Leaks, Sabotage, and Information Design*

Aaron Kolb† Erik Madsen‡

February 2, 2019

Abstract

We study optimal dynamic information disclosure by a principal to an agent of uncertain loyalty who may engage in hidden undermining, for instance through damaging leaks or sabotage. The agent requires information to correctly perform a task but may also covertly commit destructive acts which are only stochastically detectable. The principal optimally provides inconclusive incremental guidance until a deterministic time when the agent is deemed trusted and given a conclusive final report. Disloyal agents are never given incentives to feign loyalty, and in the unique time-consistent implementation undermine with variable, non-monotonic intensity over the lifetime of employment.

JEL Classification: C70, D82, D83, D86, M51

Keywords: information leaks, sabotage, principal-agent model, information design

1 Introduction

An organization has found itself the victim of information leaks and sabotage. Sensitive documents have been leaked to the media, corporate secrets have been sold to competitors, obscure vulnerable points in production lines have been discovered and sabotaged. An insider with access to privileged information must be undermining the organization — but who? Halting the distribution of sensitive data would staunch the bleeding, but also leave employees paralyzed and unable to act effectively. Limited information could be circulated

*The authors thank Laurent Mathevet and seminar audiences at Brown University and the 2018 NSF/NBER/CEME Conference at the University of Chicago for helpful conversations.

†Department of Business Economics and Public Policy, Kelley School of Business, Indiana University (Email: kolba@indiana.edu).

‡Department of Economics, New York University (Email: emadsen@nyu.edu).

to goad the perpetrator into further incriminating acts, but he may simply lie low and feign loyalty until business as usual resumes. How should the organization regulate information flow to optimally assess employee loyalty and limit damage to operations?

Abuse of organizational secrets is a common concern in a broad range of organizations. Leaks are one important instance — for instance, of governmental or military secrets to the enemy during times of war, of damaging or sensitive information about government policies to the press, and of technology or strategic plans to rival firms. Another is straightforward sabotage of sensitive assets — for instance, destruction of a key railway line transporting troops and supplies for an upcoming operation, or malicious re-writing of software code causing defects or breakdowns on a production line. (We discuss several prominent examples of such behavior in Section 1.1.) We take these occurrences as evidence that disloyal agents occasionally seek to undermine their organization through abuse of privileged information.

An important feature of the examples discussed above is that undermining is typically a dynamic process whose source can be identified only with substantial delay. A mole within an organization typically leaks many secrets over a long period, during which time the organization may realize that information is leaking, but not know who is doing the leaking. Similarly, a saboteur may disrupt a production line multiple times or bomb several installations before being caught. (In Section 1.1 we highlight the ongoing nature of undermining activities in several of the examples discussed there.) We therefore study environments in which agents have many chances to undermine their organization, with each decision to undermine being only occasionally discovered and traced back to the agent.

In our model, a principal hires an agent to perform a task, consisting of a choice between one of several actions, repeatedly over time. The optimal task choice depends on a persistent state variable initially known only to the principal, who may commit to disclose information about the state according to any dynamic policy. Disclosing information improves the agents task performance. However, the agent may secretly choose to undermine the principal at any point in time. The damage inflicted by each such act is increasing in the precision of the agent's information about the state of the world. Further, each act of undermining is detected only probabilistically. The agent's loyalty is private information, and a disloyal agent wishes to inflict maximum damage to the principal over the lifetime of his employment. The principal's goal is to regulate information flow so as to aid loyal agents, while limiting the damage inflicted by disloyal agents before they are uncovered.

We characterize the principal's optimal dynamic disclosure rule, which has several important qualitative features. First, the principal should optimally set loyalty tests — that is, it should design information flow to encourage disloyal agents to undermine at all times. In particular, any disclosure policy giving a disloyal agent strict incentives to feign loyalty at any point in time is suboptimal. Second, the principal should divide agents into two classes, trusted and untrusted, with trusted agents given perfect knowledge of the state while untrusted agents are left with uncertainty about the state which is bounded away from zero. Agents should be promoted from the untrusted pool to the trusted circle only following a trial period of set length during which they have not been caught undermining. Third, untrusted agents should first be treated to a quiet period in which no information is shared, and then be gradually informed about the state via periodic unreliable reports which sometimes deliberately mislead the agent. Finally, in the unique time-consistent implementation of the optimal policy, a disloyal agent may follow a variable-intensity policy with non-monotonic intensity over the lifetime of employment. In particular, he undermines at maximum intensity during the quiet period, then becomes more cautious and undermines with an interior intensity which increases until he achieves trusted status, following which he undermines with certainty forever. Such an undermining strategy ensures that at each instant, the agent's reputation justifies the level of information the principal has promised him.

The design of the optimal disclosure rule is crucially shaped by the agent's ability to strategically time undermining. Ideally, the principal would like to subject each agent to a low-stakes trial period, during which disloyal agents are rooted out with high enough certainty that the principal can comfortably trust the survivors. However, disloyal agents will look ahead to the end of the trial phase and decide to feign loyalty until they become trusted, so such a scheme will fail to effectively identify loyalists. Indeed, an ineffective trial period merely delays the time until the principal begins building trust, reducing lifetime payoffs from the project. Agents must therefore be given enough information prior to becoming trusted that they prefer to undermine during the trial period rather than waiting to build trust. A gradual rise in the agent's knowledge of the state via unreliable reports emerges as the cost-minimizing way to ensure disloyal agents choose to undermine at all times during the trial period.

We also show how to implement the principal's optimal disclosure rule via a sequence of noisy signals about the state. In particular, in our benchmark model with a binary state space, the principal optimally informs the agent about the state via a combination of binary signals and an inconclusive disconfirming news process.[1] Both the binary signals and the news reports are unreliable, meaning that the principal occasionally provides deliberately incorrect information about the state to the agent to constrain his knowledge. Meanwhile in

[1] A disconfirming news process informs the agent via Poisson signals that arrive at a higher rate when the agent's conjecture about the most likely state is incorrect. The process is inconclusive if signals still arrive at some rate when the agent's conjecture is correct.

an alternative specification with a continuous state space, the optimal disclosure rule can be implemented via a Brownian signal process which reports the state at each moment in time, garbled by noise whose magnitude varies over time. As in the binary state case, unreliable reports of the state are crucial to properly controlling the agent's knowledge of the true state.

The remainder of the paper is organized as follows. Section 1.1 presents real world examples of leaks and sabotage, while Section 1.2 reviews related literature. Section 2 presents the model. We discuss the solution of the model in Section 3. Section 4 interprets the model's main features in the context of the application. Section 5 discusses comparative statics, Section 6 shows robustness of the main qualitative results to modeling changes along various dimensions, and Section 7 concludes. All proofs are relegated to an appendix.

1.1 Evidence of leaks and sabotage

Wartime provides a classic setting in which leaks, sabotage, and undermining are ubiquitous. These acts occur in military, business, and political contexts alike, and many historical examples have been well-documented. We highlight a few representative anecdotes occurring the Second World War. The United States fell victim to a severe information leak when Communist sympathizers Ethel and Julius Rosenberg supplied Soviet intelligence with confidential documents about American nuclear research, allowing the Soviet Union to quickly advance its nuclear program and perform testing by 1949. The Rosenbergs passed classified documents to Soviet handlers for years before they were eventually discovered, tried and convicted in one of the most high profile espionage cases in history, and ultimately executed (Ziegler and Jacobson 1995; Haynes and Klehr 2006). In turn, the United States aided and abetted acts of sabotage by sympathizers in belligerent nations. In the now declassified 1944 brief "Simple Sabotage Field Manual," the Office of Strategic Services (the predecessor to the CIA), provided detailed suggestions to sympatheizers on how to engage in destructive acts over time in both military and industrial contexts: "Slashing tires, draining fuel tanks, starting fires, starting arguments, acting stupidly, short-circuiting electrical systems, [and] abrading machine parts will waste materials, manpower, and time." The brief also provided advice on how to minimize suspicion and risk of detection: "The potential saboteur should discover what types of faulty decisions and non-cooperation are normally found in his kind of work and should then devise his sabotage so as to enlarge that 'margin of error.' . . . If you commit sabotage on your job, you should naturally stay at your work." This advice contemplated continued and repeated acts of sabotage, often relying on intimate knowledge of organizational operations obtained trust and experience.

Leaks and sabotage are also common during peacetime, and are reported regularly in the news. We highlight several significant contemporary incidents. Technology firms have lately been popular targets for leaks by disgruntled employees. In the lead-up to the 2016 election, Facebook attracted negative attention when one of its employees leaked information via laptop screenshots to a journalist at the technology news website Gizmodo (Thompson and Vogelstein 2018). On one occasion, the screenshot involved an internal memo from Mark Zuckerberg admonishing employees about appropriate speech. On a second occasion, it showed results of an internal poll indicating that a top question on Facebook employees' minds was about the company's responsibility in preventing Donald Trump from becoming president. Other articles followed shortly after, citing inside sources, about Facebook's suppression of conservative news and other manipulation of headlines in its Trending Topics section by way of "injecting" some stories and "blacklisting" others (Isaac 2016).

Tesla has fallen victim to both information leaks and deliberate acts of sabotage. In June 2018, CEO Elon Musk sent an internal email to employees about a saboteur who had been caught "making direct code changes to the Tesla Manufacturing Operating System under false usernames and exporting large amounts of highly sensitive Tesla data to unknown third parties." In this case, the saboteur was a disgruntled employee who was passed over for a promotion, but Musk warned employees of a broader problem: "As you know, there are a long list of organizations that want Tesla to die. These include Wall Street short-sellers. Then there are the oil & gas companies. . ." (Kolodny 2018).

Meanwhile in politics, the Trump administration has suffered a series of high-profile leaks of sensitive or damaging information. These leaks include reports on Russian hacking in the United States election; discussions about sanctions with Russia; memos from former FBI director James Comey; transcripts of telephone conversations between Trump and Mexican President Peña Nieto regarding financing of a border wall (Kinery 2017); a timeline for military withdrawal from Syria; and a characterization of some third-world nations as "shithole countries" (Rothschild 2018). Past administrations have also suffered serious leaks, for instance the disclosure of classified documents by Edward Snowden and Chelsea Manning under the Obama administration.

In response to damaging leaks,[2] the Trump administration has reportedly begun circulating fabricated stories within the White House in order to identify leakers (Cranley 2018). Elon Musk has also reportedly used this technique in an attempt to identify leakers. The practice is known in the intelligence community as a "barium meal test" or a "canary trap".[3] As a wartime example, British intelligence in the Second World War planted false documents

[2] Trump has said on Twitter: "Leakers are traitors and cowards, and we will find out who they are!"

[3] Relatedly, Amazon strategically plants dummy packages in its delivery trucks to identify drivers who steal packages which deliver error messages (Peterson 2018), and the company is using dummy packages with GPS devices to help local police departments prevent theft of delivered packages (Associated Press 2018).

on a corpse discovered by nominally neutral Spain and leaked to Germany, indicating that the Allies intended to invade Greece and Sardinia. This deception campaign, dubbed Operation Mincemeat, diverted German forces and attention away from a successful invasion of Sicily (Macintyre 2010). Both examples indicate the important role played by sharing incomplete and sometimes intentionally incorrect information.

As the examples above illustrate, undermining occurs in a broad range of settings. Often, these acts have an ongoing nature and deal damage repeatedly over time. Leaders in military, business, and political settings are acutely aware of the risk of such activities, and in response formulate defensive strategies which may include circulation of dubious information.

1.2 Related Literature

Our paper contributes to the growing literature on information design sparked by Kamenica and Gentzkow (2011). Our model features persuasion in a dynamic environment, with a privately informed receiver taking actions repeatedly over time subject to imperfect monitoring by the sender. A small set of recent papers, notably Ely, Frankel, and Kamenica (2015), Ely (2017), Ely and Szydlowski (2018), Ball (2018), and Orlov, Skrzypacz, and Zryumov (2018), have studied models of dynamic persuasion without private receiver information. Another emerging strand of the literature has studied the effects of combining persuasion with a privately informed receiver. Kolotilin et al. (2017), Kolotilin (2018), and Guo and Shmaya (2018) study static persuasion problems. In each paper, the receiver makes a binary decision to act or not, and the focus is on settings in which the sender has a (possibly stochastic) bias toward action relative to the receiver. Meanwhile Au (2015) and Basu (2018) involve dynamic persuasion. In both papers the receiver chooses the timing of a single public game-ending action, and the sender's preferences are independent of both the state and the receiver's information. To the best of our knowledge, ours is the first paper studying dynamic persuasion of a privately informed receiver who acts repeatedly over time. And none of the papers cited above feature imperfect monitoring of receiver actions.

Our paper is also related to a set of papers studying long-term relationships with variable stakes where one party can betray the other or expropriate gains, notably Watson (2002), Thomas and Worrall (1994), Albuquerque and Hopenhayn (2004), Rayo and Garicano (2017), and Fudenberg and Rayo (2018). In each of these papers the focus is either on the timing of a single exogenously relationship-ending act, or (as in Thomas and Worrall (1994)) a publicly observed act which is punished by permanent reversion to autarky in equilibrium. In contrast, our model focuses on *recurring* undermining which is inflicted repeatedly and observed only stochastically. We are unaware of any other work which studies dynamic moral hazard of this sort in the context of a variable-stakes relationship. Further, we microfound the stakes of a relationship as stemming from disclosure of information about a payoff-relevant state, in contrast to the typical interpretation of the stakes curve by the papers cited as physical or human capital. This interpretation focuses our model on a different set of applications, and allows us to investigate how a desired stakes curve is implemented through dynamic signaling.

Most of the papers in this group also abstract from screening problems by assuming the preferences and outside options of all parties are known. Watson (2002), the most closely related paper to ours, is the only paper in the group to feature private information about preferences, and in particular uncertainty over whether each agent benefits from cooperation.[4] The one-time nature of betrayal in that paper yields very different predictions about optimal stakes curves, notably no quiet periods and no discrete jumps in stakes after time zero. Further, in that model optimal equilibria exhibit betrayal at exactly two points in time — at time zero and when stakes have just reached their maximum level. By contrast, in our model the time-consistent undermining path features a positive probability of undermining at all times and richer undermining dynamics.

