(b\,dt + d\Phi[\Lambda]_s)\right]\right] \\ &= \mathbb{E}\left[\int_0^{\widetilde{\tau}[\Lambda]} e^{-\rho t} r_G\,dt - \int_{\widetilde{\tau}[\Lambda]}^{\tau[\Lambda]} e^{-\rho t} b\,dt - \int_0^{\tau[\Lambda]} e^{-\rho t} d\Phi[\Lambda]_t\right].\end{aligned}$$

By comparison, the firm's profits under $\mathcal{C}$ are

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

Thus

$$\Pi[\widetilde{\mathcal{C}}] - \Pi[\mathcal{C}] = -\mathbb{E}\left[\int_{\widetilde{\tau}[\Lambda]}^{\tau[\Lambda]} e^{-\rho t}(r_B + b)\,dt\right].$$

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

Finally, I establish that $\tau[t] \wedge t = \tau[\infty] \wedge t$ for all $t \in \mathbb{R}_+ \cup \{\infty\}$. This follows because for each $\omega \in \Omega$ either $\tau[\infty](\omega) < t$, in which case $\tau[t](\omega) = \tau[\infty](\omega)$ by the no-foresight property; or else $\tau[\infty](\omega) \geq t$, in which case $\tau[t](\omega) \geq t$ and therefore $\tau[t](\omega) = t$ given $\tau[t] \leq t$.

C.3 Proof of Lemma 3

Fix an IC-B contract $\mathcal{C} = (\Phi, \tau)$ satisfying $\tau[t] \leq t$ for all $t \in \mathbb{R}_+$. Define a stochastic process F by

$$F_t = e^{\rho t}\int_0^t e^{-\rho u} d\Phi[t]_u.$$

I first show that F is $\mathbb{F}^Y$-progressively measurable. Fixing $T \in \mathbb{R}_+$, it is sufficient to show that the map $(t,\omega) \mapsto F_t(\omega)$ on $[0,T] \times \Omega$ is $\mathcal{B}([0,T]) \otimes \mathcal{F}_T^Y$-measurable. Use integration by parts to write

$$F_t = \Phi[t]_t - e^{\rho t}\int_0^t e^{-\rho u}\rho\Phi[t]_u\,du.$$

As $\tau[t] \leq t$, $\Phi[t]_t = \Phi[t]_T$ for all $t \in [0,T]$. Further, $\Phi[t]_u = \Phi[\infty]_u$ whenever $t \geq u$. Hence for all $t \leq T$,

$$F_t = \Phi[t]_T - e^{\rho t}\int_0^t e^{-\rho u}\rho\Phi[\infty]_u\,du.$$

By assumption of admissibility the first term is $\mathcal{B}([0,T])\otimes\mathcal{F}^Y_T$-measurable; and the stochastic process defined by the second term is $\mathcal{F}^Y$-adapted and pathwise continuous, hence progressively measurable, in particular $\mathcal{B}([0,T])\otimes\mathcal{F}^Y_T$-measurable considered as a function of (t,ω) on $[0,T]\times\Omega$. Then their difference is as well, establishing the result.

Define a payment process $\widetilde{\Phi}$ by $\widetilde{\Phi}[t]_s = F_{\tau[t]}\mathbf{1}\{\tau[t] \leq s\}$ for $t \in \mathbb{R}_+$ and $\widetilde{\Phi}[\infty] = 0$. Since $\tau[t] < \infty$ for all finite t by assumption, this construction is well-defined. Let $\widetilde{\mathcal{C}} = (\widetilde{\Phi}, \tau)$. Admissibility of $\widetilde{\mathcal{C}}$ follows from Lemma 4 after invoking the result of Lemma 2 that $\tau[t] = \tau[\infty] \wedge t$ for each t.

The agent's expected utility $\widetilde{\mathcal{C}}$ and reporting policy Λ' is then just

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

This is weakly lower than the agent's expected utility from Λ' under $\mathcal{C}$, and identical whenever $\Lambda' < \infty$ a.s. Thus $\widetilde{\mathcal{C}}$ satisfies IC-B, and the firm's profits under $\widetilde{\mathcal{C}}$ and truthful reporting are

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

C.4 Proof of Lemma 4

Fix $T \in \mathbb{R}_+$, and let $\Sigma = \mathcal{B}([0,T])\otimes\mathcal{F}^Y_T$. As $\mathbf{1}\{\tau[t] \leq T\} = 1$ for every $t \leq T$, the map $(t,\omega) \mapsto \mathbf{1}\{\tau[t](\omega) \leq T\}$ restricted to $[0,T]\times\Omega$ is trivially Σ-measurable. Meanwhile, $\Phi[t]_T(\omega) = F_{\tau^Y\wedge t}(\omega)$ for $(t,\omega) \in [0,T]\times\Omega$. So it is sufficient to show that $(t,\omega) \mapsto F_{\tau^Y\wedge t}(\omega)$ restricted to $[0,T]\times\Omega$ is Σ-measurable. But as F is progressively measurable, so is the stopped process F^{τ^Y}, implying the result.

C.5 Proof of Lemma 5

Fix a contract (F, τ^Y). IC-B holds iff

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

for every $\mathbb{F}$-stopping time $\Lambda' \geq \Lambda$. Rearranging this inequality yields

$$0 \geq \mathbb{E}\left[\mathbf{1}\{\tau^Y > \Lambda\} e^{-\rho \Lambda}\left(\int_\Lambda^{\tau^Y \wedge \Lambda'} e^{-\rho(t-\Lambda)} b\, dt + e^{-\rho(\tau^Y \wedge \Lambda' - \Lambda)} F_{\tau^Y \wedge \Lambda'} - F_\Lambda\right)\right].$$

Then by applying the law of iterated expectations, IC-B is equivalent to the condition that

$$0 \geq \mathbb{E}\left[\mathbf{1}\{\tau^Y > \Lambda\} e^{-\rho \Lambda}\left(\mathbb{E}_\Lambda\left[\int_\Lambda^{\tau^Y \wedge \Lambda'} e^{-\rho(t-\Lambda)} b\, dt + e^{-\rho(\tau^Y \wedge \Lambda' - \Lambda)} F_{\tau^Y \wedge \Lambda'}\right] - F_\Lambda\right)\right]$$

for every $\Lambda' \geq \Lambda$. I will rely on this alternate characterization of IC-B throughout this proof.

Now suppose that for each t and every $\mathbb{F}^Y$-stopping time $\tau' \geq t$, equation (2) holds on the event $\{\tau^Y > t\}$. Fix an $\mathbb{F}$-stopping time $\Lambda' \geq \Lambda$. For each $t \in \mathbb{R}_+$, there exists an $\mathbb{F}^Y$-stopping time $\tau'[t] \geq t$ such that $\tau'[t](\omega) = \Lambda'(\omega)$ for every $\omega \in \{\Lambda = t\}$. So fix $t \in \mathbb{R}_+$. By assumption

$$0 \geq \mathbf{1}\{\tau^Y > t\}\left(\mathbb{E}_t^B\left[\int_t^{\tau^Y \wedge \tau'[t]} e^{-\rho(s-t)} b\, ds + e^{-\rho(\tau^Y \wedge \tau'[t] - t)} F_{\tau^Y \wedge \tau'[t]}\right] - F_t\right).$$

Note that the value of $\int_t^{\tau^Y \wedge \tau'[t]} e^{-\rho(s-t)} b\, ds + e^{-\rho(\tau^Y \wedge \tau'[t] - t)} F_{\tau^Y \wedge \tau'[t]}$ depends only on the history of output. Then on $\{\Lambda = t\}$ the conditional distribution of this expression under $\mathbb{P}^B$ given $\mathcal{F}_t^Y$ is the same as under $\mathbb{P}$ given $\mathcal{F}_t$. Hence for $\omega \in \{\Lambda = t\}$,

$$\begin{aligned}&\mathbb{E}_t^B\left[\int_t^{\tau^Y \wedge \tau'[t]} e^{-\rho(s-t)} b\, ds + e^{-\rho(\tau^Y \wedge \tau'[t] - t)} F_{\tau^Y \wedge \tau'[t]}\right](\omega)\\ = {} &\mathbb{E}_t\left[\int_t^{\tau^Y \wedge \tau'[t]} e^{-\rho(s-t)} b\, ds + e^{-\rho(\tau^Y \wedge \tau'[t] - t)} F_{\tau^Y \wedge \tau'[t]}\right](\omega).\end{aligned}$$

Thus

$$0 \geq \mathbf{1}\{\tau^Y > t\}\mathbf{1}\{\Lambda = t\}\left(\mathbb{E}_t\left[\int_t^{\tau^Y \wedge \tau'[t]} e^{-\rho(s-t)} b\, ds + e^{-\rho(\tau^Y \wedge \tau'[t] - t)} F_{\tau^Y \wedge \tau'[t]}\right] - F_t\right).$$

Clearly by definition of $\tau'[t]$,

$$\begin{aligned}&\mathbf{1}\{\Lambda = t\}\left(\int_t^{\tau^Y\wedge\tau'[t]} e^{-\rho(s-t)}b\,ds + e^{-\rho(\tau^Y\wedge\tau'[t]-t)}F_{\tau^Y\wedge\tau'[t]}\right)\\ &= \mathbf{1}\{\Lambda = t\}\left(\int_\Lambda^{\tau^Y\wedge\Lambda'} e^{-\rho(s-\Lambda)}b\,ds + e^{-\rho(\tau^Y\wedge\Lambda'-\Lambda)}F_{\tau^Y\wedge\Lambda'}\right).\end{aligned}$$

Then given $\mathcal{F}_t$-measurability of $\mathbf{1}\{\Lambda = t\}$,

$$\begin{aligned}&\mathbf{1}\{\Lambda = t\}\mathbb{E}_t\left[\int_t^{\tau^Y\wedge\tau'[t]} e^{-\rho(s-t)}b\,ds + e^{-\rho(\tau^Y\wedge\tau'[t]-t)}F_{\tau^Y\wedge\tau'[t]}\right]\\ &= \mathbb{E}_t\left[\mathbf{1}\{\Lambda = t\}\left(\int_t^{\tau^Y\wedge\tau'[t]} e^{-\rho(s-t)}b\,ds + e^{-\rho(\tau^Y\wedge\tau'[t]-t)}F_{\tau^Y\wedge\tau'[t]}\right)\right]\\ &= \mathbb{E}_t\left[\mathbf{1}\{\Lambda = t\}\left(\int_\Lambda^{\tau^Y\wedge\Lambda'} e^{-\rho(s-\Lambda)}b\,ds + e^{-\rho(\tau^Y\wedge\Lambda'-\Lambda)}F_{\tau^Y\wedge\Lambda'}\right)\right]\\ &= \mathbf{1}\{\Lambda = t\}\mathbb{E}_t\left[\int_\Lambda^{\tau^Y\wedge\Lambda'} e^{-\rho(s-\Lambda)}b\,ds + e^{-\rho(\tau^Y\wedge\Lambda'-\Lambda)}F_{\tau^Y\wedge\Lambda'}\right].\end{aligned}$$

Further, for each $\omega \in \{\Lambda = t\}$ it is a basic property of conditional expectations that

$$\begin{aligned}&\mathbb{E}_t\left[\int_\Lambda^{\tau^Y\wedge\Lambda'} e^{-\rho(s-\Lambda)}b\,ds + e^{-\rho(\tau^Y\wedge\Lambda'-\Lambda)}F_{\tau^Y\wedge\Lambda'}\right](\omega)\\ &= \mathbb{E}_\Lambda\left[\int_\Lambda^{\tau^Y\wedge\Lambda'} e^{-\rho(s-\Lambda)}b\,ds + e^{-\rho(\tau^Y\wedge\Lambda'-\Lambda)}F_{\tau^Y\wedge\Lambda'}\right](\omega).\end{aligned}$$

Combining this result with the identity $\mathbf{1}\{\Lambda = t\}F_t = \mathbf{1}\{\Lambda = t\}F_\Lambda$ yields the inequality

$$0 \geq \mathbf{1}\{\tau^Y > t\}\mathbf{1}\{\Lambda = t\}\left(\mathbb{E}_\Lambda\left[\int_\Lambda^{\tau^Y\wedge\Lambda'} e^{-\rho(s-\Lambda)}b\,ds + e^{-\rho(\tau^Y\wedge\Lambda'-\Lambda)}F_{\tau^Y\wedge\Lambda'}\right] - F_\Lambda\right).$$

As this result holds for every $t \in \mathbb{R}_+$, it implies

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

Taking expectations establishes that F satisfies IC-B.

In the other direction, suppose (F, τ^Y) satisfies IC-B. Define a new golden parachute F'

as follows: for each $(t,\omega) \in \mathbb{R}_+ \times \Omega$, set $F'_t(\omega) = F_t(\omega)$ if $t \geq \tau^Y(\omega)$ and

$$F'_t(\omega) = \sup_{\tau' \geq t} \mathbb{E}^B_t \left[\int_t^{\tau^Y \wedge \tau'} e^{-\rho(s-t)} b\, ds + e^{-\rho(\tau^Y \wedge \tau' - t)} F_{\tau^Y \wedge \tau'} \right] (\omega)$$

if $\tau^Y(\omega) > t$, with the supremum ranging over all $\mathbb{F}^Y$-stopping times $\tau' \geq t$. Clearly $F' \geq F$ by taking $\tau' = t$ on the rhs. I claim that for every t and $\mathbb{F}^Y$-stopping time $\tau' \geq t$,

$$F'_t = \sup_{\tau' \geq t} \mathbb{E}^B_t \left[\int_t^{\tau^Y \wedge \tau'} e^{-\rho(s-t)} b\, ds + e^{-\rho(\tau^Y \wedge \tau' - t)} F'_{\tau^Y \wedge \tau'} \right]$$

when $\tau^Y > t$, from which it follows that (F', τ^Y) satisfies IC-B by the converse just proven. The rhs is weakly greater than the lhs given $F' \geq F$ and the definition of F', so it is sufficient to establish that

$$F'_t \geq \sup_{\tau' \geq t} \mathbb{E}^B_t \left[\int_t^{\tau^Y \wedge \tau'} e^{-\rho(s-t)} b\, ds + e^{-\rho(\tau^Y \wedge \tau' - t)} F'_{\tau^Y \wedge \tau'} \right].$$

Fix $(t,\omega) \in \mathbb{R}_+ \times \Omega$ such that $\tau^Y(\omega) > t$. For each $\varepsilon > 0$ there exists an $\mathbb{F}^Y$-stopping time $\tau^\varepsilon \geq t$ such that

$$\begin{aligned} &\sup_{\tau' \geq t} \mathbb{E}^B_t \left[\int_t^{\tau^Y \wedge \tau'} e^{-\rho(s-t)} b\, ds + e^{-\rho(\tau^Y \wedge \tau' - t)} F'_{\tau^Y \wedge \tau'} \right] (\omega) \\ &\leq \mathbb{E}^B_t \left[\int_t^{\tau^Y \wedge \tau^\varepsilon} e^{-\rho(s-t)} b\, ds + e^{-\rho(\tau^Y \wedge \tau^\varepsilon - t)} F'_{\tau^Y \wedge \tau^\varepsilon} \right] (\omega) + \frac{\varepsilon}{2}. \end{aligned}$$

And by definition of F', there exists a $\widetilde{\tau}^\varepsilon \geq \tau^Y \wedge \tau^\varepsilon$ such that

$$F'_{\tau^Y \wedge \tau^\varepsilon}(\omega) \leq \mathbb{E}^B_{\tau^Y \wedge \tau^\varepsilon} \left[\int_{\tau^Y \wedge \tau^\varepsilon}^{\tau^Y \wedge \widetilde{\tau}^\varepsilon} e^{-\rho(s - \tau^Y \wedge \tau^\varepsilon)} b\, ds + e^{-\rho(\tau^Y \wedge \widetilde{\tau}^\varepsilon - \tau^Y \wedge \tau^\varepsilon)} F_{\tau^Y \wedge \widetilde{\tau}^\varepsilon} \right] (\omega) + e^{\rho((\tau^Y \wedge \tau^\varepsilon)(\omega) - t)} \frac{\varepsilon}{2}.$$

(This stopping time may be constructed by using the definition of F' on each $E^s = \{\tau^Y > \tau^\varepsilon = s\}$ for $s \geq t$ to choose $\widetilde{\tau}^\varepsilon$ on E^s. The inequality is trivially true when $\tau^Y \leq \tau^\varepsilon$.) Thus by the law of iterated expectations

$$\sup_{\tau'\geq t}\mathbb{E}_t^B\left[\int_t^{\tau^Y\wedge\tau'} e^{-\rho(s-t)}b\,ds+e^{-\rho(\tau^Y\wedge\tau'-t)}F'_{\tau^Y\wedge\tau'}\right]$$
$$\leq \mathbb{E}_t^B\left[\int_t^{\tau^Y\wedge\widetilde{\tau}^\varepsilon} e^{-\rho(s-t)}b\,ds+e^{-\rho(\tau^Y\wedge\widetilde{\tau}^\varepsilon-t)}F_{\tau^Y\wedge\widetilde{\tau}^\varepsilon}\right]+\varepsilon.$$

And as $\widetilde{\tau}^\varepsilon\geq t$, the definition of F' implies

$$\sup_{\tau'\geq t}\mathbb{E}_t^B\left[\int_t^{\tau^Y\wedge\tau'} e^{-\rho(s-t)}b\,ds+e^{-\rho(\tau^Y\wedge\tau'-t)}F'_{\tau^Y\wedge\tau'}\right]\leq F'_t+\varepsilon.$$

As this inequality holds for all $\varepsilon>0$, the desired inequality must hold.

Finally, I show that $F'_{\Lambda\wedge\tau^Y}=F_{\Lambda\wedge\tau^Y}$ a.s. For each $\varepsilon>0$, construct an $\mathbb{F}$-stopping time $\sigma^\varepsilon\geq\Lambda$ by setting $\sigma^\varepsilon=\tau^{t,\varepsilon}$ when $\Lambda=t$, where $\tau^{t,\varepsilon}\geq t$ is an $\mathbb{F}^Y$-stopping time satisfying

$$\mathbb{E}_t^B\left[\int_t^{\tau^Y\wedge\tau^{t,\varepsilon}} e^{-\rho(s-t)}b\,ds+e^{-\rho(\tau^Y\wedge\tau^{t,\varepsilon}-t)}F'_{\tau^Y\wedge\tau^{t,\varepsilon}}\right]\geq F'_t-\varepsilon$$

whenever $\tau^Y>t$. (Such a $\tau^{t,\varepsilon}$ must exist given earlier results about F'.) Then by construction

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

whenever $\tau^Y>\Lambda$. Now as (F',τ^Y) satisfies IC-B, the earlier characterization of this property implies

$$0\geq\mathbb{E}\left[\mathbf{1}\{\tau^Y>\Lambda\}e^{-\rho\Lambda}\left(F'_\Lambda-F_\Lambda-\varepsilon\right)\right].$$

This inequality holds for all $\varepsilon>0$, hence

$$0\geq\mathbb{E}\left[\mathbf{1}\{\tau^Y>\Lambda\}e^{-\rho\Lambda}\left(F'_\Lambda-F_\Lambda\right)\right].$$

As $F'\geq F$, it must be that $F'_\Lambda=F_\Lambda$ a.e. on $\{\tau^Y>\Lambda\}$. And by construction $F'_{\tau^Y}=F_{\tau^Y}$, hence $F'_{\tau^Y\wedge\Lambda}=F_{\tau^Y\wedge\Lambda}$ a.s.