2 The model

This section presents the model and defines the associated optimal contracting problem. To clarify exposition and streamline the flow of the paper, our initial presentation of the model omits commentary on its various components. In Section 4 we provide detailed motivations and microfoundations for each modeling assumption.

2.1 The environment

A (female) principal hires a (male) agent to perform a task over a potentially infinite horizon in continuous time. The payoff to this task depends on the agent's knowledge of a persistent binary state $\omega \in \{L, R\}$.[5] Let $\mu_t = \Pr_t(\omega = R)$ be the agent's time-t beliefs about the state. The learning process generating the posterior belief process μ is detailed in Section 2.2; for the moment it may be taken to be an arbitrary martingale. The better informed the agent is about the state, the more effectively he performs the task. More precisely, when the agent's time-t beliefs are $\mu_t \in [0, 1]$, he provides an expected flow payoff of $|2\mu_t - 1|$ to the principal

[4] The companion paper Watson (1999) studies the same model but focuses on design of renegotiation-proof stakes curves, so is less closely related to our analysis.

[5] We consider alternative specifications of the state space in Sections 6.2 and 6.3.

from performing the task at time t. Performance of this task is perfectly monitored and not subject to moral hazard.

In addition to performing the perfectly monitored task, the agent may secretly take actions to undermine the principal. Specifically, at each moment in time the agent decides whether to undermine by choosing $b_t \in \{0, 1\}$ in addition to performing their publicly observed task.[6] Whenever $b_t = 1$ the agent takes actions to undermine the principal, and when the agent's posterior beliefs are μ_t he inflicts an expected flow loss of $K|2\mu_t - 1|$ on the principal, where $K > 1$. The principal discounts payoffs at rate $r > 0$, and so given an undermining action profile b and a belief process μ her expected payoffs from employing the agent until the (possibly random or infinite) time T are

$$\Pi(b, \mu, T) = \mathbb{E} \int_0^T e^{-rt}(1 - Kb_t)|2\mu_t - 1| \, dt.$$

Meanwhile, the agent's motives depend on a preference parameter $\theta \in \{G, B\}$. When $\theta = G$, the agent is *loyal*. A loyal agent has preferences which are perfectly aligned with the principal's. That is, given an undermining action profile b and a belief process μ, the loyal agent's expected payoff from being employed until time T is $\Pi(b, \mu, T)$. On the other hand, when $\theta = B$ the agent is *disloyal*. The disloyal agent has interests totally opposed to the principal's, and his expected payoff is $-\Pi(b, \mu, T)$. Neither agent type incurs any direct costs of task performance or undermining; they care only about the payoffs the principal receives. If the agent is terminated or does not accept employment, he receives an outside option normalized to 0.

2.2 The information structure

Each player possesses private information about a portion of the environment. Prior to accepting employment with the principal, the agent is privately informed of his type θ. The principal believes that $\theta = G$ with probability $q \in (0, 1)$. Meanwhile ω is initially unobserved by either party, each of whom assigns probability $1/2$ that $\omega = R$. Upon hiring the agent, the principal becomes perfectly informed of ω while the agent receives a public signal $s \in \{L, R\}$ of ω which is correct with probability $\rho \geq 1/2$; the case $\rho = 1/2$ corresponds to a setting in which the agent receives no exogenous information about ω.

Once the agent is employed, the principal perfectly observes the agent's task performance but only imperfectly observes whether the agent has undermined. In particular, whenever the agent chooses $b_t = 1$ over a time interval dt the principal immediately receives definitive

[6] Formally, the agent chooses a Lebesgue-measurable undermining path b.

confirmation of this fact with probability $\gamma\, dt$, where $\gamma > 0$. Otherwise the act of undermining goes permanently undetected.

If the agent chooses action sequence b, the cumulative probability that his undermining has gone undetected by time t is therefore $\exp\left(-\gamma \int_0^t b_s\, ds\right)$. Note that the detection rate is *not* cumulative in past undermining — the principal has one chance to detect an act of undermining at the time it occurs. To ensure consistency with this information structure, we assume the principal does not observe her ex post payoffs until the end of the game.

The agent receives no further exogenous news about ω after observing the signal s. However, the principal has access to a disclosure technology allowing her to send noisy public signals of the state to the agent over time. This technology allows the principal to commit to signal structures inducing arbitrary posterior belief processes about the state. Formally, in line with the Bayesian persuasion literature, the principal may choose any $[0, 1]$-valued martingale process μ, where μ_t represents the agent's time-t posterior beliefs that $\omega = R$. The only restriction on μ is that $\mathbb{E}[\mu_0]$ be equal to the agent's posterior beliefs after receipt of the signal s. We will refer to any such process μ as an *information policy.*

2.3 Contracts

The principal hires the agent by offering a menu of contracts committing to employment terms over the duration of the relationship. Each contract specifies an information policy as well as a termination policy and a recommended unermining action profile for the agent. No transfers are permitted.[7] Without loss, the principal commits to fire the agent the first time she observes an act of undermining.[8] Further, it is not actually necessary to specify a recommended undermining policy. This is because for each agent type, all undermining policies maximizing the agent's payoff yield the same principal payoff. Thus a contract can be completely described by an information policy, i.e. a martingale belief process μ for the agent's posterior beliefs. Note that every contract automatically satisfies the participation constraint, as there exists an undermining policy guaranteeing the disloyal agent a nonnegative payoff.

In our model the agent possesses private information about his preferences prior to accepting a contract. As a result, as is well-recognized in the mechanism design literature, the principal might benefit by offering distinct contracts targeted at each agent type. However, it turns out that in our setting this is never necessary. As the following lemma demonstrates,

[7]See Section 6.1 for a discussion of how transfers affect the analysis.

[8]We assume for now that the principal does not fire the agent before he has been revealed to be disloyal. We discuss the possibility of pre-emptive firing in section 3.6.

no menu of contracts can outperform a single contract accepted by both agent types.[9]

Lemma 1. *Given any menu of contracts* (μ^1, μ^2), *there exists an* $i \in \{1, 2\}$ *such that offering* μ^i *alone weakly increases the principal's payoffs.*

The intuition for this result is simple and relies on the zero-sum nature of the interaction between the principal and a disloyal agent. If the principal offers a menu of contracts which successfully screens the disloyal type into a different contract from the loyal one, she must be increasing the disloyal agent's payoff versus forcing him to accept the same contract as the one taken by the loyal agent. But increasing the disloyal agent's payoff decreases the principal's payoff, so it is always better to simply offer one contract and pool the disloyal agent with the loyal one; the principal must then use information to screen dynamically.

In light of the developments in this subsection, the principal's problem is to choose a single information policy μ maximizing the payoff function

$$\begin{aligned}\Pi[\mu] = q\,\mathbb{E}\int_0^\infty & e^{-rt}|2\mu_t - 1|\,dt \\ & - (1-q)\sup_b \mathbb{E}\int_0^\infty \exp\left(-rt - \gamma\int_0^t b_s\,ds\right)(Kb_t - 1)|2\mu_t - 1|\,dt,\end{aligned} \tag{1}$$

where expectations are with respect to uncertainty in μ and b (in case the disloyal agent employs a mixed strategy). The term $|2\mu_t - 1|$ appears both as a benefit of informing the loyal agent and as a cost of informing the disloyal agent, in case the latter has not yet been detected and would currently choose to undermine.

3 Solving the model

In this section we derive an optimal information policy and discuss its features in detail. Sections 3.1 and 3.2 present key observations that simplify the principal's problem. In particular, they reduce the problem from choosing an optimal stochastic process to designing a deterministic one-dimensional disclosure path. Section 3.3 states a proposition characterizing the optimal disclosure path, and provides economic intuition for its qualitative features. Section 3.4 traces the formal derivation of an optimal disclosure path, outlining its main steps and explaining their logic. Section 3.5 characterizes the optimal information policy, and Section 3.6 characterizes when the principal pre-emptively fires the agent without offering an operating contract. Section 3.7 establishes existence of a unique undermining path under which the optimal contract is time-consistent.

[9] Interestingly, the same result holds in the static persuasion model of Guo and Shmaya (2018) under a set of regularity conditions, though for different reasons.

3.1 Disclosure paths

The structure of payoffs in the model is such that from an ex ante point of view, flow payoffs depend only on the precision of the agent's beliefs, and not on whether the beliefs place higher weight on the event $\omega = R$ or $\omega = L$. As a result, our analysis will focus on designing the process $|2\mu_t - 1|$, which measures how precisely the agent's beliefs pin down the true state.

Definition 1. *Given an information policy μ, the associated* disclosure path *x is the function defined by $x_t \equiv \mathbb{E}|2\mu_t - 1|$.*

The disclosure path traces the ex ante expected precision of the agent's beliefs over time. The fact that x must be generated by a martingale belief process places several key restrictions on its form. First, the martingality of μ implies that the ex post precision process $|2\mu_t - 1|$ is a submartingale by Jensen's inequality. The disclosure path must therefore be an increasing function, reflecting the fact that on average the agent must become (weakly) more informed over time. Second, ex post precision can never exceed 1, so $x_t \leq 1$ at all times.

Finally, the agent receives an exogenous signal of the state upon hiring which is correct with probability $\rho \geq 1/2$. Hence the initial precision of his beliefs prior to any information disclosure is $\phi \equiv |2\rho - 1|$. This places the lower bound $x_0 \geq \phi$ on the average precision of the agent's beliefs at the beginning of the disclosure process, allowing for the possibility of initial time-zero disclosure.

The following lemma establishes that the properties just outlined, plus a technical right-continuity condition, are all of the restrictions placed on a disclosure path by the requirement that it be generated by a martingale belief process.

Lemma 2. *A function x is the disclosure path for some information policy iff it is càdlàg, monotone increasing, and $[\phi, 1]$-valued.*

This characterization isn't quite enough to pass from analyzing information policies to disclosure paths, as the principal might well choose an information policy which leads to a stochastic ex post precision process, and such stochasticity will in general impact the disloyal agent's best response and thus the principal's payoff.

Definition 2. *An information policy μ is* deterministic *if its disclosure path x satisfies $x_t = |2\mu_t - 1|$ at all times and in all states of the world.*

Deterministic information policies are those for which the disclosure path is a sufficient statistic for the agent's information at each point in time.[10] The following lemma justifies focusing on such policies, by showing that the principal can pass from an arbitrary information policy to a deterministic policy yielding the same disclosure path and a weakly higher payoff. The basic intuition is that varying the precision of the agent's beliefs allows the disloyal agent to tailor her undermining strategy more precisely to the state of the world. Pooling these states reduces the disloyal agent's flexibility without reducing the average effectiveness of the loyal agent's actions.

Lemma 3. *Given any information policy μ, there exists a deterministic information policy μ' with the same disclosure path. All such policies μ' yield a weakly higher payoff to the principal than μ.*

When the principal employs a deterministic information policy, her payoff function may be written

$$\Pi[x] = q \int_0^\infty e^{-rt} x_t \, dt - (1-q) \sup_b \mathbb{E} \int_0^\infty \exp\left(-rt - \gamma \int_0^t b_s \, ds\right) (K b_t - 1) x_t \, dt,$$

i.e. a function of the disclosure path x alone rather than the full ex post precision process. In light of Lemma 3, we will restrict attention to deterministic information policies and focus on the design of the disclosure path x.

3.2 Loyalty tests

We will be particularly interested in studying information policies which do not provide incentives for disloyal agents to feign loyalty, i.e. refrain from undermining, at any point during their employment.

Definition 3. *An information policy is a* loyalty test *if undermining unconditionally at all times is an optimal strategy for the disloyal agent.*

The first main result of our analysis is that the principal cannot benefit by using the promise of future information disclosure to induce loyalty by a disloyal agent early on. Instead, the principal optimally employs only loyalty tests to screen disloyal from loyal agents. The following proposition states the result.

Proposition 1. *Suppose an information policy μ is not a loyalty test. Then there exists a loyalty test μ' yielding a strictly higher payoff to the principal than μ.*

[10] Importantly, deterministic information policies still induce stochastic posterior belief processes — the precision of the agent's beliefs is non-random, but the likelihood of a particular state being the true one is a random variable.

An important ingredient of this result is the opposition of interests of the principal and disloyal agent. Fixing a disclosure rule, if the disloyal agent finds it optimal to feign loyalty for a time rather than undermine immediately, the gains from remaining employed and exploiting future information must outweigh the losses from failing to undermine the principal early on. As gains to the disloyal agent are losses to the principal, the principal suffers from the disloyal agent's decision to defer undermining.

Of course, the principal cannot force the disloyal agent to undermine the project. The best she can do is design a loyalty test to ensure the disloyal agent cannot raise his payoff by feigning loyalty. Thus testing loyalty constrains the set of implementable information policies, and so it is not immediate from the logic of the previous paragraph that loyalty tests are optimal.

The key additional insight is the observation that the payoff of any disclosure rule which is not a loyalty test can be improved by releasing information more quickly. Intuitively, if over some time interval the disloyal agent strictly prefers to feign loyalty, then the principal may bring forward the release of some amount of information from the end to the beginning of this interval without disturbing the optimality of feigning loyalty during the interval. By bringing forward enough information, the principal can make the disloyal agent indifferent between undermining or not over the interval. This modification also makes the policy locally a loyalty test over the interval, without improving the disloyal agent's payoff, since feigning loyalty remains (weakly) optimal and thus the early information release can't be used to harm the principal. And the modification improves the principal's payoff when the loyal agent is present, as the loyal agent benefits from the earlier information release.

The proof of Proposition 1 builds on this insight by modifying a given disclosure policy beginning at the first time feigning loyalty becomes optimal. The modified policy releases as much additional information as possible in a lump at that instant without improving the disloyal agent's continuation utility. Afterward, it releases information smoothly at the maximum rate consistent with the constraints of a loyalty test. This process is not guaranteed to yield more precise information disclosure at every point in time versus the original policy. However, the proof shows that enough information release is brought forward to improve the loyal agent's payoff, and therefore the principal's.

3.3 The optimal disclosure path

In this section we state Proposition 2, which characterizes the optimal disclosure path, and provide economic intuition for its features. To aid exposition, we defer formal arguments to the following section. The optimal disclosure path takes one of several forms, depending on the parameter values. In all forms, the disclosure path grows deterministically, and the disloyal type of agent weakly prefers to undermine the principal at all times (and sometimes strictly). In Section 3.5, we characterize the information policy which implements the optimal disclosure path.