C.6 Proof of Lemma 6

Fix an IC-B contract $\mathcal{C} = (F, \tau^Y)$. I first claim that $\tau^Y \leq \inf\{t \;:\; F_t = 0\}$. For suppose not; then there exists some time t and state of the world in which $F_t = 0$, $\tau^Y > t$, and $\theta_t = B$. In this case the agent's utility from reporting a state change is 0, while his expected utility from never reporting the change is strictly positive given that he collects strictly positive expected flow rents and receives a non-negative golden parachute. Hence IC-B is violated, a contradiction of our assumption.

I now show that the contract $\mathcal{C}' = (F', \tau^Y)$ with $F' = \min\{F, b/\rho\}$ satisfies IC-B. Consider first the agent's incentives under $\mathcal{C}'$ at time t in any state of the world such that $\theta_t = B$ and $F_t \leq b/\rho$. The agent's utility from reporting the state change is the same under F and F', while his expected utility from any delayed reporting strategy is weakly lower under F' than under F. Thus given that $\mathcal{C}$ satisfies IC-B, so does $\mathcal{C}'$.

Next consider the agent's incentives at time t in any state of the world such that $\theta_t = B$ and $F_t > b/\rho$. In this case his expected utility under F' from any delayed reporting strategy is at most b/ρ, as F' is bounded above by b/ρ. Meanwhile his expected utility from reporting the state change immediately is exactly b/ρ. Thus IC-B is satisfied for F' everywhere. Note that as $\mathcal{C}'$ implements the same stopping strategy and pays the agent weakly less than $\mathcal{C}$, it must be weakly more profitable for the firm.

Now I construct a new contract $\widetilde{\mathcal{C}} = (\widetilde{F}, \widetilde{\tau}^Y)$ as follows. Whenever $t \leq \tau^Y$, set $\widetilde{F}_t = F'_t$. For $t > \tau^Y$ set

$$\widetilde{F}_t = \max\left\{b/\rho - (b/\rho - F'_{\tau^Y})e^{\rho(t-\tau^Y)}, 0\right\}.$$

For the termination policy, let $\widetilde{\tau}^Y = \inf\{t \;:\; \widetilde{F}_t = 0\}$.

Suppose that in some state of the world $\tau^Y \geq \widetilde{\tau}^Y$. This means that at $\widetilde{F}_{\widetilde{\tau}^Y} = F'_{\widetilde{\tau}^Y}$, but also by definition $\widetilde{F}_{\widetilde{\tau}^Y} = 0$, meaning $F'_{\widetilde{\tau}^Y} = 0$ and thus $\tau^Y \leq \widetilde{\tau}^Y$. Hence $\tau^Y \geq \widetilde{\tau}^Y$ implies $\tau^Y = \widetilde{\tau}^Y$, meaning $\widetilde{\tau}^Y \geq \tau^Y$ in all cases.

Next I check incentive compatibility. First consider times t and states of the world in which $\theta_t = B$ and $\widetilde{\tau}^Y > t \geq \tau^Y$. In this case, note that $\widetilde{F}_t$ is pathwise differentiable, and $\frac{d}{dt}\widetilde{F}_t = \rho\widetilde{F}_t - b$. the change in the agent's utility utility from waiting a moment dt to report is

$$b\,dt - \rho\widetilde{F}_t\,dt + \frac{d}{dt}\widetilde{F}_t\,dt = 0.$$

As this equality holds everywhere until $\widetilde{\tau}^Y$, the agent is indifferent between reporting now and any delayed reporting strategy. Thus IC-B holds in these cases. Next consider states of the world in which $\theta_t = B$ and $\widetilde{\tau}^Y > t$. In this case no delayed reporting strategy τ' such that $\tau^Y \geq \tau'$ is more profitable than immediate reporting given that $\mathcal{C}'$ is IC-B. And on the other hand any delayed reporting strategy τ' yields exactly the same expected utility to the agent as the reporting strategy $\tau' \wedge \tau^Y$, as we've seen that the agent's expected utility is unchanged past that point. Thus IC-B holds in all cases.

Finally, I claim that $\widetilde{\mathcal{C}}$ is at least as profitable for the firm as $\mathcal{C}'$, and therefore as $\mathcal{C}$. Because $\widetilde{\mathcal{C}}$ operates for at least as long as $\mathcal{C}'$ in the good state in all cases and pays the agent no more than $\mathcal{C}'$ whenever $\widetilde{\tau}^Y = \tau^Y$, I need only check the cases in which $\widetilde{\tau}^Y > \tau^Y$. But in this case the path of $\widetilde{F}_t$ is constructed to be (weakly) declining given $b/\rho \geq F'$, and so the firm pays the agent no more stopping at $\widetilde{\tau}^Y$ than he would have stopping at τ^Y; in other words, $\widetilde{F}_{\widetilde{\tau}^Y} \leq \widetilde{F}_{\tau^Y} = F'_{\tau^Y}$. Thus payments under $\widetilde{\mathcal{C}}$ are weakly lower than under $\mathcal{C}'$, proving the claim.

C.7 Proof of Lemma 7

Fix an $\mathbb{F}^Y$-stopping time τ^Y. One feasible reporting policy for the agent at time t is to to never report a state change, i.e. to implement the stopping time $\tau' = \infty$. Then if F is any IC-B golden parachute implementing τ^Y, the IC-B constraint implies

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

It follows that if the golden parachute

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

satisfies IC-B, then it minimizes fees paid to the agent state by state over all IC-B contracts implementing τ^Y, and is therefore profit-maximizing.

Fix a time t such that $\tau^Y > t$ and a delayed reporting policy $\tau' \geq t$, and rewrite the agent's expected utility under τ' using the law of iterated expectations as

$$\begin{aligned}
&\mathbb{E}_t^B\left[\int_t^{\tau'\wedge\tau^Y} e^{-\rho(s-t)}b\,ds + e^{-\rho(\tau'\wedge\tau^Y-t)}F^*_{\tau'\wedge\tau^Y}\right]\\
&= \mathbb{E}_t^B\left[\int_t^{\tau'\wedge\tau^Y} e^{-\rho(s-t)}b\,ds + e^{-\rho(\tau'\wedge\tau^Y-t)}\mathbb{E}^B_{\tau'\wedge\tau^Y}\left[\int_{\tau'\wedge\tau^Y}^{\tau^Y} e^{-\rho(s-\tau'\wedge\tau^Y)}b\,ds\right]\right]\\
&= \mathbb{E}_t^B\left[\int_t^{\tau'\wedge\tau^Y} e^{-\rho(s-t)}b\,ds + \mathbb{E}^B_{\tau'\wedge\tau^Y}\left[\int_{\tau'\wedge\tau^Y}^{\tau^Y} e^{-\rho(s-t)}b\,ds\right]\right]\\
&= \mathbb{E}_t^B\left[\int_t^{\tau^Y} e^{-\rho(s-t)}b\,ds\right] = F_t^*.
\end{aligned}$$

Hence every delayed reporting policy τ' yields the same expected utility under F^* as truthful reporting, i.e. F^* is IC-B.

C.8 Proof of Lemma 8

Given truthful reporting, under any contract F the firm collects flow payments dY until $\tau^Y\wedge\Lambda$ and pays out a lump sum of $F_{\tau^Y\wedge\Lambda}$ at project termination. Its expected profits under F are therefore

$$\Pi = \mathbb{E}\left[\int_0^{\tau^Y\wedge\Lambda} e^{-\rho s}dY_s - e^{-\rho(\tau^Y\wedge\Lambda)}F_{\tau^Y\wedge\Lambda}\right] = \mathbb{E}\left[\int_0^{\tau^Y\wedge\Lambda} e^{-\rho s}r_G\,ds - e^{-\rho(\tau^Y\wedge\Lambda)}F_{\tau^Y\wedge\Lambda}\right].$$

Now I write out the expectation wrt uncertainty in the switching time explicitly, using the fact that the conditional measure on output paths up to time t given $\Lambda\geq t$ is equal to $\mathbb{P}^G$:

$$\Pi = \int_0^\infty dt\,\alpha\,e^{-\alpha t}\,\mathbb{E}^G\left[\int_0^{\tau^Y\wedge t} e^{-\rho s}r_G\,ds - e^{-\rho(\tau^Y\wedge t)}\widetilde{F}_{\tau^Y\wedge t}\right].$$