To gain intuition, consider a relaxed problem in which the disloyal agent *exogenously* undermines at all times, regardless of the information provided. Given a solution to this relaxed problem, if the disloyal agent does in fact prefer to undermine at all times, it is also a solution to the original problem. Assuming that the disloyal agent is always undermining, the agent's reputation at any time, conditional on not having been caught undermining, is calculated by Bayes' rule as $\pi_t = \frac{q}{q+(1-q)e^{-\gamma t}}$, where $\pi_0 = q$ is the starting reputation. Define

$$t^* \equiv \max\left\{0, \frac{1}{\gamma}\ln\left((K-1)\frac{1-q}{q}\right)\right\}$$

to be the first time the agent's reputation is at least $\frac{K-1}{K}$, and note that the principal's payoff in (1) is maximized pointwise by setting $x_t = 1$ if $t \geq t^*$ and $x_t = 0$ otherwise. If the starting reputation of the agent is above $\frac{K-1}{K}$, then it is best to fully inform the agent from the start, and under this policy, the disloyal agent has no incentive to wait and strictly prefers to undermine at all times, so the policy solves the original problem as well.

If the starting reputation is below $\frac{K-1}{K}$, then $t^* > 0$, and it must be checked that the disloyal agent weakly prefers to undermine at all times. Now the disloyal agent's strongest temptation *not* to undermine is just prior to time t^*, when his continuation payoff at stake is largest relative to his flow payoff from undermining. Provided that the jump in information is not too large — that is, provided that ϕ is above some threshold $\overline{x}$, the disloyal agent is willing to undermine at all times, and the solution to the relaxed problem is also the solution to the original problem. Hereafter, we refer to this as the "high knowledge" case.

If the starting level of information is below $\overline{x}$, then the naive solution above fails because the jump in information at time t^* is so large that the disloyal agent would temporarily prefer to cooperate if he has not yet been caught. In this case some prior disclosures must be made so that the jump in the precision of the agent's information at the final disclosure time isn't too large. The following definition characterizes the profit-maximizing way to make these prior disclosures.

Definition 4. *A disclosure path x satisfies the* optimal-growth *property if there exist time* *thresholds* $\underline{t} \geq 0$ *and* $\overline{t} \geq \underline{t}$ *such that*

$$x_t = \begin{cases} \phi, & t < \underline{t} \\ x_{\underline{t}} \exp\left((r + \gamma/K)(t - \underline{t})\right), & \underline{t} \leq t < \overline{t} \\ 1, & \overline{t} \leq t \end{cases}$$

and x *is continuous at every* $t \neq \overline{t}$*. In particular,* $x_{\underline{t}} = \phi$ *if* $0 < \underline{t} < \overline{t}$*.*

One requirement imposed by the optimal-growth property is that, after an initial period of no disclosure, x must grow at rate $r+\gamma/K$ until the time $\overline{t}$ of final disclosure. The intuition for the desirability of this property is as follows. Recall that, all else equal, the principal maximizes profits by minimizing the precision of the agent's information prior to time t^*. Thus, looking earlier in the contract from t^*, the quality of the agent's information should deteriorate as quickly as possible consistent with incentive-compatibility. Equivalently, going forward in time information is released as quickly as possible. It turns out that the maximum permissible rate of disclosure is $r + \gamma/K$. This modification to the ideal disclosure path can be visualized graphically — the constant disclosure level prior to time t^* is shifted upward to ensure undermining is optimal near time t^*, and then is tilted counterclockwise as steeply as possible consistent with incentive compatibility.

The second requirement imposed by the optimal-growth property is that x not drop below ϕ, the minimum admissible disclosure. This property must be imposed because if ϕ is large, then the shift-and-tilt described previously may produce a disclosure path which drops below ϕ early in the contract. This violates the lower bound constraint on x required for implementability. As a remedy, the principal optimally irons the disclosure path so that it does not drop below ϕ. Figure 1b illustrates this ironing in the early stage of the contract. Ironing is necessary when ϕ is below $\overline{x}$ but above a threshold $\underline{x}$. We refer to this case as *moderate knowledge*. Intuitively, the principal delays information release at the beginning while she observes the agent and, absent any detected undermining, she raises her beliefs about the agent's loyalty. On the other hand, if the disclosure path traces back to a level *above* ϕ, then the principal happily jump starts the relationship with an additional atom of information above ϕ. We refer to this case, which occurs when $\phi < \underline{x}$, as *low knowledge*, and it is illustrated in Figure 1c.

The procedure just described for adapting the disclosure path to achieve a loyalty test takes as given that the time of final disclosure remains fixed at t^*. In fact, the principal will in general also wish to delay this time, choosing a final disclosure time $\overline{t}$ which is larger than t^*. The tradeoff faced by the principal in choosing $\overline{t}$ is as follows: increasing $\overline{t}$ reduces disclosure and losses early in the contract (i.e. when the principal's posterior beliefs about the agent's type are below $(K-1)/K$), but delays disclosure and profits late in the contract (when posterior beliefs are above $(K-1)/K$). The unique optimal disclosure path chooses $\bar{t}$ to balance these two forces. In particular, in both the low knowledge and moderate knowledge cases, the jump is strictly later than t^*.[11]

To summarize, in all parametric cases, the optimal contract proceeds through three phases (the second of which may be degenerate): an initial calibration phase in the form of either an initial discrete disclosure or a quiet period in which no information is released, a period of gradual information release, and finally a jump in which all remaining information is released. This contract will satisfy the optimal growth property for some time thresholds $\underline{t}$ and $\bar{t}$. The threshold $\underline{t}$ marks the transition from the first to the second phase, while $\bar{t}$ is the boundary between the second and third phases.

Proposition 2, stated below, formally characterizes an optimal disclosure path. Section 3.4 provides a detailed walk-through of its proof. Explicit expressions for the optimal thresholds appearing in the proposition are provided in Appendix A.6. Figure 1 illustrates the optimal disclosure path for the low, moderate, and high knowledge cases.

Proposition 2 (Optimal Disclosure Path). *There exists a unique optimal disclosure path x^*, which satisfies the optimal-growth property for some thresholds $\underline{t}$ and $\bar{t}$ and some x_0^*. If $q \geq (K-1)/K$, then $\underline{t} = \bar{t} = 0$. Otherwise, there exist disclosure thresholds $\underline{x}$ and $\overline{x} > \underline{x}$ in $(0,1)$, independent of ϕ, such that:*

- *If $\phi \geq \overline{x}$, then $\underline{t} = \bar{t} = t^*$,*
- *If $\underline{x} < \phi < \overline{x}$, then $0 < \underline{t} < t^* < \bar{t}$,*
- *If $\phi \leq \underline{x}$, then $0 = \underline{t} < t^* < \bar{t}$ and $x_0^* = \underline{x}$.*

3.4 Deriving the optimal disclosure path

In this section we provide an overview of the formal proof of Proposition 2. In particular, we establish a series of lemmas which lay the groundwork for the proof. The formal proof of the proposition, stated in Appendix A.6, builds on these lemmas to complete the derivation.

In light of Proposition 1, the principal optimizes her payoff among the set of information policies which are loyalty tests. We must therefore characterize the constraints that testing loyalty places on information policies. Given a disclosure path x, define $U_t \equiv$

[11] To see this, first note that if the jump is before t^*, the principal unambiguously benefits by shifting the path to the right, since this only reduces information at times when the agent's reputation is low. Moreover, if the jump is at t^*, there is positive marginal benefit and zero marginal cost to shifting the path to the right, since at time t^* the coefficient on x_t in the principal's payoff vanishes.

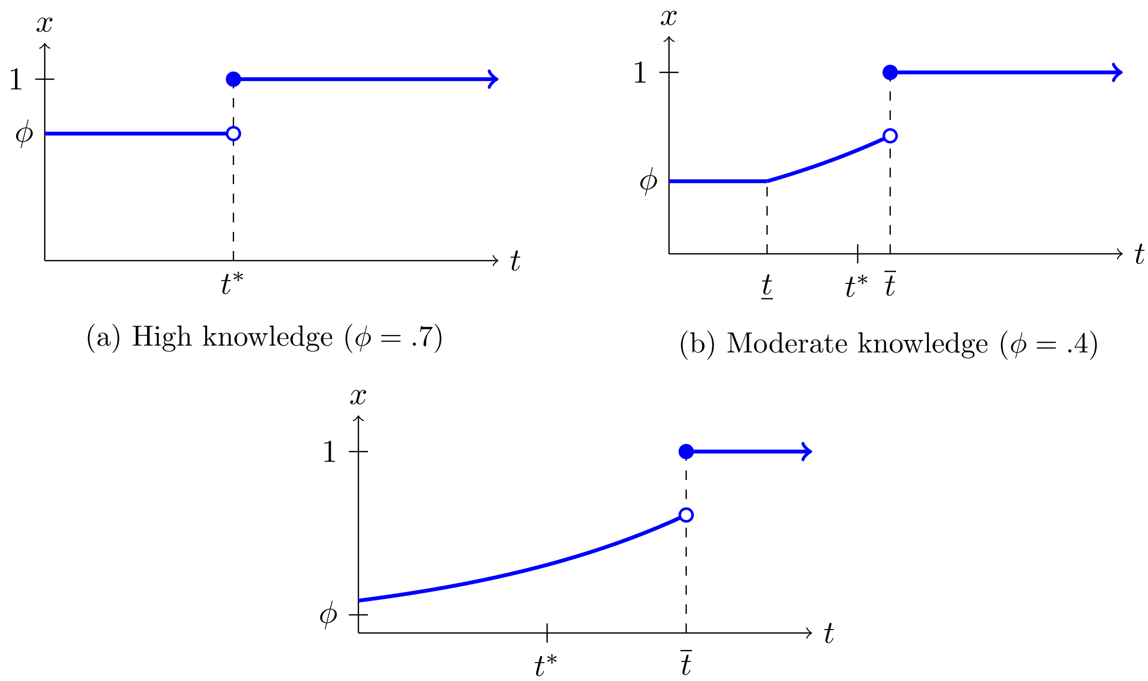

(a) High knowledge ($\phi = .7$)

(b) Moderate knowledge ($\phi = .4$)

(c) Low knowledge ($\phi = .1$)

Figure 1: Three forms of the optimal disclosure path for $K = 3$, $q = .2$, $r = .05$, $\gamma = 2$ and various values of starting knowledge ϕ.

$\int_t^\infty e^{-(r+\gamma)(s-t)}(K-1)x_s\,ds$ to be the disloyal agent's ex ante continuation utility at time t, supposing he has not been terminated before time t and undermines forever afterward. The following lemma characterizes a necessary condition satisfied by any loyalty test, which is also sufficient whenever the underlying information policy is deterministic.

Lemma 4. *Suppose x is the disclosure path of a loyalty test. Then*

$$\dot{U}_t \leq \left(r + \frac{\gamma}{K}\right) U_t \tag{2}$$

a.e. Conversely, if μ is a deterministic information policy with disclosure path x satisfying (2) a.e., then μ is a loyalty test.

Equation (2) ensures that from an ex ante perspective, at each time t the gains from undermining at that instant must outweigh the accompanying risk of discovery and dismissal. Thus it is surely at least a necessary condition for any loyalty test. The condition is not in general sufficient, as incentive-compatibility might be achieved on average but not ex post in some states of the world. However for *deterministic* information policies, ex ante and ex post incentives are identical, and equation (2) is a characterization of loyalty tests. Recall from Lemma 3 that the principal might as well choose a deterministic information policy, so

(2) may be taken as a characterization of loyalty tests for purposes of designing an optimal policy.

In light of Proposition 1 and Lemmas 3 and 4, the principal's design problem reduces to solving

$$\sup_{x\in\mathbb{X}} \int_0^\infty e^{-rt} x_t (q - (1-q)(K-1)e^{-\gamma t})\, dt \quad \text{s.t.} \quad \dot{U}_t \leq (r + \gamma/K) U_t, \tag{PP}$$

where $\mathbb{X}$ is the set of monotone, càdlàg $[\phi, 1]$-valued functions and $U_t = \int_t^\infty e^{-(r+\gamma)(s-t)}(K-1)x_s\, ds$.

The form of the constraint suggests transforming the problem into an optimal control problem for U. The following lemma faciliates this approach, using integration by parts to eliminate x from the objective in favor of U.

Lemma 5. *For any $x \in \mathbb{X}$,*

$$\int_0^\infty e^{-rt} x_t (q - (1-q)(K-1)e^{-\gamma t})\, dt = -\left(1 - \frac{K}{K-1} q\right) U_0 + \frac{q\gamma}{K-1} \int_0^\infty e^{-rt} U_t\, dt.$$

The final step in the transformation of the problem is deriving the admissible set of utility processes corresponding to the admissible set $\mathbb{X}$ of disclosure paths. The bounds on x immediately imply that $U_t \in [\underline{U}, \overline{U}]$ for all t, where $\overline{U} \equiv (K-1)/(r+\gamma)$ and $\underline{U} \equiv \phi \overline{U}$. It is therefore tempting to conjecture that problem (PP) can be solved by solving the auxiliary problem

$$\sup_{U\in\mathbb{U}} \left\{ -\left(1 - \frac{K}{K-1} q\right) U_0 + \frac{q\gamma}{K-1} \int_0^\infty e^{-rt} U_t\, dt \right\} \quad \text{s.t.} \quad \dot{U}_t \leq (r + \gamma/K) U_t, \tag{PP'}$$

where $\mathbb{U}$ is the set of absolutely continuous $[\underline{U}, \overline{U}]$-valued functions. Any solution to this problem can be mapped back into a disclosure policy via the identity

$$x_t = \frac{1}{K-1}\left((r+\gamma)U_t - \dot{U}_t\right), \tag{3}$$

obtainable by differentiating the expression $U_t = \int_t^\infty e^{-(r+\gamma)(s-t)} x_s\, ds$.