The expression in brackets is bounded, so I use Fubini's theorem to exchange the order of integration, yielding

$$\begin{aligned}
\Pi &= \mathbb{E}^G\left[\int_0^\infty dt\,\alpha\,e^{-\alpha t}\left(\int_0^{\tau^Y\wedge t} e^{-\rho s} r_G\,ds - e^{-\rho(\tau^Y\wedge t)} F_{\tau^Y\wedge t}\right)\right] \\
&= \mathbb{E}^G\left[\int_0^{\tau^Y} dt\,\alpha\,e^{-\alpha t}\left(\int_0^t e^{-\rho s} r_G\,ds - e^{-\rho t}\widetilde{F}_t\right) + e^{-\alpha\tau^Y}\left(\int_0^{\tau^Y} e^{-\rho s} r_G\,ds + e^{-\rho\tau^Y} F_{\tau^Y}\right)\right] \\
&= \mathbb{E}^G\left[\int_0^{\tau^Y} dt\,\alpha\,e^{-\alpha t}\left(\int_0^t e^{-\rho s} r_G\,ds - e^{-\rho t}\widetilde{F}_t\right) + e^{-\alpha\tau^Y}\int_0^{\tau^Y} e^{-\rho s} r_G\,ds\right].
\end{aligned}$$

(In the last line I have used the fact that $\widetilde{F}_{\tau^Y} = 0$.) Using integration by parts, the first term can be evaluated as

$$\int_0^{\tau^Y} dt\,\alpha\,e^{-\alpha t}\int_0^t e^{-\rho s} r_G\,ds = -e^{-\alpha\tau^Y}\int_0^{\tau^Y} e^{-\rho s} r_G\,ds + \int_0^{\tau^Y} e^{-(\rho+\alpha)t} r_G\,dt.$$

Hence

$$\Pi = \mathbb{E}^G\left[\int_0^{\tau^Y} e^{-(\rho+\alpha)t}(r_G - \alpha\widetilde{F}_t)\,dt\right].$$

C.9 Proof of Lemma 9

I prove this lemma with the aid of Lemma 22 in the technical appendix, which establishes that a contract (F, τ^Y) satisfies IC-B iff M^{τ^Y} is a $\mathbb{P}^B$-supermartingale, where

$$M_t = \frac{b}{\rho}(1 - e^{-\rho t}) + e^{-\rho t} F_t.$$

Fix a contract (F, τ^Y) evolving as the Ito process

$$dF_t = \gamma_t\,dt + \beta_t\,dZ_t^G = \left(\gamma_t - \frac{\Delta r}{\sigma}\beta_t\right)dt + \beta_t\,dZ_t^B$$

for $\mathbb{F}^Y$-adapted, progressively measurable processes γ and β. For every $n = 1, 2, ...$, define $\tau^n \equiv \inf\left\{t\ :\ \int_0^t |\gamma_s|\,ds + \int_0^t \beta_s^2\,ds \geq n\right\}$. And for every t and $m = 1, 2, ...$, define $\sigma^{t,m} \equiv \inf\{s \geq t\ :\ F_s + |\gamma_s| + |\beta_s| \geq m\}$. (Right-continuity of $\mathbb{F}^Y$ under the usual conditions ensures that each $\sigma^{t,m}$ is an $\mathbb{F}^Y$-stopping time.)

Suppose first that F satisfies

$$b - \rho F_t + \gamma_t - \frac{\Delta r}{\sigma}\beta_t \leq 0$$

whenever $t < \tau^Y$. Fix times t and $s > t$, and let n be large enough that $\tau^n > t$. Ito's lemma implies that

$$\mathbb{E}_t^B[M_{s\wedge\tau^n}^{\tau^Y}] = M_t^{\tau^Y} + \mathbb{E}_t^B\left[\int_{t\wedge\tau^Y}^{s\wedge\tau^Y\wedge\tau^n} e^{-\rho u}\left(b - \rho F_u + \gamma_u - \frac{\Delta r}{\sigma}\beta_u\right)du + \int_{t\wedge\tau^Y}^{s\wedge\tau^Y\wedge\tau^n} e^{-\rho u}\beta_u\, dZ_u^B\right].$$

Now, given the regularization imposed by τ^n, the last term is a $\mathbb{P}^B$-martingale whose expectation vanishes. And the remaining term within the expectation is non-positive a.s. by assumption. Hence $\mathbb{E}_t^B[M_{s\wedge\tau^n}^{\tau^Y}] \leq M_t^{\tau^Y}$. This holds for all n sufficiently large, so also

$$\liminf_{n\to\infty} \mathbb{E}_t^B[M_{s\wedge\tau^n}^{\tau^Y}] \leq M_t^{\tau^Y}.$$

Given that M is non-negative, Fatou's lemma allows the limit and expectation on the lhs to be swapped, yielding $\mathbb{E}_t^B[M_s^{\tau^Y}] \leq M_t^{\tau^Y}$. So M^{τ^Y} is a supermartingale, implying by the previous theorem that F satisfies IC-B.

Conversely, suppose that F is an IC-B contract. For this direction, assume additionally that γ and β are pathwise right-continuous. Fix a time $t < \tau^Y$ and take m large enough that $\sigma^{t,m} > t$. (Right-continuity of F, γ, and β ensures that such an m exists.) Use Lemma 22 in the technical appendix to conclude that for all $s > t$,

$$\mathbb{E}_t^B\left[\frac{b}{\rho}(1 - e^{-\rho(s\wedge\tau^Y\wedge\sigma^{t,m})}) + e^{-\rho(s\wedge\tau^Y\wedge\sigma^{t,m})}F_{s\wedge\tau^Y\wedge\sigma^{t,m}}\right] \leq \frac{b}{\rho}(1 - e^{-\rho t}) + e^{-\rho t}F_t.$$

Using Ito's lemma, this is equivalent to

$$\mathbb{E}_t^B\left[\int_t^{s\wedge\tau^Y\wedge\sigma^{t,m}} e^{-\rho u}\left(b - \rho F_u + \gamma_u - \frac{\Delta r}{\sigma}\beta_u\right)du + \int_t^{s\wedge\tau^Y\wedge\sigma^{t,m}} e^{-\rho u}\beta_u\, dZ_u^B\right] \leq 0.$$

Due to the regularization imposed by $\sigma^{t,m}$, the final term is a $\mathbb{P}^B$-martingale whose expectation vanishes, yielding

$$\mathbb{E}_t^B\left[\int_t^{s\wedge\tau^Y\wedge\sigma^{t,m}} e^{-\rho u}\left(b - \rho F_u + \gamma_u - \frac{\Delta r}{\sigma}\beta_u\right)du\right] \leq 0.$$

Now, given regularization under $\sigma^{t,m}$, the integral $\int_t^s e^{-\rho u}\mathbf{1}\{\tau^Y \wedge \sigma^{t,m} \geq u\} \left|b + \gamma_u - \frac{\Delta r}{\sigma}\beta_u\right| \, du$ is uniformly bounded state by state and thus in expectation, and so Fubini's theorem permits swapping the order of integration for this term. Meanwhile the non-negativity of F allows the swapping of the order of integration in the remaining term by Tonelli's theorem. This procedure yields

$$\int_t^s e^{-\rho u}\mathbb{E}_t^B \left[\mathbf{1}\{\tau^Y \wedge \sigma^{t,m} \geq u\} \left(b - \rho F_u + \gamma_u - \frac{\Delta r}{\sigma}\beta_u\right)\right] du \leq 0.$$

Given $\tau^Y \wedge \sigma^{t,m} > t$ by assumption, $\mathbf{1}\{\tau^Y \wedge \sigma^{t,m} \geq u\}$ must converge pathwise to 1 as $u \downarrow t$. Hence given pathwise right-continuity of γ and β, the expression

$$e^{-\rho u}\mathbf{1}\{\tau^Y \wedge \sigma^{t,m} \geq u\} \left(b - \rho F_u + \gamma_u - \frac{\Delta r}{\sigma}\beta_u\right)$$

is pathwise right-continuous at t. Further, $\mathbf{1}\{\sigma^{t,m} \geq u\} \left(b - \rho F_u + \gamma_u - \frac{\Delta r}{\sigma}\beta_u\right)$ is uniformly bounded by $b + \left(1 + \rho + \frac{\Delta r}{\sigma}\right) m$ for all $u \geq t$. Therefore the dominated convergence theorem applies, and

$$e^{-\rho u}\mathbb{E}_t^B \left[\mathbf{1}\{\tau^Y \wedge \sigma^{t,m} \geq u\} \left(b - \rho F_u + \gamma_u - \frac{\Delta r}{\sigma}\beta_u\right)\right]$$

is right-continuous at t.

Now I apply the fundamental theorem of calculus to conclude that

$$\int_t^s e^{-\rho u}\mathbb{E}_t^B \left[\mathbf{1}\{\tau^Y \wedge \sigma^{t,m} \geq u\} \left(b - \rho F_u + \gamma_u - \frac{\Delta r}{\sigma}\beta_u\right)\right] du$$

is (right-)differentiable wrt s at t with derivative $e^{-\rho t}\left(b - \rho F_t + \gamma_t - \frac{\Delta r}{\sigma}\beta_t\right)$. Hence for every $\varepsilon > 0$,

$$\begin{aligned} &\frac{1}{s-t}\int_t^s e^{-\rho u}\mathbb{E}_t^B \left[\mathbf{1}\{\tau^Y \wedge \sigma^{t,m} \geq u\} \left(b - \rho F_u + \gamma_u - \frac{\Delta r}{\sigma}\beta_u\right)\right] du \\ &\geq e^{-\rho t}\left(b - \rho F_t + \gamma_t - \frac{\Delta r}{\sigma}\beta_t\right) - \varepsilon \end{aligned}$$

for s sufficiently close to t. As the lhs is bounded above by zero, this implies

$$e^{-\rho t}\left(b - \rho F_t + \gamma_t - \frac{\Delta r}{\sigma}\beta_t\right) \leq \varepsilon$$

for all $\varepsilon > 0$, i.e. $b - \rho F_t + \gamma_t - \frac{\Delta r}{\sigma}\beta_t \leq 0$.