It turns out that this conjecture is correct if ϕ is not too large, but for large ϕ the resulting solution fails to respect the lower bound $x \geq \phi$. To correct this problem equation (3) can be combined with the lower bound $x_t \geq \phi$ to obtain an additional upper bound $\dot{U}_t \leq (r+\gamma)U_t - (K-1)\phi$ on the growth rate of U. This bound suggests solving the modified auxiliary problem

$$\begin{aligned} &\sup_{U \in \mathbb{U}} \left\{ -\left(1 - \frac{K}{K-1} q\right) U_0 + \frac{q\gamma}{K-1} \int_0^\infty e^{-rt} U_t \, dt \right\} \\ &\text{s.t.} \quad \dot{U}_t \le \min\{(r+\gamma/K)U_t, (r+\gamma)U_t - (K-1)\phi\} \end{aligned} \tag{PP''}$$

The following lemma verifies that any solution to this relaxed problem for fixed $U_0 \in [\underline{U}, \overline{U}]$ yields a disclosure path respecting $x \ge \phi$. Thus a full solution to the problem allowing U_0 to vary solves the original problem.

Lemma 6. *Fix $u \in [\underline{U}, \overline{U}]$. There exists a unique solution U^* to problem (PP'') subject to the additional constraint $U_0 = u$, and the disclosure policy x^u defined by*

$$x_t^u = \frac{1}{K-1} \left((r+\gamma)U_t^* - \dot{U}_t^* \right)$$

satisfies $x^u \in \mathbb{X}$.

We prove Proposition 2 by solving problem (PP'') subject to $U_0 = u \in [\underline{U}, \overline{U}]$, and then in a final step optimize over u to obtain a solution to problem (PP'') and thus to problem (PP). When U_0 is held fixed, the objective is increasing in U pointwise, so optimizing the objective amounts to maximizing the growth rate $\dot{U}$ subject to the control and the upper bound $U \le \overline{U}$. The result is that U solves the ODE $\dot{U}_t = \min\{(r+\gamma/K)U_t, (r+\gamma)U_t - (K-1)\phi\}$ until the point at which $U_t = \overline{U}$, and then is constant afterward. Solving problem PP then reduces to a one-dimensional optimization over U_0, which can be accomplished algebraically.

The forces shaping an optimal continuation utility path can be cleanly mapped onto the forces shaping an optimal disclosure path identified in Section 3.3. An important property of the solution to problem PP'' is that only a single constraint on $\dot{U}_t$ binds at a given time, with the constraint $\dot{U}_t \le (r+\gamma)U_t - (K-1)\phi$ binding early in the optimal contract while $\dot{U}_t \le (r+\gamma/K)U_t$ binds later on. The first constraint binding corresponds to a constant disclosure level $x_t = \phi$, while the second constraint binding corresponds to x growing at rate $r + \gamma/K$. Thus, holding fixed the time of final disclosure, the optimal disclosure path declines at rate $r + \gamma/K$ as one rewinds the contract toward time 0, until the lower bound $x = \phi$ is reached and the agent cannot be made any less informed. The growth rate $r + \gamma/K$ can be interpreted as the "fastest possible" growth rate for x, in a global average sense. (It is of course possible to achieve locally higher growth rates of x by designing a very concave U, with the extreme case of a concave kink corresponding to a discrete disclosure.)

Meanwhile the jump in disclosure at the final disclosure time arises due to the kink in the continuation utility curve when it reaches the upper bound $\overline{U}$. The size of the jump depends on the slope of the utility curve to the left of the kink, which is optimally made as large as possible consistent with the IC constraint and $x \geq \phi$. This exactly accords with the considerations shaping the size of the jump discussed in Section 3.3. Finally, the timing of the final disclosure is controlled by the choice of the initial continuation utility U_0. The smaller the choice of U_0, the longer continuation utility must grow before it reaches $\overline{U}$, and thus the later the time of final disclosure. So the optimization over the disclosure time discussed informally in Section 3.3 is achieved formally by optimizing U_0 in the control problem PP''.

3.5 The optimal information policy

Proposition 2 characterizes the unique optimal disclosure path x^*, but does not actually describe an optimal information policy implementing x^*. Recall that Lemmas 2 and 3 guarantee the existence of such a policy, and ensure that it can be taken to be deterministic. However, a more explicit characterization is necessary to understand how optimal information disclosure should be implemented in practice. In this subsection we characterize an optimal information policy, and describe a signal process the principal may commit to which induces the desired belief process.

The basic problem is to find a martingale belief process μ^* satisfying $|2\mu_t^* - 1| = x_t^*$ at all times. Equivalently, at all times $\mu_t^* \in \{(1 - x_t^*)/2, (1 + x_t^*)/2\}$, meaning that pathwise μ^* is composed of piecewise segments of the two disclosure envelopes $(1 - x_t^*)/2$ and $(1 + x_t^*)/2$. In light of this fact, the requirement that μ^* be a martingale implies that it must be a (time-inhomogeneous) Markov chain transitioning between the upper and lower disclosure envelopes, and uniquely pins down transition probabilities for this process. Thus the process μ^* implementing x^* is uniquely defined, and can be explicitly constructed by calculating the transition rates between the envelopes which make μ^* a martingale.

We will now describe a signal process to which the principal may commit in order to induce the posterior belief process μ^*. To begin, define the agent's *state conjecture* at a given time t to be the state they currenty view as mostly likely; that is, their state conjecture is R if $\mu_t^* \geq 1/2$, and L otherwise. The principal influences the agent's beliefs by sending periodic recommendations about the correct state conjecture. The signal structure is designed such that each time a recommendation is received, the agent changes his state conjecture. (In the microfoundation given in Section 4.1, this recommendation corresponds to a change in the agent's optimal task action, and indeed the principal's optimal policy could be equivalently cast as a sequence of recommended task actions.) The overall arrival rate of recommendations is calibrated to induce the appropriate pace of disclosure, as follows.

First, at any time corresponding to a discontinuity of the disclosure path, the principal recommends a change in state conjecture with a probability which is higher when the agent's state conjecture is incorrect than when it is correct. In particular, when the agent reaches trusted status, the principal recommends a change in state conjecture if and only if the agent's conjecture is wrong, revealing the state.

Second, during the gradual disclosure phase the principal recommends a change in state conjecture via a Poisson process whose intensity is higher when the agent's conjecture is currently wrong. This sort of Poisson signal process is sometimes referred to as an inconclusive contradictory news process.[12] Interestingly, the expected intensity is constant over the gradual disclosure phase, so that the agent switches conjectures at a constant average rate throughout this phase.[13] However, as the phase progresses the informativeness of the Poisson signals, as measured by the gap in the arrival rate when the agent's conjecture is wrong versus right, increases.

The following proposition formally establishes the claims made above. It also provides additional details of the signal process implementing the optimal belief process μ^*.

Proposition 3 (Optimal Information Policy)**.** *There exists a unique deterministic information policy μ^* implementing x^*. It is induced by the following signal process:*

- *At time $t = 0$, a signal arrives with probability $\frac{(1+x_0^*)(x_0^*-\phi)}{2x_0^*(1-\phi)}$ if the agent's time $t = 0-$ state conjecture is incorrect, and with probability $\frac{(1-x_0^*)(x_0^*-\phi)}{2x_0^*(1+\phi)}$ otherwise,*

- *On the time inteval $(\underline{t}, \overline{t})$, a signal arrives with intensity $\overline{\lambda}_t \equiv \frac{1}{2}\left(r + \frac{\gamma}{K}\right)\frac{1+x_t^*}{1-x_t^*}$ if the agent's state conjecture is incorrect, and with intensity $\underline{\lambda}_t \equiv \frac{1}{2}\left(r + \frac{\gamma}{K}\right)\frac{1-x_t^*}{1+x_t^*}$ otherwise. Conditional on the agent's information, the expected intensity of signal arrival is $\frac{1}{2}\left(r + \frac{\gamma}{K}\right)$ at all times on this interval.*

- *At time $t = \overline{t}$, a signal arrives with probability 1 if the agent's time $t = \overline{t}-$ state conjecture is incorrect, and with probability 0 otherwise.*

Figure 2 shows a sample path of μ^*, the agent's posterior belief that $\omega = R$, in an optimal contract for the low knowledge case.

[12] This news process has been featured, for instance, as one of the learning processes available to a decision-maker in the optimal learning model of Che and Mierendorff (2017). In that paper, the news process has fixed informativeness, as measured by the ratio of arrival rates in each state of the world. By contrast, our optimal disclosure process features an informativeness which rises over time. Such inhomogeneous contradictory news processes are considered, for instance, in the more general learning model of Zhong (2018).

[13] Our optimal information policy during the gradual disclosure phase bears some resemblance to the suspense-maximizing policy in Ely, Frankel, and Kamenica (2015), which also induces a deterministic belief precision path and features inconclusive contradictory news in the form of plot twists. However, suspense-maximization dictates a *decreasing* rate of plot twists, whereas our optimal policy features a constant arrival rate of contradictory news.

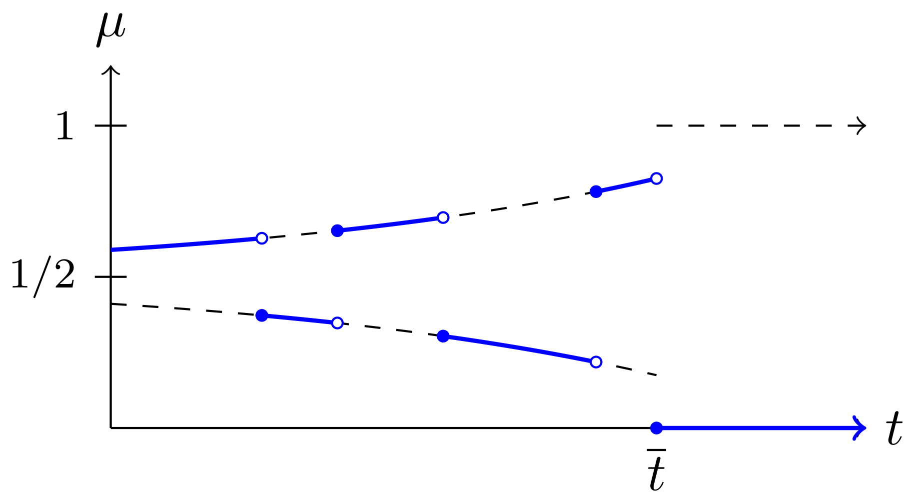

Figure 2: Sample path of information in the optimal contract (low knowledge case). Parameter values match the third panel of Figure 1.

3.6 Pre-emptive firing

We have so far ignored the possibility that the principal might fire the agent even if he is not known to be disloyal. This possibility is relevant if the agent's initial reputation q is small and his initial knowledge ϕ is large. If the agent is ever fired while his type is uncertain, termination will optimally take place at time zero when his reputation is as low as possible, i.e. before the contract is signed. In this subsection we discuss when such pre-emptive firing is optimal. Equivalently, we identify when the optimal contract characterized above is profitable for the principal.

Figure 3 displays a representative division of the (q, ϕ)-parameter space into hire and no-hire regions. Regions I, II, and III represent parameter regions where $q < (K-1)/K$ and the agent begins with high, moderate, and low knowledge, respectively. Meanwhile region IV is the set of parameters for which $q \geq (K-1)/k$. The shaded region represents the parameter values for which the principal prefers to pre-emptively fire the agent immediately rather than offer an operating contract. Note that this outcome occurs only for a subset of regions I and II, when the agent begins with moderate or high knowledge. By contrast, if $q \geq (K-1)/K$ or the agent begins with low knowledge, the optimal operating contract always yields strictly positive payoffs, and pre-emptive firing does not occur. This is unsurprising given that in these situations, the optimal contract involves voluntary disclosure of information by the principal at time zero.

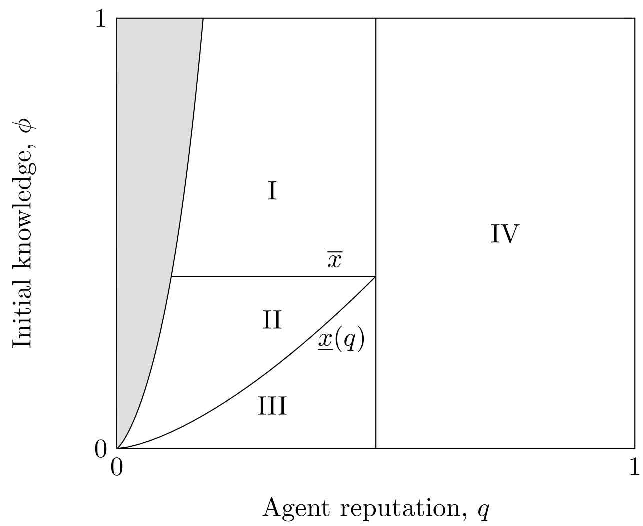

Figure 3: Optimal contract regimes in (q, ϕ)-space. Regions I, II, III represent the high, moderate, and low knowledge cases when $q < (K-1)/K$, while in region IV $q \geq (K-1)/K$. Pre-emptive firing is optimal in the shaded region.

3.7 Time-consistent undermining

As noted in Section 2.3, when characterizing the principal's optimal contract it is not necessary to specify a recommended undermining policy for the disloyal agent. This is because the zero-sum nature of the game between principal and disloyal agent implies that indifferences for the agent correspond to indifferences for the principal. There may then exist many profit-maximizing agent-optimal undermining policies for a given information policy. In particular, under the optimal information policy μ^*, throughout the gradual disclosure phase the agent is indifferent between undermining or feigning loyalty. In this subsection we show that if we wish the contract to additionally be *time-consistent*, then the intensity of the agent's undermining at each moment of time is uniquely determined.

To allow for variable intensity of undermining, we extend the strategy space available to the agent. The agent may now choose any measurable path $\beta : \mathbb{R}_+ \to [0, 1]$, with β_t specifying the intensity of undermining at time t. Undermining at intensity $\beta_t \in [0, 1]$ incurs expected flow costs $\beta_t |2\mu_t - 1|$ and is detected by the principal at rate $\gamma \beta_t$. This extended strategy space is the continuous-time equivalent to allowing behavioral mixed strategies in discrete time.[14]

Fixing an undermining policy β, let π be the principal's induced posterior belief process

[14]Directly defining behavioral mixed strategies in continuous time is not possible due to measurability issues. An alternative approach is to allow mixtures over entire paths. In this alternative specification, it can be shown that there exists a time-consistent mixed strategy involving randomization over the time at which undermining begins during the gradual disclosure phase.

that the agent is loyal. Given an information policy μ^*, an undermining policy β is *time-consistent* if for each time t, the continuation policy $(\mu^*_s)_{s\geq t}$ is an optimal information policy for the principal's problem under initial belief π_t that the agent is loyal and initial agent beliefs μ^*_t about the state. Informally, at no point in time should the principal wish to renege on the information policy previously promised, and substitute a new information policy in its place going forward, based on her updated beliefs about the agent's type. (Of course, any such rewritten policy must respect any information already disclosed to the agent.) To side-step difficult issues of modeling continuous-time games without commitment,[15] we do not formally characterize the time-consistent solution as an equilibrium of a dynamic game satisfying some notion of sequential rationality; however, we see it as very much in the same spirit.