C.10 Proof of Lemma 10

See Propositions 8.10 and 8.11 of Harrison (2013).

C.11 Proof of Lemma 11

Throughout this proof I freely invoke well-known properties of $U(m, n, z)$. In particular, I will often invoke properties holding when $m > 0$ and $n > m + 2$, as is the case when $m = \beta - 1$ and $n = 2\beta - \frac{2\alpha}{k^2}$ given $\beta > 1 + \frac{2\alpha}{k^2}$.

I begin by deriving a general solution to the ODE

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

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

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

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

$$\rho v_H = \left(-\alpha z + k^2 - \alpha - \frac{k^2}{z+1}\right) z\frac{dv_H}{dz} + \frac{k^2}{2}z^2\frac{d^2 v_H}{dz^2}.$$

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

$$\left(\alpha(\beta-1)z + \rho + \alpha + (\alpha - k^2)(\beta - 1) - \frac{1}{2}k^2(\beta-1)(\beta-2)\right) w_H = (k^2\beta - \alpha - \alpha z)z\frac{dw_H}{dz} + \frac{k^2}{2}z^2\frac{d^2 w_H}{dz}.$$

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

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

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

$$(\beta - 1)w_H = \left(2\beta - \frac{2\alpha}{k^2} - t\right)\frac{dw_H}{dt} + t\frac{d^2 w_H}{dt^2}.$$

This is Kummer's differential equation, which has general solution

$$w_H(t) = C_1 U(m, n, t) + C_2 M(m, n, t),$$

where U and M are Tricomi's and Kummer's confluent hypergeometric functions and $m \equiv \beta - 1$ and $n \equiv 2\beta - \frac{2\alpha}{k^2}$ are both strictly positive. 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 $M(m, n, t)$ is known to diverge in the limit as $t \to \infty$. As the leading term of v_H also diverges in this limit, no solution with $D \neq 0$ can satisfy the desired regularity conditions at $x = 1$. This leaves solutions to the ODE in the lemma statement 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

$$V_H(x) \equiv x^{\beta}(1-x)^{1-\beta}U\left(\beta - 1, 2\beta - \frac{2\alpha}{k^2}, \frac{2\alpha}{k^2}\frac{x}{1-x}\right)$$

for $x \in (0, 1)$. $U(m, n, \cdot)$ is known to be a strictly positive function on $(0, \infty)$ when $m > 0$, so the function $\Gamma(x) \equiv \log V_H(x)$ is also well-defined on $(0, 1)$. As $U(m, n, \cdot)$ is an analytic function on $(0, \infty)$, V_H and Γ are C^2 functions on $(0, 1)$. En route to solving the boundary value problem, I establish several basic facts about V_H, Γ, and their derivatives in the limits

$x \to 0$ and $x \to 1$.

Lemma 17. *The following limiting expressions hold as $x \to 1$:*

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

Proof. Let $m \equiv \beta - 1$, $n \equiv 2\beta - \frac{2\alpha}{k^2}$, $\gamma \equiv \frac{k^2}{2\alpha}$, and $z \equiv \gamma^{-1}\frac{x}{1-x}$. Then I may write V_H as

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

Now I invoke the well-known asymptotic expansion of U as $z \to \infty$ when $z \in \mathbb{R}$. The full expansion is $U(m,n,z) \sim z^{-m}{}_2F_0(m, m-n+1;;-1/z)$, where ${}_2F_0$ is the well-known generalized hypergeometric series. To third order this expansion says that

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

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 V_H to obtain

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

taking $x \to 1$ implies $\lim_{x\to 1} V_H(x) = a_0\gamma^m$, which is the first formula in the lemma. Now differentiate wrt x, noting that $\frac{dz}{dx} = \frac{\gamma^{-1}}{(1-x)^2} = \gamma^{-1}(1+\gamma z)^2$. I obtain

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

Thus

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

Writing this explicitly in terms of model parameters, this is

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

Using the quadratic formula characterizing β, this simplifies to

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

which is the second formula in the lemma statement. The limiting expression for $\Gamma'(x) = V_H'(x)/V_H(x)$ may then be obtained by combining previous results.

Finally, the second derivative of V_H is

$$\begin{aligned} V_H''(x) &= (x+1)\gamma^{m-1}(1+\gamma z)^2\left(a_1 z^{-2} - a_2 z^{-3} + \Phi'(z)\right) \\ &\quad + 2x\gamma^{m-1}(1+\gamma z)^3\left(a_1 z^{-2} - a_2 z^{-3} + \Phi'(z)\right) \\ &\quad + x\gamma^{m-2}(1+\gamma z)^4\left(-2a_1 z^{-3} + 3a_2 z^{-4} + \Phi''(z)\right). \end{aligned}$$

The first term converges to $2\gamma^{m+1}a_1$ as $x \to \infty$. As for the rest, pull out the common factor of $x\gamma^{m-2}(1+\gamma z)^3$ and expand

$$2\gamma\left(a_1 z^{-2} - a_2 z^{-3} + \Phi'(z)\right) + (1+\gamma z)\left(-2a_1 z^{-3} + 3a_2 z^{-4} + \Phi''(z)\right).$$

The result is

$$(-2a_1 + \gamma a_2)z^{-3} + 3a_2 z^{-4} + 2\gamma\Phi'(z) + (1+\gamma z)\Phi''(z).$$

All but the firm term are $O(z^{-4})$, which when multiplied by $(1+\gamma z)^3$ die out as $z \to \infty$. I conclude that

$$\lim_{x\to 1} V_H''(x) = 2\gamma^{m+1}a_1 + \gamma^{m+1}(-2a_1 + \gamma a_2) = \gamma^{m+2}a_2,$$

which is the final formula in the lemma statement. □

Lemma 18. *The following limiting expressions hold as* $x \to 0$:

$$\lim_{x\to 0} V_H(x) = \infty,$$
$$\liminf_{x\to 0} \Gamma'(x) = -\infty.$$

Proof. Let $m \equiv \beta - 1$, $n \equiv 2\beta - \frac{2\alpha}{k^2}$, $\gamma \equiv \frac{k^2}{2\alpha}$, and $z \equiv \gamma^{-1}\frac{x}{1-x}$. Then I may write V_H as

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

A standard property of U is $\lim_{z\to 0} z^{m+1}U(m,n,z) = \infty$ when $m > 0$ and $n > m+2$, proving the result.

Given $\lim_{x\to 0} V_H(x) = \infty$, clearly also $\lim_{x\to 0}\Gamma(x) = \infty$. 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 $\lim_{x\to 0}\Gamma(x) = \infty$. □

These lemmas provide the tools to demonstrate the existence of constants $\underline{\pi} \in (0,1)$ and $C > 0$ such that

$$V(x;C) = \frac{r_G + b}{\rho + \alpha}x - \frac{b}{\rho} + CV_H(x)$$

satisfies $V(\underline{\pi};C) = V'(\underline{\pi};C) = 0$. These boundary conditions are equivalent to $CV_H(\underline{\pi}) = -\frac{r_G+b}{\rho+\alpha}x + \frac{b}{\rho}$ and $CV_H'(\underline{\pi}) = -\frac{r_G+b}{\rho+\alpha}$. 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 $\phi(x) \equiv -\frac{1}{\overline{x}-x}$ and $\overline{x} \equiv \frac{b(\rho+\alpha)}{(r_G+b)\rho} > 0$. If such a solution $\underline{\pi}$ exists, a corresponding C is $C = -\frac{1}{V_H'(\underline{\pi})}\frac{r_G+b}{\rho+\alpha}$. And if $\Gamma'(\underline{\pi}) < 0$, then given strict positivity of V_H it must be that $V_H'(\underline{\pi}) < 0$ and so $C > 0$.

Suppose first that $\overline{x} \leq 1$. Then $\lim_{x\uparrow\overline{x}}\phi(x) = -\infty$ while $\phi(0) = -\overline{x}^{-1}$. Given the continuity of ϕ, it must be bounded on $[0,x_0]$ for any $x_0 < \overline{x}$. Thus given $\liminf_{x\to 0}\Gamma'(x) = -\infty$, there exists an $x_1 \in (0,\overline{x})$ 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 $[x_1,\overline{x}]$. Thus in particular there exists an $x_2 \in (x_1,\overline{x})$ 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 $x^* \in (x_1,x_2)$ such that $\Gamma'(x^*) = \phi(x^*)$. Further, $\Gamma'(x^*) < 0$ given the negativity of ϕ on $[0,\overline{x})$.

Now suppose that $\overline{x} > 1$. In this case ϕ is continuous and decreasing on $[0,1]$, with $\phi(1) = -1/(\overline{x}-1)$. 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

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

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 $x^* \in (x_1, x_2)$ to $\Gamma'(x) = \phi(x)$. Given the negativity of ϕ, it must be that $\Gamma'(x^*) < 0$.