We begin by noting when $q \geq (K-1)/K$ or the agent begins with high knowledge, there is no gradual disclosure phase and the unique optimal undermining policy is $\beta^* = 1$. This policy turns out to be time-consistent. If $q \geq (K-1)/K$, then under μ^* the principal releases all information about the state immediately, meaning the information policy is trivially time-consistent. So consider the case that $q < (K-1)/K$ but the agent begins with high knowledge. Under μ^* the principal releases all information at time t^*, hence after this point the contract is trivially time-consistent. So it suffices to consider time-consistency for time prior to t^*. Let π be the principal's posterior belief process that the agent is loyal. By definition, t^* is the time at which π reaches $(K-1)/K$, since the disloyal agent undermines at all times. Hence prior to time t^* the agent's reputation remains sufficiently low that the principal's optimal continuation policy is to continue releasing no information until time t^*.

When $q < (K-1)/K$ and the agent begins with low or moderate knowledge, the optimal information policy has a non-degenerate gradual disclosure phase. In these cases any undermining policy on the gradual disclosure phase is optimal. Critically, the choice of undermining policy determines the path of the principal's posterior belief process that the agent is loyal. Undermining by the disloyal agent at all times turns out not to be time-consistent. This is because under such a policy, the principal's posterior beliefs reach the critical threshold $\frac{K-1}{K}$ at which full information release is optimal at time t^*; but μ^* specifies that the principal hold some information back until time $\bar{t} > t^*$. Thus for times on $[t^*, \bar{t})$ the optimal continuation contract disagrees with μ^*.

To restore time-consistency, the agent must undermine with lower intensity in order to ensure that the principal's posterior beliefs reach $(K-1)/K$ precisely at time $\bar{t}$. In fact,

[15] See, for instance, Orlov, Skrzypacz, and Zryumov (2018) and Daley and Green (2018) for recent papers tackling these issues. Additionally, in our setting the principal is informed of the true state once the project begins. A full no-commitment model would then need to reconcile the principal's persuasion technology with the potential signaling of private information by the principal's persuasion actions.

it must do more - it must raise beliefs at precisely the rate which ensures that at no point during the gradual disclosure phase would the principal prefer to release an additional atom of information, or cease releasing information for some time. There turns out to be a unique undermining policy achieving this property. The following proposition formalizes this result. It also highlights the important fact that the unique time-consistent undermining policy may be non-monotonic: when the agent begins with moderate knowledge, undermining initially occurs at the maximum rate, then drops to an interior level at the start of the gradual disclosure phase and rises steadily thereafter.

Proposition 4. *Under the information policy μ^*, there exists an essentially unique time-consistent optimal undermining path β^*. $\beta^*_t = 1$ for $t \in [0, \underline{t}) \cup [\overline{t}, \infty)$, while on $[\underline{t}, \overline{t})$, $\beta^* < 1$ is a strictly increasing, continuous function.*

4 Interpreting the model

Our model is designed to capture, in stylized fashion, the incentives and informational environment in an organization experiencing leaks or sabotage. In this subsection we provide additional commentary on several key modeling choices, interpreting and microfounding them in the context of our application.

4.1 The principal's payoffs

The principal's payoffs in our model are designed to capture the fact that information sharing is both necessary for the agent to productively perform her job, and harmful on net when used maliciously.

The flow payoff $|2\mu_t - 1|$ deriving from the agent's task performance can be microfounded as follows. Suppose that at each time t the agent's task is to match a binary action $a_t \in \{L, R\}$ with a persistent state $\omega \in \{L, R\}$ at each time $t \in \mathbb{R}_+$. The principal's ex post flow payoff at time t from task performance is

$$\pi(a_t, \omega) \equiv \mathbf{1}\{a_t = \omega\} - \mathbf{1}\{a_t \neq \omega\}.$$

Under this payoff specification, the principal-optimal task action given the agent's posterior belief $\mu_t = \Pr_t(\omega = R)$ about the state is $a^*_t = R$ if $\mu_t > 1/2$, and $a^*_t = L$ if $\mu_t < 1/2$. The principal's expected profits at time t, conditional on the agent's information set, are then $|2\mu_t - 1|$. Under this microfoundation the task action a is observable, and so optimal task performance (i.e. implementation of a^*) can be costlessly enforced by the threat of immediate termination.

The flow damage $K|2\mu_t - 1|$ can be microfounded in a similar way by introducing a hidden action which inflicts damage only when it is matched to the state. The assumption that $K > 1$ implies that the expected flow cost from undermining outweighs the expected flow benefit from optimal task performance at all information levels. Hence were the agent known to be disloyal the principal would prefer not to employ him. (The case $K \leq 1$ is trivial, since the disloyal agent would obtain a nonpositive payoff from accepting the contract and would therefore be willing to screen himself out.)

The specification of task payoffs and undermining damage could be generalized to arbitrary affine functions of $|2\mu_t - 1|$ without substantially changing the analysis. Such payoff structures would be generated by, for instance, a task action whose payoff from matching the state is greater or less than its harm from mismatch; or an undermining action which inflicts damage when it matches the state, but is harmless when mismatching. We have used a linear payoff specification merely for clarity of exposition.

Generalizing further to nonlinear functions of $|2\mu_t - 1|$ would introduce a new force into the model — payoff-relevance of stochasticity in the precision of the agent's information. The principal would then face a non-trivial design problem not only over how quickly to disclose information to the agent, but also over how much randomness to introduce into the timing of disclosure.

Specifically, under linear payoffs a key step in our analysis (Lemma 3) established that an arbitrary information policy can always be replaced by a deterministic policy weakly increasing the principal's payoffs. Were the task payoff convex in the precision of beliefs over some range, this result may not hold; the convexity would create an incentive for introducing stochasticity to the disclosure path. Of course, such randomization might also increase a disloyal agent's payoffs, if the harm inflicted by undermining were also convex in the precision of beliefs. Allowing for convexity in both task and undermining payoffs would thus require trading off these factors, building in an additional layer of analysis into the design problem. We have chosen to abstract from these considerations to clarify the dynamics of trust-building and timing of undermining.

4.2 The agent's payoffs

We have assumed that the agent has one of two rather stark motives when working for the principal — he either wants to benefit or harm the principal as much as possible. Our goal in writing payoffs this way is to isolate and study the implications for information design of screening out agents whose incentives cannot be aligned. For one thing, we believe this to be a quite realistic description of several of the cases of leaks and sabotage documented in the introduction. We also see methodological value in abstracting from considerations of at what times and to what extent the agent should be given incentives to work in the interest of the firm. These optimizations are often complex, and have been studied extensively elsewhere in the literature. By eliminating such factors we obtain the cleanest analysis of the impact of a pure screening motive.

An important implication of our preference specification is that the agent faces no uncertainty about whether he would prefer to oppose the principal, only about how to effectively do so. In other words, the hidden state ω captures not the overarching goals of the principal, but rather the means by which the goal can be attained. We see this as a realistic assumption in settings in which a mole or saboteur is inserted into the organization expressly to undermine it, or in which an employee becomes disgruntled for personal rather than ideological reasons (for instance, displeasure at being passed over for a promotion).

However, it is plausible that in some contexts an employee might decide to undermine their organization only after learning about the organization's long-run goals or values. In such cases, our model could be considered the second stage of a larger game in which the agent first becomes informed about the principal's goals. Indeed, in such a setting disloyal employees may naturally enter the second stage with information about ω obtained incidentally in the course of ascertaining the principal's motives, motivating the initial signal received by the agent in our model.

4.3 The task monitoring structure

We have assumed that the agent's task action is not subject to moral hazard, while any steps he takes to undermine the principal are detected only stochastically. We view this as a reasonable first approximation to the basic dilemma faced by a leaker or saboteur, who must remain on the job nominally aiding the organization while simultaneously undertaking damaging acts which might be discovered. The inability of the principal to detect all acts of undermining can be justified in an organizational context where the principal employs many agents. In this setting, an undetected act of undermining may be interpreted as an instance which the principal is unable to trace back to a particular agent.[16]

What if the agent's task action were also subject to moral hazard? One alternative specification would be to make all deviations of the agent observed only stochastically, including all task actions underlying the flow payoff collected by the principal. If deviations from

[16] In particular, in the limit of an organization with a continuum of agents, our model would remain unchanged if the principal were allowed to observe its aggregate flow payoffs, since this process could not identify undermining by individual agents even statistically.

the desired task action are detected at the same rate as undermining, then the agent always coordinates improper task performance and undermining. The result is a model in which the agent contributes a flow payoff of $|2\mu_t - 1|$ when cooperating and inflicts net flow damages of $(K+1)|2\mu_t - 1|$ when not cooperating. So by rescaling K, this model can be mapped onto one in which the task action is directly observable.

Another specification would allow improper task performance to go undetected, in contrast to the stochastic detection of explicit undermining.[17] In this environment, the agent can effectively undermine using both low-stakes and high-stakes methods. So long as low-stakes undermining is undetectable, the disloyal agent always chooses to undermine at some level. If high-stakes undermining is detected with high probability, i.e. γ is large, the agent might choose never to undermine at that level, even with no future information forthcoming, so as to remain employed and continue low-stakes undermining. Otherwise, the agent makes a dynamic choice whether to engage in high stakes undermining, taking into account the rate of future information arrival. The analysis of such a model would be substantially similar to ours, differing only in details.[18]

The most natural model specifications with hidden task actions lead to analyses which differ in only minor ways from our model. We have therefore chosen to focus on the simpler model with no moral hazard over the task action to simplify exposition.

4.4 The agent's initial information

Our specification of the information structure allows for the agent to observe an exogenous time-zero signal of the project state. This signal is designed to capture several different pre-play channels of information potentially available to the agent. First, in some settings organizations may not be able to exclude new hires from obtaining *some* sense of the nature of their project simply by being employed. For instance, information could be gleaned from discussions with coworkers, company-wide e-mails, or from initial training the employee must be given to be effective at his job.

Second, exogenous initial information arises naturally in a larger model in which an organization becomes aware it is being undermined only after employees have already been on the job for some time and had access to project-related information. Such a supergame leads naturally to a partially informed cohort when the principal begins the process of rooting out disloyal elements. (In fact, it might be natural in some cases to assume employees know

[17] We thank Piotr Dworczak for suggesting this specification.

[18] In particular, a "loyalty test" result analogous to Proposition 1 would hold, with the principal always making the agent at least weakly willing to pursue high-stakes undermining. The reason is that, just as in our model, if low-stakes undermining is strictly optimal then the principal can bring forward information to benefit the loyal agent without increasing the utility of the disloyal agent.

everything about the project prior to a loyalty scare. In such applications the organization must rely on changes in strategy that make existing knowledge out of date to limit the damage of an underminer, an extension discussed in Section 6.3.)

In case none of these information channels is relevant in a particular application, the corresponding solution can be recovered as a special case of our model by setting $\rho = 1/2$.

5 Comparative statics

In this section we discuss comparative statics for important features of the optimal contract with respect to changes in the key model parameters. For clarity of exposition we will focus on the moderate knowledge case; the comparative statics for other cases are similar and are provided in Section A.12, along with proofs for all results reported in this section.

Table 1 below reports comparative statics for each model output (by column) with respect to each model input (by row). Here Π are the principal's expected profits under the optimal contract; $\underline{t}$ and $\bar{t}$ are the threshold times denoting the end of the quiet period and the time at which the agent reaches trusted status, respectively; $\Delta \equiv \bar{t} - \underline{t}$ is the length of the gradual disclosure phase; and $\overline{x}$ is the precision of the agent's beliefs just prior to reaching trusted status.

Table 1: Comparative statics

	$d\overline{x}/$	$d\Delta/$	$d\underline{t}/$	$d\bar{t}/$	$d\Pi/$
$d\gamma$	$+$	$\pm$	$-$	$-$	$+$
dq	0	0	$-$	$-$	$+$
dK	$+$	$+$	$+$	$+$	$-$
dr	$-$	$-$	$+$	$-$	$-$
$d\phi$	0	$-$	$+$	$-$	$-$

An increase in the damage K from undermining has a negative direct effect on the principal. This direct effect implies that the principal's payoff under the optimal contract unambiguously decreases with K: fixing an information policy, loyalty test or otherwise, and a disloyal agent strategy, the principal's expected flow payoff is made lower at each instant in time, and hence her ex ante payoff, after minimizing over disloyal agent strategies and maximizing over information policies, must decrease. The impact of higher K on the information policy is a bit more subtle, due to effects on the agent's incentives. Despite that the disloyal agent enjoys inflicting larger damage on the principal, the optimal contract must provide *stronger* incentives to make the disloyal agent willing to undermine. With higher K, the agent obtains both more value from undermining today but also a higher continuation value from being employed and able to undermine in the future. Importantly, the former value is proportional to K while the latter value is proportional to $K-1$, which grows faster in relative terms, and hence the net effect on the disloyal agent is a stronger incentive to feign loyalty. In response, fixing the final jump time, the principal must front load information by starting the smooth release sooner and release it more slowly thereafter, which is reflected in both a smaller final jump and a longer smooth release period. In a final optimization, the principal then delays both the start and end of the smooth release period, as this delay now has a larger (helpful) effect of reducing information in the early part of the contract, when the agent is likely to be disloyal, and a smaller (harmful) effect of reducing information in the later part of the contract.