Finally, I must establish that $\widetilde{V}$ can be extended to a C^2 function on $[\underline{\pi}, 1]$. Once I have done this, it is automatic that $\widetilde{V}$ satisfies the ODE on the extended domain simply by taking limits of each side of the ODE as $\pi_0 \to 1$. I have already shown that $\lim_{x\to 1} \widetilde{V}^{(n)}(x)$ all exist and are finite for $n = 0, 1, 2$. Defining $\widetilde{V}(1) = \lim_{x\to 1} \widetilde{V}(x)$ extends $\widetilde{V}$ continuously to $[\underline{\pi}, 1]$. But I must check that the first two derivatives of the resulting function exist at 1 and are continuous. To do this, I invoke the mean value theorem. For any $x \in (\underline{\pi}, 1)$ the MVT ensures existence of a $y(x) \in (x, 1)$ such that

$$\frac{\widetilde{V}(1) - \widetilde{V}(x)}{1 - x} = \widetilde{V}'(y(x)).$$

Taking $x \to 1$ implies $y(x) \to 1$ by the squeeze theorem. Then as the limit of $\widetilde{V}'$ exists, so does the derivative of $\widetilde{V}$ at $x = 1$, and $\widetilde{V}'(1) = \lim_{x\to 1} \widetilde{V}'(x)$. Thus $\widetilde{V}$ is a C^1 function on $[\underline{\pi}, 1]$. Another application of the MVT in exactly the same fashion ensures that $\widetilde{V}$ is C^2.

C.12 Proof of Lemma 12

Lemma 19. *f is a continuous, strictly increasing function on $[\underline{\pi}, 1]$ satisfying $f(\underline{\pi}) = 0$ and $f(1) = F_0^*$.*

By definition, each p^x solves the SDE

$$p_t = p_0 + \int_0^t \left(-\alpha p_s - \left(\frac{\Delta r}{\sigma}\right)^2 p_s^2 (1 - p_s) \right) ds + \int_0^t \frac{\Delta r}{\sigma} p_s (1 - p_s)\, dZ_s^B$$

with initial condition $p_0 = x$. To avoid the possibility that this SDE does not have a unique strong solution, I assume in particular that p^x solves the regularized SDE

$$p_t = p_0 + \int_0^t \left(-\alpha \phi(p_s) - \left(\frac{\Delta r}{\sigma}\right)^2 \phi(p_s)^2 (1 - \phi(p_s)) \right) ds + \int_0^t \frac{\Delta r}{\sigma} \phi(p_s)(1 - \phi(p_s))\, dZ_s^B,$$

where $\phi(y) = \min\{\max\{y, 0\}, 1\}$. (I will show in a moment that a solution to the regularized SDE is also a solution to the original one.) The coefficients of this SDE are continuous, bounded, continuously differentiable on $[0, 1]$ (taking one-sided derivatives at the boundaries), and constant outside $[0, 1]$. Hence they are globally Lipschitz continuous and satisfy a quadratic growth condition, so by Theorems 5.2.5 and 5.2.9 of Karatzas and Shreve (1991) there exists a unique strong solution of the SDE for every initial condition $x \in \mathbb{R}$ for all time.

In particular, strong uniqueness ensures that for each $x \in (0, 1]$, $0 < p_t^x < 1$ for all $t > 0$ a.s. For note that $p = 0$ is a solution to the SDE as well, and so if $p_t = 0$ then $p_0 = 0 < x$. And similarly $p_t = p_0 - \alpha t$ is a solution to the SDE for any $p_0 > 1$ and $t \leq (p_0 - 1)/\alpha$. Thus if $p_t \geq 1$ for $t > 0$ then $p_0 = p_t + \alpha t > x$. This verifies the earlier claim that a solution to the regularized SDE is also a solution to the original one.

In addition, strong uniqueness ensures that for every $x' > x$ in $(0, 1]$, $p_t^{x'} > p_t^x$ for all t a.s. For suppose two solutions p and p' to the SDE satisfy $p_t = p'_t$ for some t. Then strong uniqueness ensures that $p = p'$ a.s. In particular $p_0 = p'_0$ a.s., establishing that $p_t^x \neq p_t^{x'}$ for all time a.s. As solutions to the SDE are pathwise continuous and $p_0^{x'} > p_0^x$, this proves the claim. This result establishes that $\tau_{x'}^Y > \tau_x^Y$ a.s., and so f is strictly increasing in x.

As for continuity, use the fact that solutions to SDEs with globally Lipschitz coefficients are continuous flows wrt time and initial data, i.e. the map $(x, t) \to p_x^t(\omega)$ is continuous in (x, t) for a.e. $\omega \in \Omega$. Fix a state $\omega \in \Omega$ and for each x consider the path $p^x(\omega)$ with associated stopping time $\tau_x^Y(\omega)$. (Explicit references to ω will be suppressed to conserve on notation.) By the reasoning in the previous paragraph, τ_x^Y is strictly increasing in x. So fix x and consider first taking $x' \uparrow x$. Let $\lim_{x' \uparrow x} \tau_{x'}^Y = L$ (guaranteed to exist given monotonicity of τ_x^Y); then $\lim_{x' \uparrow x} p^{x'}_{\tau^Y_{x'}} = p_L^x$. And as $p^{x'}_{\tau^Y_{x'}} = \underline{\pi}$ for every x', this means that $p_L^x = \underline{\pi}$ and therefore that $\tau_x^Y \leq L$. So by monotonicity $\lim_{x' \uparrow x} \tau_{x'}^Y = \tau_x^Y$.

Next take $x' \downarrow x$. Given $p^x_{\tau^Y_x} = \underline{\pi}$, the non-zero quadratic variation of p^x and the strong Markov property imply that for every $\varepsilon > 0$ there exists a $t' \in (\tau_x^Y, \tau_x^Y + \varepsilon)$ such that $p_{t'}^x < \underline{\pi}$. Now, $p_{t'}^{x'} \to p_{t'}^x$ as $x' \to x$, so $\lim_{x' \downarrow x} \tau_{x'}^Y < t' < \tau_x^Y + \varepsilon$ for every ε. Thus $\lim_{x' \downarrow x} \tau_{x'}^Y = \tau_x^Y$ pointwise. This establishes that τ_x^Y is continuous in x a.s. The bounded convergence theorem then implies continuity of f.

C.13 Proof of Lemma 13

By definition,

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

for all $t \leq \tau^*$, where $\tau^* = \inf\{t \ : \ \pi_t \leq \underline{\pi}\}$. Given that $\pi_t = \mathbb{E}_t^Y[\mathbf{1}\{\Lambda > t\}]$, this may equivalently be written

$$\widetilde{V}(\pi_t) = \mathbb{E}_t^Y\left[\int_t^{t\vee(\tau^*\wedge\Lambda)} e^{-\rho(s-t)} r_G\, ds - \int_{t\vee(\tau^*\wedge\Lambda)}^{\tau^*} e^{-\rho(s-t)} b\, ds\right].$$

On $\{\tau^* \geq t\}$ this latter expression may be partitioned as

$$\begin{aligned}
&\mathbb{E}_t^Y\left[\int_t^{t\vee(\tau^*\wedge\Lambda)} e^{-\rho(s-t)} r_G\, ds - \int_{t\vee(\tau^*\wedge\Lambda)}^{\tau^*} e^{-\rho(s-t)} b\, ds\right] \\
&= \mathbb{E}_t^Y\left[\mathbf{1}\{\Lambda > t\}\left(\int_t^{\tau^*\wedge\Lambda} e^{-\rho(s-t)} r_G\, ds - \int_{\tau^*\wedge\Lambda}^{\tau^*} e^{-\rho(s-t)} b\, ds\right)\right] \\
&\quad - \mathbb{E}_t^Y\left[\mathbf{1}\{t \geq \Lambda\}\int_t^{\tau^*} e^{-\rho(s-t)} b\, ds\right].
\end{aligned}$$

Now, Lemma 8 and the argument in the proof of Proposition 4 imply that $V(F_t^*)$ may be written

$$V(F_t^*) = \mathbb{E}_t\left[\int_t^{\tau^*\wedge\Lambda} e^{-\rho(s-t)} r_G\, ds - e^{-\tau^*\wedge\Lambda} F^*_{\tau^*\wedge\Lambda}\right] = \mathbb{E}_t\left[\int_t^{\tau^*\wedge\Lambda} e^{-\rho(s-t)} r_G\, ds - \int_{\tau^*\wedge\Lambda}^{\tau^*} e^{-\rho(s-t)} b\, ds\right]$$

on $\{\tau^* \wedge \Lambda > t\}$. And on $\{\tau^* \geq t \geq \Lambda\}$,

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

Therefore by the law of iterated expectations

$$\begin{aligned}
\widetilde{V}(\pi_t) &= \mathbb{E}_t^Y\left[\mathbf{1}\{\Lambda > t\} V(F_t^*)\right] - \mathbb{E}_t^Y\left[\mathbf{1}\{t \geq \Lambda\} F_t^*\right] \\
&= \pi_t V(F_t^*) - (1-\pi_t) F_t^* \\
&= \pi_t V(f(\pi_t)) - (1-\pi_t) f(\pi_t)
\end{aligned}$$

on $\{\tau^* \geq t\}$. In particular, this holds whenever $\pi_t = x$ for each $x \in [\underline{\pi}, 1]$, proving the lemma.

D Technical appendix

This section contains technical lemmas used in the proofs of results appearing in the main text.

Lemma 20. *Fix a contract $\mathcal{C} = (F, \tau^Y)$. Suppose that for each t,*

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

whenever $\tau^Y \wedge \Lambda > t$. Then $\mathcal{C}$ satisfies IC-G.

Proof. Fix a contract $\mathcal{C} = (F, \tau^Y)$. IC-G holds iff

$$\mathbb{E}\left[\int_0^{\tau^Y \wedge \Lambda} e^{-\rho t} b\, dt + e^{-\rho(\tau^Y \wedge \Lambda)} F_{\tau^Y \wedge \Lambda}\right] \geq \mathbb{E}\left[\int_0^{\tau^Y \wedge \sigma} e^{-\rho t} b\, dt + e^{-\rho(\tau^Y \wedge \sigma)} F_{\tau^Y \wedge \sigma}\right]$$

for every $\mathbb{F}$-stopping time $\sigma \leq \Lambda$. This inequality may be rearranged to obtain

$$\mathbb{E}\left[\mathbf{1}\{\tau^Y \wedge \Lambda > \sigma\} e^{-\rho\sigma}\left(\int_\sigma^{\tau^Y \wedge \Lambda} e^{-\rho(t-\sigma)} b\, dt + e^{-\rho(\tau^Y \wedge \Lambda - \sigma)} F_{\tau^Y \wedge \Lambda} - F_\sigma\right)\right] \geq 0.$$

Then by applying the law of iterated expectations, IC-G is equivalent to the condition that

$$\mathbb{E}\left[\mathbf{1}\{\tau^Y \wedge \Lambda > \sigma\} e^{-\rho\sigma}\left(\mathbb{E}_\sigma\left[\int_\sigma^{\tau^Y \wedge \Lambda} e^{-\rho(t-\sigma)} b\, dt + e^{-\rho(\tau^Y \wedge \Lambda - \sigma)} F_{\tau^Y \wedge \Lambda}\right] - F_\sigma\right)\right] \geq 0.$$

for every $\sigma \leq \Lambda$. Note that on for every time t, on the event $\{\sigma = t\}$ the expression inside the expectation is equal to

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

which by assumption is non-negative. As this holds for every choice of t, it must be that

$$\mathbf{1}\{\tau^Y \wedge \Lambda > \sigma\} e^{-\rho\sigma}\left(\mathbb{E}_\sigma\left[\int_\sigma^{\tau^Y \wedge \Lambda} e^{-\rho(t-\sigma)} b\, dt + e^{-\rho(\tau^Y \wedge \Lambda - \sigma)} F_{\tau^Y \wedge \Lambda}\right] - F_\sigma\right) \geq 0$$

for every $\sigma \leq \Lambda$. Hence $\mathcal{C}$ satisfies IC-G. □

Lemma 21. *Fix a contract $\mathcal{C} = (F, \tau^Y)$. Suppose there exist progressively measurable processes γ and β such that for all t, each of the following holds a.e. on $\{\tau^Y > t\}$:*

- $F_t = F_0 + \int_0^t \gamma_s \, ds + \int_0^t \beta_s \, dZ_s^G$,
- $b - \rho F_t + \gamma_t - \frac{\Delta r}{\sigma} \beta_t = 0$,
- $\mathbb{E}^G \left[\int_0^t \beta_s^2 \, ds \right] < \infty$,
- $\beta_t \geq 0$.

Then $\mathcal{C}$ is an IC contract.

Proof. The proof of Lemma 9 establishes that $\mathcal{C}$ is IC-B, as right-continuity of γ and β are unnecessary for the sufficiency proof of that lemma. I check IC-G using the sufficiency condition of Lemma 20 in the technical appendix. Define a process U by

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

when $\tau^Y \wedge \Lambda > t$, and $U_t = U_{\tau^Y \wedge \Lambda}$ otherwise. Then IC-G obtains if $U_t \geq F_t$ whenever $\tau^Y \wedge \Lambda > t$.

When $\tau^Y \wedge \Lambda > t$, U_t may be written

$$U_t = \int_t^\infty du \, \alpha e^{-\alpha(u-t)} \mathbb{E}_t^G \left[\int_t^{u \wedge \tau^Y} e^{-\rho(s-t)} b \, ds + e^{-\rho(u \wedge \tau^Y - t)} F_{u \wedge \tau^Y} \right].$$

Use Ito's lemma to expand the final term as

$$e^{-\rho(u \wedge \tau^Y - t)} F_{u \wedge \tau^Y} = F_t + \int_t^{u \wedge \tau^Y} e^{-\rho(s-t)} (\gamma_s - \rho F_s) \, ds + \int_t^{u \wedge \tau^Y} e^{-\rho(s-t)} \beta_s \, dZ_s^G.$$

Due to the regularity condition imposed on β in the lemma statement, the final term is a $\mathbb{P}^G$-martingale. Thus U_t reduces to

$$U_t = F_t + \int_t^\infty du \, \alpha e^{-\alpha(u-t)} \mathbb{E}_t^G \left[\int_t^{u \wedge \tau^Y} e^{-\rho(s-t)} (\gamma_s + b - \rho F_s) \, ds \right].$$

And for each $s \in [t, \tau^Y)$, $\gamma_s + b - \rho F_s = \frac{\Delta r}{\sigma} \beta_s \geq 0$ a.s. So the final term is non-negative, yielding $U_t \geq F_t$. □

Lemma 22. *Fix a contract $\mathcal{C} = (F, \tau^Y)$ and define the $\mathbb{F}^Y$-adapted process M by*

$$M_t = \frac{b}{\rho}(1 - e^{-\rho t}) + e^{-\rho t} F_t.$$

Then $\mathcal{C}$ is an IC-B contract iff M^{τ^Y} is a $\mathbb{P}^B$-supermartingale.

Proof. Fix any t and $\mathbb{F}^Y$-stopping time $\tau' \geq t$, and consider states of the world in which $\tau^Y > t$. Then the definition of M implies

$$\begin{aligned}\mathbb{E}_t^B[M_{\tau^Y \wedge \tau'}] &= \mathbb{E}_t^B\left[\frac{b}{\rho}(1 - e^{-\rho(\tau^Y \wedge \tau')}) + e^{-\rho(\tau^Y \wedge \tau')} F_{\tau^Y \wedge \tau'}\right] \\ &= \frac{b}{\rho}(1 - e^{-\rho t}) + e^{-\rho t}\mathbb{E}_t^B\left[\frac{b}{\rho}(1 - e^{-\rho(\tau^Y \wedge \tau' - t)}) + e^{-\rho(\tau^Y \wedge \tau' - t)} F_{\tau^Y \wedge \tau'}\right].\end{aligned}$$

Now, suppose $\mathcal{C}$ is an IC-B contract. Then by Lemma 5 the final term is at most F_t, meaning

$$\mathbb{E}_t^B[M_{\tau^Y \wedge \tau'}] \leq \frac{b}{\rho}(1 - e^{-\rho t}) + e^{-\rho t} F_t = M_t,$$

and in particular this holds when $\tau' = s > t$. Thus M^{τ^Y} is a supermartingale.

On the other hand, suppose M^{τ^Y} is a $\mathbb{P}^B$-supermartingale. Then $\mathbb{E}_t^B[M_{\tau^Y \wedge \tau'}]$ is at most $M_t^{\tau^Y} = M_t$, implying

$$F_t \geq \mathbb{E}_t^B\left[\frac{b}{\rho}(1 - e^{-\rho(\tau^Y \wedge \tau' - t)}) + e^{-\rho(\tau^Y \wedge \tau' - t)} F_{\tau^Y \wedge \tau'}\right].$$

Thus by Lemma 5 $\mathcal{C}$ is an IC-B contract. □

Lemma 23. *Fix an IC-B contract $\mathcal{C} = (F, \tau^Y)$ satisfying $\tau^Y = \inf\{t : F_t = 0\}$ and $F \leq b/\rho$. Then there exists another IC-B contract $\mathcal{C}' = (F', \tau')$ which is cadlag and satisfies $F' \leq F$ and $\tau' = \tau^Y$ a.s. Thus $\mathcal{C}'$ yields weakly higher expected profits than $\mathcal{C}$.*

Proof. Fix any IC-B contract (F, τ^Y) satisfying $F \leq b/\rho$. I first establish that F^{τ^Y} is a $\mathbb{P}^B$-supermartingale. Let

$$M_t = \frac{b}{\rho}(1 - e^{-\rho t}) + e^{-\rho t} F_t.$$

By Lemma 22, M^{τ^Y} is a $\mathbb{P}^B$-supermartingale. Thus for all t such that $\tau^Y > t$ and all $s > t$,

$$\mathbb{E}_t^B[F_s^{\tau^Y}] = \mathbb{E}_t^B\left[e^{\rho(s \wedge \tau^Y)} M_s^{\tau^Y} - \frac{b}{\rho}(e^{\rho(s \wedge \tau^Y)} - 1)\right] = \mathbb{E}_t^B\left[(e^{\rho(s \wedge \tau^Y)} - 1)\left(M_s^{\tau^Y} - \frac{b}{\rho}\right)\right] + \mathbb{E}_t^B[M_s^{\tau^Y}].$$

As $F \leq b/\rho$ by assumption, $M \leq b/\rho$ as well. This establishes the bound

$$(e^{\rho(s\wedge\tau^Y)} - 1)\left(M_s^{\tau^Y} - \frac{b}{\rho}\right) \leq (e^{\rho t} - 1)\left(M_s^{\tau^Y} - \frac{b}{\rho}\right).$$

Hence

$$\mathbb{E}_t^B[F_s^{\tau^Y}] \leq -\frac{b}{\rho}(e^{\rho t} - 1) + e^{\rho t}\mathbb{E}_t^B[M_s^{\tau^Y}] \leq -\frac{b}{\rho}(e^{\rho t} - 1) + e^{\rho t} M_t^{\tau^Y} = F_t^{\tau^Y}.$$

Thus F^{τ^Y} is a $\mathbb{P}^B$-supermartingale.