An increase in the quality γ of the monitoring technology directly benefits the principal through faster detection of disloyal types. Moreover, as with K, this benefit occurs independently of the information policy and disloyal agent's strategy, which allows us to conclude without further calculation that the principal's payoff increases with γ. In addition to the direct benefit, the principal is also indirectly affected in two ways through the disloyal agent's incentives. While better monitoring makes the disloyal agent more likely to be detected if he undermines today and this increases the cost of undermining, it also increases the agent's effective discount rate — remaining employed in the future becomes less valuable — which decreases the cost of undermining. The first effect dominates just prior to the final information jump, and to keep the policy a loyalty test, the final jump must be smaller. But the second effect dominates during the smooth release period (when much of the agent's continuation value depends on future information release) and the agent's IC constraint is loosened, reflected in the growth condition $\dot{x}_t \leq (r + \gamma/K)x_t$. After optimizing the start time of this smooth release period, we find that the period both starts and ends sooner with higher γ, but its length may increase or decrease.

Using similar arguments to those above about the principal's flow payoff, we see that higher proportion of loyal agents q improves the principal's payoff. It has no direct effect on the agent's incentives, and thus $\overline{x}$ is unchanged, as is the length of the smooth release period. But with lower risk of facing a disloyal agent, the principal finds it optimal to start this period (and therefore reach the final jump) sooner. A higher initial knowledge requirement ϕ hurts the principal, and it does so by simply restricting the space of feasible information policies. As with q, it has no direct effect on the agent's incentives, and thus $\overline{x}$ is unchanged, but since $x_0 \geq \phi$ is a binding constraint in the moderate knowledge case, the smooth release period is shortened and its start is delayed. In response, the principal shifts the smooth release period slightly earlier in time (the cost of increasing information prior to t^* affects a smaller interval of times), resulting in an earlier final jump with higher ϕ.

An increase in the discount rate has two opposing effects on the principal's payoff.[19] Supposing that an information policy remains a loyalty test as the discount rate changes, the increase in the discount rate hurts the principal as it increases the weight she places on the flow payoffs earned in the early part of the contract, which are negative since the agent's reputation is low. Unlike the arguments about the direct effect on the principal's flow payoffs of changes in γ, K and q, which apply independent of the information policy and disloyal agent strategy, this claim uses the fact that information policy is a loyalty test, and hence we must consider the effects of an increase in the discount rate on the set of loyalty tests over which the principal optimizes. The accompanying increase in the agent's discount rate makes the disloyal agent more willing to undermine today at the risk of being detected, forgoing continuation utility. Consequently, the disloyal agent's IC constraint is relaxed, which enlarges the set of loyalty tests. Although this effect runs opposite the first one, we show analytically that the first effect dominates — the principal is always made better off when the discount rate increases. The principal adjusts to the agent's increased impatience by making the contract more similar to the unconstrained first best: information release begins later but is faster and ends sooner.

6 Robustness

In this section we study the robustness of our results to various changes in the environment.

6.1 Transfers

Suppose that in addition to regulating the rate of information release, the principal could commit to a schedule of transfers to the agent. Could it use this additional instrument to improve contract performance? To study this possibility, we consider an augmentation of the model with transfers in which the agent cares about money as well as the firm's payoffs. Our main finding is that while the principal may utilize up-front transfers as part of a menu of contracts designed to screen out the disloyal agent, the contract intended for the loyal agent will be qualitatively similar to the optimal contract in the baseline model.

Specifically, if the firm's expected profits (including any spending on transfers) are Π, and the expected net present value of transfers made to the agent are T, then we model the loyal agent's total payoff as

$$U^G = (1-\delta)\Pi + \delta T$$

[19] For comparative statics with respect to the discount rate, we study the normalized payoff $r\Pi$ to eliminate the mechanical diminuation of flow payoffs at all times due to a rise in the discount rate.

while the disloyal agent's payoff is

$$U^B = -(1-\delta)\Pi + \delta T$$

for some $\delta \in (0,1)$. This specification is intended to capture a world in which agents are homogeneous in their relative interest in money versus firm outcomes, but differ in their preferred firm outcomes in the same way as in the baseline model.[20] In particular, the limiting case $\delta = 0$ can be thought of as nesting the baseline model, as transfers cannot be used to usefully align incentives and the optimal contract is as in the baseline.

When $\delta > 0$, cash transfers can be used to improve outcomes by screening out the disloyal agent. In particular, suppose that under a particular operating contract involving no transfers, the disloyal agent receives a payoff of U^B. Note that a buyout contract providing an immediate transfer of $(1-\delta)U^B$ followed by termination yields the same payoff U^B to the disloyal agent, since $\Pi = -(1-\delta)U^B$ for this contract. So the disloyal agent is willing to take the buyout over the operating contract. And offering this contract reduces the total cost to the principal in case the disloyal agent is present, from U^B to $(1-\delta)U^B$. So when transfers are allowed, a menu of contracts offering both an operating contract and a buyout will generally be optimal.

To better understand the optimal form of the operating contract in this environment, assume for simplicity that $\delta \leq 1/2$. This assumption ensures that the loyal agent would rather take the operating contract than a buyout.[21] In this case the optimal menu of contracts involves a buyout calibrated as in the previous paragraph, with the operating contract designed to solve

$$\sup_{x \in \mathbb{X}} \int_0^\infty e^{-rt} x_t (q - (1-q)(1-\delta)(K-1)e^{-\gamma t})\, dt \quad \text{s.t.} \quad \dot{U}_t \leq (r + \gamma/K) U_t,$$

where now the disloyal agent's contribution to the principal's payoff is depressed by a factor

[20] Moreover, this specification captures situations in which disloyal agents have a higher marginal value for money than the principal. This is because the disloyal agent enjoys a transfer from the principal through two channels: the direct benefit of the transfer, as well as the cost to the principal of the transfer.

[21] It can be shown that this result also holds for all δ if $K \leq 2$. When both assumptions fail, then the IC constraint of the loyal agent potentially binds. To ensure that the loyal agent chooses the operating contract, large payments in the distant future can be used to boost the value of the contract to the loyal agent without making it more valuable to the disloyal agent (who discounts future payments more heavily due to termination).

The optimal contract then either involves an operating contract with eventual payments and a buyout, or else a single operating contract with no buyout. In particular screening contracts may not be employed if both K and q are large, so that offering a buyout rarely saves money on disloyal agents but forces the principal to make large payments when the loyal agent is present.

$1-\delta$. Using integration by parts as in Lemma 5, the objective may be rewritten

$$-\left((1-\delta)(1-q)-\frac{1}{K-1}q\right)U_0+\frac{q\gamma}{K-1}\int_0^\infty e^{-rt}U_t\,dt.$$

Notice that conditional on U_0, the objective is exactly as in the problem without transfers, and has the same solution. Equivalently, conditional on x_0 the path of x is unchanged by the presence of transfers. The downweighting of the disloyal agent's payoff by factor $1-\delta$ merely makes a delivery of a given utility U_0 to the bad agent less costly, increasing the optimal size of U_0. So transfers change the operating contract by boosting the optimal initial grant of information, and otherwise leave the optimal path of information release unchanged.

6.2 A continuous state space model

We model uncertainty about the correct task action using a binary state space. This choice affords significant tractability, as the agent's posterior belief process is then one-dimensional. Richer state spaces would typically require a more complex analysis considering the optimal influence of several moments of the agent's beliefs over time. However, when preferences take a quadratic-loss form, payoffs depend only on the second moment of beliefs in general, and much larger state spaces can be accommodated within the existing analysis with minor changes. In this subsection we provide a brief overview of this extension. The main finding is that while, naturally, the details of the optimal information policy may change depending on the state space, the optimal disclosure path x^* is robust and continues to capture the optimal rate at which information should be disclosed.

Concretely, suppose that the state space is the real line $\Omega=\mathbb{R}$, with the agent initially possessing a diffuse prior over Ω. Upon becoming employed, the agent observes an exogenous public signal $s_i \sim N(\omega,\eta^2)$ for some $\eta>0$, inducing time-zero posterior beliefs $\omega|s_i \sim N(s_i,\eta^2)$. (Since preferences will depend only on the second moment of beliefs, the form of the agent's prior and the exogenous signal are not very important and chosen mainly to simplify exposition.)

Principal payoffs depend on the agent's posterior beliefs as follows. At each moment in time the agent undertakes a (perfectly monitored) task action $a_t\in\mathbb{R}$, yielding a flow payoff to the principal of

$$\pi(a_t,\omega)=1-(a_t-\omega)^2/C$$

for some constant $C\geq\eta^2$. (This inequality is without loss, as otherwise the principal would optimally immediately release information to avoid negative expected flow payoffs.) When the agent's posterior beliefs have mean μ_t, the optimal task action is $a_t=\mu_t$. If the variance of the agent's posterior beliefs is σ_t^2, then the induced expected flow payoffs to the principal from this choice of task action are

$$\pi(\sigma_t^2) = 1 - \sigma_t^2/C.$$

We will take this expression to be the principal's flow payoff function given an induced agent belief process. As in the baseline model, by undermining the agent inflicts a flow loss of $K\pi(\sigma_t^2)$ on the principal.

Now define the disclosure path $x_t \equiv \mathbb{E}[1 - \sigma_t^2/C]$ to be the principal's ex ante time-t flow payoff under a given information policy. This process plays the same role as the x process defined in the binary case. The following lemma proves an analogue of Lemma 2 in the baseline model, establishing the set of disclosure paths corresponding to posterior belief processes which can be induced by the principle.

Lemma 7. *For any information policy, x_t is càdlàg, monotone increasing, and $[\phi, 1]$-valued, where $\phi = 1 - \eta^2/C$.*

Now suppose that any disclosure path satisfying the necessary conditions of the lemma above can be implemented by some deterministic information policy, i.e. with the posterior variance process σ^2 a deterministic function. In this case, for the same reasons as in Lemma 3 the principal might as well choose a deterministic information policy. So the problem reduces to the optimal design of x, with an objective function and constraints identical to the baseline model. The solution x^* is thus also identical. The final step is to show that x^* can in fact be implemented by some deterministic information policy.[22] The following lemma establishes this fact.

Lemma 8. *There exists a deterministic information policy whose disclosure path is x^*.*

The proof of the lemma constructs a particular signal process whose disclosure path is x^*. We will describe the process in the low-knowledge case, with the other cases following along similar lines. The principal first releases a discrete normally distributed signal at time zero in the low-knowledge case, with mean ω and variance calibrated to achieve the desired time-zero posterior variance. Then during the smooth disclosure phase $[\underline{t}_L, \overline{t}_L]$, the principal discloses information via a continuous signal process with drift ω obscured by a Brownian

[22]This implementation problem is similar to one faced in Ball (2018), in which the agent's posterior variance process is designed in a relaxed problem and afterward it is proven that there actually exists an information policy inducing the desired sequence of posterior variances. In that paper delayed reports of an exogenous Brownian state process serve as a suitable information policy. As no such process exists in our model, our proof proceeds along different lines.

noise term whose variance declines as the phase progresses. Finally, at time $\bar{t}_L$ the principal reveals ω, eliminating all residual uncertainty.

This extension shows that the details of the optimal signal process will depend on the underlying state space, and in particular Brownian signals become a natural choice when the state space is continuous and preferences are linear in the agent's posterior variance. Nonetheless, the optimal disclosure path is robust to alternative specifications of the state space in which flow payoffs remain linear in a one-dimensional summary statistic of the agent's posterior beliefs.

6.3 An evolving task

We have so far assumed that the underlying state of the world is fixed, and therefore that the agent's information never becomes obsolete. However, in many environments a fluctuating state may be a more natural assumption, capturing environments in which the priorities of the organization change over time. These changes may be driven, for instance, by a changing competitive environment or the development of new product lines. In this subsection we explore how the optimal information policy is impacted by a fluctuating state. We find that in the case the optimal contract of the baseline model begins with a quiet period, the principal can profitably exploit this fluctuation and lets the agent's information deteriorate temporarily. In all other cases, the changing state does not change the optimal disclosure path at all.

Suppose that instead of being held fixed, ω evolves as a Markov chain which transitions between states at rate $\lambda > 0$. This change affects the model by enlarging the set of admissible disclosure paths x. Recall that when ω was fixed, Lemma 2 required that any disclosure path be a monotone function. This is because for a fixed state, the precision of the agent's beliefs cannot drift down on average over time. However, when ω fluctuates, absent disclosures the precision of the agent's beliefs naturally deteriorate toward 0. In particular, if the agent's beliefs are μ_t at some time t, and no further information is disclosed, then the agent's beliefs evolve according to the ODE

$$\dot{\mu}_t = \lambda(1 - 2\mu_t).$$

Accordingly, the disclosure path $x_t = \mathbb{E}|2\mu_t - 1|$ decays at rate 2λ over any time interval when no information is disclosed.

The model can therefore be adapted to accommodate the changing state by requiring that x, rather than being a monotone function, satisfy the weaker condition that $e^{2\lambda t}x_t$ be monotone. In the low knowledge case, when the lower bound constraint was non-binding for a fixed state, this relaxation of the problem does not change the optimal disclosure path at all. However, in the moderate and high knowledge cases, the principal takes advantage of the changing state to initially degrade the precision of the agent's information over time. A consequence is that the optimal length of the gradual disclosure period lengthens and the agent graduates to trusted status later.[23] Figure 4 illustrates how the optimal contract in the moderate-knowledge case is impacted by a changing state.

Note that the information policy implementing a given disclosure path changes somewhat under a changing state. In particular, the agent must be kept abreast of changes in the state with at least some probability in order to maintain a constant disclosure level. So an agent who has reached trusted status will receive an immediate report whenever the state has changed. Further, during the gradual disclosure phase status updates serve a dual role — they both inform the agent in case he was previously off-track, and update him about changes in the state. Of course, the agent is always left uncertain as to the motivation for a particular status update!

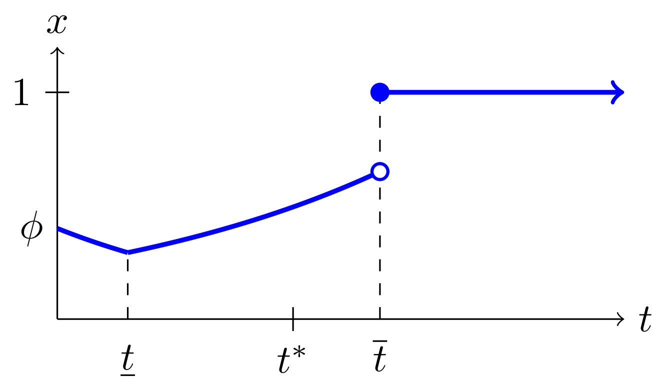

Figure 4: The optimal disclosure path in the moderate knowledge case for an evolving state.