Now I assume wlog that $F = F^{\tau^Y}$, if necessary passing to the stopped process, so that F is a $\mathbb{P}^B$-supermartingale. Then by Proposition 1.3.14 of Karatzas and Shreve (1991), F possesses pathwise right limits for all time almost surely, and the process F' constructed by setting $F'_t(\omega) = \lim_{s\downarrow t} F_t(\omega)$ for each $\omega \in \Omega$ is a $\mathbb{P}^B-$supermartingale wrt to the filtration $\{\mathcal{F}_{t+}^Y\}_{t\geq 0}$. As $\mathbb{F}^Y$ is right-continuous under the usual conditions, this means F' is a $\mathbb{P}^B-$supermartingale wrt $\mathbb{F}^Y$. Further, the proposition says that $F'_t \leq F_t$ for all time a.s.

Suppose additionally that $\tau^Y = \inf\{t : F_t = 0\}$, and define $\tau' \equiv \inf\{t : F'_t = 0\}$. I claim that $\mathcal{C}' = (F', \tau')$ satisfies the conditions of the lemma statement. First note that as F is non-negative, so is F'. So F' is a feasible payment process. Let $\tau' \equiv \inf\{t : F'_t = 0\}$. I claim that $\tau' = \tau^Y$ a.s. Given $F' \leq F$, clearly $\tau^Y \leq \tau'$. In the other direction, suppose that for some t and in some state of the world $\tau' = t$, so that in particular $\lim_{s\downarrow t} F_s = 0$. Given that F is a $\mathbb{P}^B$-supermartingale, the inequality $F_t \leq \mathbb{E}^B[F_s]$ holds for all $s > t$. As F is uniformly bounded, the bounded convergence theorem then implies that $F_t \leq \mathbb{E}^B[\lim_{s\downarrow t} F_s] = 0$. Thus $\tau^Y \leq t$, as desired.

From $F' \leq F$ and $\tau' = \tau^Y$ it follows that $\mathcal{C}'$ pays the agent weakly less than (F, τ^Y) while operating the project for the same length of time in the good state in all cases. Thus $\mathcal{C}'$ is weakly more profitable for the firm, as claimed. It remains only to show that $\mathcal{C}'$ is IC-B, which I verify using the characterization of Lemma 5.

Whenever $F'_t = F_t$ IC-B is automatic, as the agent's profits from any delayed reporting strategy are weakly lower under F'. So consider any time t such that $F_t > F'_t$. For every $s > t$ and $\mathbb{F}^Y$-stopping time $\tau'' \geq s$, incentive-compatibility of F implies that

$$F_s \geq \mathbb{E}_s^B\left[\frac{b}{\rho}(1 - e^{-\rho((\tau^Y\wedge\tau'')-s)}) + e^{-\rho((\tau^Y\wedge\tau'')-s)} F_{\tau^Y\wedge\tau''}\right].$$

Taking expectations wrt $\mathbb{E}_t^B$, multiplying through by $e^{-\rho(s-t)}$, and taking the supremum over all $\tau'' \geq s$ gives

$$\mathbb{E}_t^B[e^{-\rho(s-t)}F_s] \geq \sup_{\tau''\geq s} \mathbb{E}_t^B\left[\frac{b}{\rho}(e^{-\rho(s-t)} - e^{-\rho((\tau^Y\wedge\tau'')-t)}) + e^{-\rho((\tau^Y\wedge\tau'')-t)}F_{\tau^Y\wedge\tau''}\right].$$

Rearranging yields

$$\mathbb{E}_t^B\left[\frac{b}{\rho}(1-e^{-\rho(s-t)}) + e^{-\rho(s-t)}F_s\right] \geq \sup_{\tau''\geq s} \mathbb{E}_t^B\left[\frac{b}{\rho}(1 - e^{-\rho((\tau^Y\wedge\tau'')-t)}) + e^{-\rho((\tau^Y\wedge\tau'')-t)}F_{\tau^Y\wedge\tau''}\right].$$

Now take $s \downarrow t$. The rhs is increasing in s, so the limit of the rhs exists and is equal to the supremum taken over all $\tau' > t$. Meanwhile the interior of the expectation on the lhs is uniformly bounded by b/ρ, so the bounded convergence theorem allows the limit and expectation to be swapped. As $\lim_{s\downarrow t} F_s = F'_t$ by definition, the resulting inequality is

$$F'_t \geq \sup_{\tau''>t} \mathbb{E}_t^B\left[\frac{b}{\rho}(1 - e^{-\rho((\tau^Y\wedge\tau'')-t)}) + e^{-\rho((\tau^Y\wedge\tau'')-t)}F_{\tau^Y\wedge\tau''}\right].$$

As $F \geq F'$ and $\tau^Y = \tau'$ a.s., the inequality is preserved when F is replaced by F' and τ^Y is replaced by τ' on the rhs, yielding

$$F'_t \geq \sup_{\tau''>t} \mathbb{E}_t^B\left[\frac{b}{\rho}(1 - e^{-\rho((\tau'\wedge\tau'')-t)}) + e^{-\rho((\tau'\wedge\tau'')-t)}F'_{\tau'\wedge\tau''}\right].$$

Finally, the expected payoff of any stopping time $\tau'' \geq t$ is equal to a weighted average of F'_t and the expected payoff from some $\tau'' > t$. As no stopping time $\tau'' > t$ can improve on F'_t, neither can a stopping time $\tau'' \geq t$. It follows that $\mathcal{C}'$ satisfies IC-B. □

Remark. *The cadlag golden parachute F' is* not *necessarily a modification of F in the technical sense. That is, there may exist times t for which $F_t \neq F'_t$ with strictly positive probability. However, Lemma 23 establishes that any such disagreement preserves incentive-compatibility and weakly improves profitability of the contract.*

Lemma 24. *Suppose that for some $F^* > 0$, $V : [0, F^*] \to \mathbb{R}$ is a C^2 function satisfying $V'(0) = 0$, $V'(F^*) = 0$, and the HJB equation on $[0, F^*]$. Then:*

- *$V'(F_0) > 0$ and $V''(F_0) < 0$ for all $F_0 \in [0, F^*)$,*
- *The unique maximizers $(\gamma^*(F_0), \beta^*(F_0))$ of the HJB equation for each $F_0 \in [0, F^*)$ are*

$$\beta^*(F_0) = -\frac{\Delta r}{\sigma}\frac{V'(F_0)}{V''(F_0)}, \quad \gamma^*(F_0) = \beta^*(F_0)\frac{\Delta r}{\sigma} + \rho F - b,$$

- $(\gamma^*(F^*), \beta^*(F^*)) = (\rho F^* - b, 0)$ *are maximizers of the HJB equation for* $F_0 = F^*$.

Proof. Fix $F_0 \in [0, F^*]$ and suppose $V'(F_0) < 0$. Then γ may be taken arbitrarily large and negative without violating IC, so that the rhs of the HJB equation equals ∞. However, the lhs of the HJB equation is finite, a contradiction. So $V'(F_0) \geq 0$. An analogous argument rules out $V''(F_0) > 0$ by considering arbitrarily large and positive β. So both V' and $-V''$ must be weakly positive. Next suppose that $V'(F_0) > 0$ and $V''(F_0) = 0$. Then again the rhs of the HJB equation may be made arbitrarily large by taking both γ and β large and positive. So either $V'(F_0) > 0$ and $V''(F_0) < 0$, or else $V'(F_0) = 0$ and $V''(F_0) \leq 0$. In particular, V is weakly increasing on $[0, F^*]$.

Suppose that $V'(F_0) = 0$ and $V''(F) \leq 0$ for some $F_0 \in [0, F^*)$. Then $\beta^*(F_0)V''(F) = 0$ for any maximizer $\beta^*(F_0)$, in which case the HJB equation implies $(\rho + \alpha)V(F_0) = r_G - \alpha F$. But at F^* the same argument applies, so that $(\rho+\alpha)V(F^*) = r_G - \alpha F^* < r_G - \alpha F_0 = \rho V(F_0)$. This contradicts the fact that V is weakly increasing on $[0, F^*]$, so it must be that $V'(F_0) > 0$ and $V''(F_0) < 0$ for all $F_0 \in [0, F^*)$.

The stated maximizers of the HJB equation for $F_0 \in [0, F^*)$ follow from noting that the IC constraint must bind, else increasing γ a sufficiently small amount would increase the rhs of the HJB equation while preserving the IC constraint given $V'(F_0) > 0$. Thus $\gamma = \frac{\Delta r}{\sigma}\beta + \rho F_0 - b$ at the optimum, and substituting this expression into the rhs of the HJB equation yields a concave quadratic in β with unique maximum $\beta = -\frac{\Delta r}{\sigma}\frac{V'(F_0)}{V''(F_0)}$. And as $V''(F^*) \leq 0$ by continuity of V'' while $V'(F^*) = 0$, the HJB equation at F^* is maximized by $(\gamma, \beta) = (\rho F^* - b, 0) \in \mathrm{IC}(F^*)$. □