6.4 Replacing the agent

Our model assumes that the agent is irreplaceable, so that if the agent is discovered to be disloyal and fired the principal loses all future payoffs from performance of the agent's task. This is a reasonable assumption for settings in which the agent is a specialist and the task payoff is relative to some performance benchmark for a replacement. However, in other settings it may be reasonable to believe that the principal can replace a disloyal agent with someone of similar skills, perhaps at a cost. Our analysis can be adapted to accommodate this scenario with only minor modifications. We outline the required changes, and the associated

[23]To understand these facts, note that extending the length of the gradual disclosure phase, and the corresponding final disclosure time, incurs the same marginal cost at the end of the contract for all λ, but provides increasing savings at the start of the contract due to lower average disclosure levels when λ is larger.

implications for optimal disclosure, here. We will consider two alternative specifications: in one, the principal's post-termination payoff does not enter the disloyal agent's; in the other, it does. We find that in the first specification, the principal's solution is exactly the same as in the baseline model, while in the second, the optimal disclosure curve is quantitatively similar and differs mainly in a tighter constraint on the rate of information release during the gradual disclosure phase.

First suppose the agent cares only about firm outcomes while employed. Then the firm's continuation value following termination does not matter for his incentives, and in particular the game is no longer zero-sum. However, this turns out not to impact the optimality of offering a single contract. For it continues to be the case that any contract can be modified to create a loyalty test without changing the payoff to the disloyal agent while (weakly) improving the payoff to the principal. So without loss the principal offers only loyalty tests. Also, with a positive post-termination payoff the principal optimally recommends undermining at all times under any loyalty test, so all contracts optimally chosen by the principal must discover the disloyal agent at the same rate. Thus the contribution to the principal's profits from the post-termination option is the same for all contracts. The game is then effectively zero-sum with a disloyal agent, implying optimality of a single contract.

In light of the previous discussion, the ability to collect a continuation payoff $\Pi > 0$ following termination changes the principal's problem only by adding on a constant term $(1-q)\frac{\gamma}{r+\gamma}\Pi$ to the principal's payoff, which critically does not depend on the choice of x. So the contract design problem is completely unchanged by the possibility of replacement, and the optimal information policy will be the same. If Π is determined endogenously, say as the value of the original problem minus a search cost, it can be calculated by solving a simple fixed-point problem. Say V is the expected profit to the principal of the current agent under an optimal contract, and that new workers are drawn from the same distribution as the original, after incurring a search cost $\zeta > 0$. (We will suppose ζ is sufficiently small that the principal optimally replaces the agent.) Then Π satisfies

$$\Pi = V + (1-q)\frac{\gamma}{r+\gamma}(\Pi - \zeta),$$

which can be solved for Π to obtain an expression for the lifetime value of the position.

Alternatively, suppose the agent's preferences are over the lifetime value of the position, regardless of who holds it. In this case the game continues to be zero-sum, and the principal continues to optimally offer only a single contract which is a loyalty test. However, now the disloyal agent's incentive constraints for undermining tighten, reflecting the added value of holding off replacement through feigning loyalty. Specifically, the continuation value to the disloyal agent at a given time t is now

$$U_t = \int_t^\infty e^{-(r+\gamma)t}[(K-1)x_t - \gamma\Pi]\, dt,$$

which accounts for the arrival of the continuation profit Π to the principal upon discovery and termination. Letting $U_t^\dagger = U_t + \gamma\Pi/(\gamma + r)$ represent the gross continuation payoff to the agent before accounting for the continuation value to the principal, the principal's continuation payoff function is, up to a constant not depending on x,

$$-\left(1 - \frac{K}{K-1}q\right) U_0^\dagger + \frac{q\gamma}{K-1}\int_0^\infty e^{-rt} U_t^\dagger\, dt,$$

analogously to the baseline model. And the constraint $x \geq \phi$ may be written

$$\dot{U}_t^\dagger \leq (r+\gamma)U_t^\dagger - (K-1)\phi,$$

again as in the baseline model. However, appropriately modifying the proof of Lemma 4 yields the IC constraint

$$\dot{U}_t^\dagger \leq (r+\gamma/K)U_t^\dagger - \frac{r\gamma}{r+\gamma}\frac{K-1}{K}\Pi,$$

which is tighter than the analogous constraint in the baseline model. The bounds on $U^\dagger$ may also tighten — while its upper bound is $\overline{U}^\dagger = (K-1)/(r+\gamma)$ as before, it must satisfy the potentially tighter lower bound

$$\underline{U}^\dagger = \max\left\{\phi\overline{U}^\dagger, \left(\frac{\gamma}{r+\gamma} - \frac{\gamma/K}{r+\gamma/K}\right)\Pi\right\}.$$

This bound ensures that the upper bound on $\dot{U}^\dagger$ is non-negative in both constraints for all admissible levels of U. The new second term in the bound reflects the fact that when the principal's continuation payoff is high, and the current level of disclosure is low, the agent may prefer to remain loyal rather than trigger turnover even absent future disclosures.

These changes to the contracting problem yield several changes to the optimal contract. First, the growth rate of x during the gradual disclosure phase is slower than without replacement. Second, the size of the disclosure when the agent becomes trusted becomes smaller. Third, the optimal size of initial disclosure will grow, meaning the size of the quiet period shrinks and more information structures fall into the low-knowledge case. Otherwise, the optimal contract is qualitatively similar to the no-replacement case, and is derived similarly.

Note that all changes in the contract are driven by changes in the incentives for disloyal agents to feign loyalty. Continuation payoffs do not directly enter the principal's objective for the current agent for any given information policy. And replacement does not change the undermining policies which are optimally induced, as even without replacement the principal already optimally induces as much undermining as possible. So while post-termination options potentially boost the principal's lifetime profits, for a given agent they only restrict the set of loyalty tests which can be offered.

7 Conclusion

We study optimal information disclosure in a dynamic principal-agent problem when knowledge of an underlying state of the world can be used either to aid the principal (through better job performance) or to harm her (through leaks or sabotage). Some fraction of agents oppose the interests of the principal, and can undermine her in a gradual and imperfectly detectable manner. Ideally, the principal would like to keep the agent in the dark until she can ascertain loyalty with sufficient confidence by detecting and tracing the source of any undermining; only after such a screening phase would the principal inform the agent about the state. However, such an information structure gives a disloyal agent strong incentives to feign loyalty early on by refraining from undermining, negating the effectiveness of the scheme.

We show that the principal optimally disburses information as quickly as possible while dissuading disloyal agents from feigning loyalty at any time, maximizing the rate of detection of disloyalty. This optimal information policy can be implemented as an inconclusive contradictory news process, which gradually increases the precision of the agent's information up to a deterministic terminal date. At this date the agent is deemed trusted and brought into an inner circle with a final, fully revealing report of the state.

Our model assumes that agents are either fully aligned or totally opposed to the principal's long-run interests. This restriction allows us to focus on the principal's problem of "plugging leaks" by detecting and screening out disloyal agents. However, in some situations it may be realistic to suppose that even disloyal agents can be "converted" into loyal ones through sufficient monetary compensation, team-building exercises, or other incentives. We leave open as an interesting direction for future work the broader question of when a principal might prefer to screen versus convert agents who face temptations to undermine her.

Our model also assumes that the nature of the underlying task is exogenous, and that the agent's loyalty is independent of the task. Relaxing these assumptions would allow one to study how a firm or organizations should proceed when facing severe challenges: while some alternatives might be mild and uncontroversial, other more drastic course corrections might carry a risk of alienating a subset of agents. We find this question, of how firms should choose among alternatives with different implications for agent loyalty, to be worth further exploration.

References

Albuquerque, Rui and Hugo A. Hopenhayn (2004). Optimal Lending Contracts and Firm Dynamics. In: *The Review of Economic Studies* 71.2, pp. 285–315.

Associated Press (2018). Police are using fake Amazon packages to catch porch pirates. In: *CBS News*. URL: https://www.cbsnews.com/news/package-thieves-being-caught-with-decoy-amazon-deliveries.

Au, Pak Hung (2015). Dynamic information disclosure. In: *RAND Journal of Edconomics* 46.4, pp. 791–823.

Ball, Ian (2018). Dynamic Information Provision: Rewarding the Past and Guiding the Future. Working paper.

Basu, Pathikrit (2018). Dynamic Bayesian Persuasion with a Privately Informed Receiver. Working paper.

Capasso, Vincenzo and David Bakstein (2005). An introduction to continuous-time stochastic processes. In: *Birkhäuser Boston*, p. 1.

Che, Yeon-Koo and Konrad Mierendorff (2017). Optimal sequential decision with limited attention. Working paper.

Cranley, Ellen (2018). A senior White House aide is reportedly sharing fake stories with staffers to identify who's leaking to the press. In: *Business Insider*. URL: https://www.businessinsider.com/trump-leaks-leakers-2018-6.

Daley, Brendan and Brett S Green (2018). Bargaining and News. Working paper.

Ely, Jeffrey (2017). Beeps. In: *The American Economic Review* 107.1, pp. 31–53.

Ely, Jeffrey, Alexander Frankel, and Emir Kamenica (2015). Suspense and surprise. In: *Journal of Political Economy* 123.1, pp. 215–260.

Ely, Jeffrey and Martin Szydlowski (2018). Moving the Goalposts. In: *Journal of Political Economy*, forthcoming.

Fudenberg, Drew and Luis Rayo (2018). Training and Effort Dynamics in Apprenticeship.

Guo, Yingni and Eran Shmaya (2018). The Interval Structure of Optimal Disclosure. In: *Econometrica*, forthcoming.

Haynes, John Earl and Harvey Klehr (2006). *Early Cold War Spies: The Espionage Trials That Shaped American Politics*. Cambridge University Press.

Isaac, Mike (2016). Facebook 'Trending' List Skewed by Individual Judgment, Not Institutional Bias. In: *The New York Times*. URL: https://www.nytimes.com/2016/05/21/technology/facebook-trending-list-skewed-by-individual-judgment-not-institutional-bias.

Kamenica, Emir and Matthew Gentzkow (2011). Bayesian persuasion. In: *The American Economic Review* 101.6, pp. 2590–2615.

Kinery, Emma (2017). As Jeff Sessions Announces Leak Crackdown, Here are 12 Major Leaks of the Trump Era. In: *USA Today*. URL: https://www.usatoday.com/story/news/politics/onpolitics/2017/08/04/here-12-major-leaks-donald-trump-era/539674001.

Kolodny, Lora (2018). Elon Musk Emails Employees about 'Extensive and Damaging Sabotage' by Employee. In: *CNBC*. URL: https://www.cnbc.com/2018/06/18/elon-musk-email-employee-conducted-extensive-and-damaging-sabotage.html.

Kolotilin, Anton (2018). Optimal information disclosure: A linear programming approach. In: *Theoretical Economics* 13.2, pp. 607–635.

Kolotilin, Anton et al. (2017). Persuasion of a privately informed receiver. In: *Econometrica* 85.6, pp. 1949–1964.

Liptser, Robert S and Albert N Shiryaev (2013). *Statistics of random processes: I. General theory*. Vol. 5. Springer Science & Business Media.

Macintyre, Ben (2010). *Operation Mincemeat*. Bloomsbury.

Orlov, Dmitry, Andrzej Skrzypacz, and Pavel Zryumov (2018). Persuading the Principal To Wait. In: *Journal of Political Economy*, forthcoming.

Peterson, Hayley (2018). Amazon plants fake packages in delivery trucks as part of an undercover ploy to 'trap' drivers who are stealing. In: *Business Insider*. URL: https://www.businessinsider.com/amazon-sets-traps-for-drivers-2018-9.

Rayo, Luis and Luis Garicano (2017). Relational Knowledge Transfers. In: *The American Economic Review* 107.9, pp. 2695–2730.

Rothschild, Neal (2018). The 16 Juiciest White House Leaks of 2018. In: *Axios*. URL: https://www.axios.com/white-house-leaks-president-trump-2018-92899c04-0dfe-4b64-bc42-18fcdb09e28f.

Thomas, Jonathan and Timothy Worrall (1994). Foreign Direct Investment and the Risk of Expropriation. In: *Review of Economic Studies* 61.1, pp. 81–108.

Thompson, Nicholas and Fred Vogelstein (2018). Inside the Two Years That Shook Facebook - and the World. In: *Wired*. URL: https://www.wired.com/story/inside-facebook-mark-zuckerberg-2-years-of-hell.
Watson, Joel (1999). Starting small and renegotiation. In: *Journal of Economic Theory* 85.1, pp. 52–90.
— (2002). Starting small and commitment. In: *Games and Economic Behavior* 38.1, pp. 176–199.
Zhong, Weijie (2018). Optimal Dynamic Information Acquisition. Working paper.
Ziegler, Charles A and David Jacobson (1995). *Spying without spies: Origins of America's secret nuclear surveillance system*. Greenwood Publishing Group.

A Proofs

A.1 Technical lemmas

For the lemmas in this subsection, fix a deterministic information policy μ with associated disclosure path x.

Given any undermining policy b, define U^b to be the disloyal agent's continuation value process under b, conditional on the agent remaining employed. By definition,

$$U_t^b = \int_t^\infty \exp\left(-r(s-t) - \gamma \int_t^s b_u \, du\right)(Kb_s - 1)x_s \, ds$$

for all t. This function is absolutely continuous with a.e. derivative

$$\frac{dU^b}{dt} = (r + \gamma b_t)U_t^b - (Kb_t - 1)x_t.$$

Note that the rhs is bounded below by $f(U_t^b, t)$, where

$$f(u,t) \equiv \min\{(r+\gamma)u - (K-1)x_t, ru + x_t\}.$$

Lemma A.1. *Suppose $g(u,t)$ is a function which is strictly increasing in its first argument. Fix $T \in \mathbb{R}_+$, and suppose there exist two absolutely continuous functions u_1 and u_2 on $[0,T]$ such that $u_1(T) \geq u_2(T)$ while $u_1'(t) = g(u_1(t),t)$ and $u_2'(t) \geq g(u_2(t),t)$ on $[0,T]$ a.e. Then:*

1. *$u_1 \geq u_2$.*
2. *If in addition $u_1(T) = u_2(T)$ and $u_2'(t) = g(u_2(t),t)$ on $[0,T]$ a.e., then $u_1 = u_2$.*

3. If in addition $u_2'(t) > g(u_2(t), t)$ on some positive-measure subset of $[0, T]$, then $u_1(0) > u_2(0)$.

Proof. Define $\Delta(t) \equiv u_2(t) - u_1(t)$, and suppose by way of contradiction that $\Delta(t_0) > 0$ for some $t_0 \in [0, T)$. Let $t_1 \equiv \inf\{t \geq t_0 \ : \ \Delta(t_0) \leq 0\}$. Given continuity of Δ, it must be that $t_1 > t_0$. Further, $t_1 \leq T$ given $u_1(T) \geq u_2(T)$. And by continuity $\Delta(t_1) = 0$. But also by the fundamental theorem of calculus

$$\Delta(t_1) = \Delta(t_0) + \int_{t_0}^{t_1} \Delta'(t)\, dt.$$

Now, given that $\Delta(t) > 0$ on (t_0, t_1), it must be that

$$\Delta'(t) = u_2'(t) - u_1'(t) \geq g(u_2(t), t) - g(u_1(t), t) > 0$$

a.e. on (t_0, t_1) given that g is strictly increasing in its first argument. Hence from the previous identity $\Delta(t_1) > \Delta(t_0)$, a contradiction of $\Delta(t_1) = 0$ and $\Delta(t_0) > 0$. So it must be that $\Delta \leq 0$, i.e. $u_2 \leq u_1$.

Now, suppose further that $u_1(T) = u_2(T)$ and $u_2'(t) = g(u_2(t), t)$ on $[0, T]$ a.e. Trivially $u_2(T) \geq u_1(T)$ and $u_1'(t) \geq g(u_1(t), t)$ on $[0, T]$ a.e., so reversing the roles of u_1 and u_2 in the proof of the previous part establishes that $u_1 \leq u_2$. Hence $u_1 = u_2$.

Finally, suppose that $u_2'(t) > g(u_2(t), t)$ on some positive-measure subset of $[0, T]$. We have already established that $u_1 \geq u_2$. Assume by way of contradiction that $u_1(0) = u_2(0)$. Let $t_0 \equiv \inf\{t \ : \ u_1'(t), u_2'(t) \text{ exist and } u_1'(t) = g(u_1(t), t), u_2'(t) > g(u_2(t), t)\}$. The fact that $u_1'(t) = g(u_1(t), t)$ a.e. while $u_2'(t) > g(u_2(t), t)$ on a set of strictly positive measure ensures that $t_0 \in [0, T)$. By construction $u_2'(t) = g(u_2(t), t)$ on $[0, t_0)$, in which case a minor variant of the argument used to prove the first two parts of this lemma shows that $u_1 = u_2$ on $[0, t_0]$. In particular, $u_1(t_0) = u_2(t_0)$. But then by construction $\Delta'(t_0)$ exists and

$$\Delta'(t_0) = u_2'(t_0) - u_1'(t_0) > g(u_2(t_0), t_0) - g(u_1(t_0), t_0) = 0.$$

In particular, for sufficiently small $t > t_0$ it must be that $\Delta(t) > 0$, which is the desired contradiction. □

Lemma A.2. *Given any $T \in \mathbb{R}_+$ and $\overline{u} \in \mathbb{R}$, there exists a unique absolutely continuous function u such that $u(T) = \overline{u}$ and $u'(t) = f(u(t), t)$ on $[0, T]$ a.e.*

Proof. Suppose first that x is a simple function; that is, x takes one of at most a finite number of values. Given that x is monotone, this means it is constant except at a finite set of jump points $D = \{t_1, ..., t_n\}$, where $0 < t_1 < ... < t_n < T$. In this case $f(u, t)$ is uniformly

Lipschitz continuous in u and is continuous in t, except on the set D. Further, f satisfies the bound $|f(u,t)| \leq (K-1)+(r+\gamma)|u|$. Then by the Picard-Lindelhöf theorem there exists a unique solution to the ODE between each t_k and t_{k+1} for arbitrary terminal condition at t_{k+1}. Let u^0 be the solution to the ODE on $[t_n, T]$ with terminal condition $u^0(T) = \overline{u}$, and construct a sequence of functions u^k inductively on each interval $[t_{n-k}, t_{n-k+1}]$ by taking $u^k(t_{n-k+1}) = u^{k-1}(t_{n-k+1})$ to be the terminal condition for u^k. Then the function u defined by letting $u(t) = u^k(t)$ for $t \in [t_{n-k}, t_{n-k+1}]$ for $k = 0, ..., n$, with $t_0 = 0$ and $t_{n+1} = T$, yields an absolutely continuous function satisfying the ODE everywhere except on the set of jump points D.

Now consider an arbitrary x. As x is non-negative, monotone, and bounded, there exists a increasing sequence of monotone simple functions x^n such that $x^n \uparrow x$ uniformly on $[0,T]$. Further, there exists another decreasing sequence of monotone simple functions $\widetilde{x}^m$ such that $\widetilde{x}^m \downarrow x$ uniformly on $[0,T]$. For each n, let u^n be the unique solution to the ODE $u'(t) = f^n(u,t)$ with terminal condition $u(T) = \overline{u}$, where

$$f^n(u,t) \equiv \min\{(r+\gamma)u - (K-1)\widetilde{x}^n_t, ru + x^n_t\}.$$

(The arguments used for the case of x simple ensure such a solution exists.) By construction $f^n(u,t)$ is increasing in n for fixed (u,t), hence $du^m/dt \geq f^n(u^m(t),t)$ for every $m > n$. Also observe that each f^n is strictly increasing in its first argument, hence Lemma A.1 establishes that the sequence u^n is pointwise decreasing. Then the pointwise limit of u^n as $n \to \infty$ is well-defined (though possibly infinite), a function we will denote u^∞. Further, $f^n(u,t) \leq ru + 1$, hence by Grönwall's inequality $u^n(t) \geq -1/r + (\overline{u} + 1/r)\exp(r(t-T))$. Thus the u^n are uniformly bounded below on $[0,T]$, meaning u^∞ is finite-valued.

Now define an absolutely continuous function u^* on $[0,T]$ pointwise by setting

$$u^*(t) = \overline{u} - \int_t^T f(u^\infty(s), s)\, ds.$$

Our goal is to show that u^* is the desired solution to the ODE. Once that is established, Lemma A.1 ensures uniqueness given that f is strictly increasing in its first argument.

Use the fundamental theorem of calculus to write

$$u^n(t) = \overline{u} - \int_t^T \frac{du^n}{ds}\, ds = \overline{u} - \int_t^T f^n(u^n(s), s)\, ds,$$

and take $n \to \infty$ on both sides. Note that $|f^n(u^n(s), s)| \leq (K-1) + (r+\gamma)|u^n(s)|$, and recall that the sequence u^n is pointwise decreasing and each u^n satisfies the lower bound

$u^n(t) \geq -1/r + (\overline{u} + 1/r)\exp(r(t-T))$. Hence a uniform lower bound on the $u^n(t)$ is $\min\{-1/r, \overline{u}\}$. So $|u^n|$ may be bounded above by $|u^n(s)| \leq \max\{1/r, |\overline{u}|, |u^1(s)|\}$. As u^1 is continuous and therefore bounded on $[0,T]$, this bound implies that $(K-1) + (r+\gamma)\max\{1/r, |\overline{u}|, |u^1(s)|\}$ is an integrable dominating function for the sequence $|f^n(u^n(s), s)|$. The dominated convergence theorem then yields

$$u^\infty(t) = \overline{u} - \int_t^T \lim_{n\to\infty} f^n(u^n(s), s)\, ds = \overline{u} - \int_t^T f(u^\infty(s), s)\, ds = u^*(t).$$

So u^∞ and u^* are the same function. Meanwhile, differentiating the definition of u^* yields

$$\frac{du^*}{dt} = f(u^\infty(t), t) = f(u^*(t), t).$$

Then as $u^*(T) = \overline{u}$, u^* is the desired solution to the ODE. □

Lemma A.3. *An undermining policy b^* maximizes the disloyal agent's payoff under μ iff* $\frac{dU^{b^*}}{dt} = f(U_t^{b^*}, t)$ *a.e.*

Proof. Suppose first that b^* is an undermining policy whose continuation value process U^{b^*} satisfies $\frac{dU^{b^*}}{dt} = f(U_t^{b^*}, t)$ a.e. Let $x_\infty \equiv \lim_{t\to\infty} x_t$, which must exist given that x is an increasing function. We first claim that $\limsup_{t\to\infty} U_t^{b^*} = (K-1)x_\infty/(r+\gamma)$. First, given that $x_t \leq x_\infty$ for all time it must also be that $U_t^{b^*} \leq (K-1)x_\infty/(r+\gamma)$, as the latter is the maximum achievable payoff when $x_t = x_\infty$ for all time. Suppose $\overline{U}_\infty \equiv \limsup_{t\to\infty} U_t^{b^*} < (K-1)x_\infty/(r+\gamma)$. Then $\limsup_{t\to\infty} f(U_t^{b^*}, t) \leq (r+\gamma)\overline{U}_\infty - (K-1)x_\infty < 0$. This means that for sufficiently large t, dU^{b^*}/dt is negative and bounded away from zero, meaning U^{b^*} eventually becomes negative. This is impossible, so it must be that $\overline{U}_\infty = (K-1)x_\infty/(r+\gamma)$.

Now fix an arbitrary undermining policy b with continuation value process U^b. Suppose first that at some time T, $U_T^{b^*} \geq U_T^b$. Then by Lemma A.1, $U_t^{b^*} \geq U_t^b$ for all $t < T$, and in particular the lifetime value to the disloyal agent of b^* is at least as high as that of b. So consider the remaining possibility that $U_t^{b^*} < U_t^b$ for all t. Let $\Delta U \equiv U_0^b - U_0^{b^*} > 0$. As $f(u,t)$ is strictly increasing in u and $dU^b/dt \geq f(U_t^b, t)$ for all time, it must therefore be the case that $dU^b/dt > dU^{b^*}/dt$ a.e. Thus

$$U_t^b = U_0^b + \int_0^t \frac{dU^b}{ds}\, ds > U_0^b + \int_0^t \frac{dU^{b^*}}{ds}\, ds = U_t^{b^*} + \Delta U$$

for all t. But then

$$\limsup_{t\to\infty} U_t^b \geq \overline{U}_\infty + \Delta U = (K-1)x_\infty/(r+\gamma) + \Delta U,$$

and as $\Delta U > 0$ it must therefore be that $U_t^b > (K-1)x_\infty/(r+\gamma)$ for some t sufficiently large. This contradicts the bound $U_t^b \leq (K-1)x_\infty/(r+\gamma)$ for all t, hence it cannot be that $U_t^{b^*} < U_t^b$ for all time. This proves that b^* is an optimal undermining policy for the disloyal agent.

In the other direction, fix an undermining policy b such that $\frac{dU^b}{dt} > f(U_t^b, t)$ on some positive-measure set of times. Then there must exist a time T such that the equality fails on a positive-measure subset of $[0, T]$. Let U^* be the unique solution to the ODE $du/dt = f(u, t)$ on $[0, T]$ with terminal condition $u(T) = U_T^b$. Such a solution is guaranteed to exist by Lemma A.2. And by Lemma A.1, $U_0^* > U_0^b$. So define an undermining profile b^* by setting $b_t^* = b_t$ for $t \geq T$, and

$$b_t^* = \begin{cases} 1, & Kx_t \geq \gamma U_t^* \\ 0, & Kx_t < \gamma U_t^* \end{cases}$$

for $t < T$. Note that the ODE characterizing U^* may be written

$$\frac{dU^*}{dt} = (r + \gamma b_t^*)U_t^* - (Kb_t^* - 1)x_t.$$

Multiplying through by $\exp\left(-rt - \gamma \int_0^t b_s^*\, ds\right)$ allows the ODE to be written as a total differential:

$$\frac{d}{dt}\left(\exp\left(-rt - \gamma\int_0^t b_s^*\, ds\right)U_t^*\right) = -\exp\left(-rt - \gamma\int_0^t b_s^*\, ds\right)(Kb_t^* - 1)x_t.$$

Integrating from 0 to T yields

$$U_0^* = \int_0^T \exp\left(-rt - \gamma\int_0^t b_s^*\, ds\right)(Kb_t^* - 1)x_t\, dt + \exp\left(-rT - \gamma\int_0^T b_t^*\, dt\right)U_T^*.$$

Since $U_T^* = U_T^b = U_T^{b^*}$, this means that U_0^* is the lifetime value to the disloyal agent of undermining policy b^*. So b is not an optimal undermining policy. $\square$

A.2 Proof of Lemma 1

First note that given any menu of contracts, it is optimal for an agent to accept some contract. For the agent can always secure a value at least equal to their outside option of 0 by choosing their myopically preferred task action at each instant.

Suppose without loss that the disloyal agent prefers contract 1 to contract 2. If the loyal agent also prefers contract 1, then the principal's payoff from offering contract 1 alone is the same as from offering both contracts. On the other hand, suppose the loyal agent prefers contract 2. In this case if the principal offers contract 2 alone, her payoff when the agent is loyal is unchanged. Meanwhile by assumption the disloyal agent's payoff decreases when accepting contract 1 in lieu of contract 2. As the disloyal agent's preferences are in direct opposition to the principal's, the principal's payoff when the agent is disloyal must (weakly) increase by eliminating contract 1.

A.3 Proof of Lemma 2

We first prove the "only if" direction. Let x be the disclosure path for some information policy μ. By assumption, μ is càdlàg and thus so is x. Let $\mathcal{F}$ denote the filtration generated by μ. By the martingale property, for all $0 \leq s < t$, $2\mu_s - 1 = \mathbb{E}[2\mu_t - 1|\mathcal{F}_s]$, and by the triangle inequality, $|2\mu_s - 1| \leq \mathbb{E}[|2\mu_t - 1||\mathcal{F}_s]$. Taking expectations at time 0 and using the law of iterated expectations, $\mathbb{E}[|2\mu_s - 1|] \leq \mathbb{E}[|2\mu_t - 1|]$, and thus by definition $x_s \leq x_t$, so x is monotone increasing. By the same logic, we have $\phi \leq x_t$ for all t, and since $\mu_t \in [0, 1]$ for all t, $x_t \leq 1$. This concludes the proof for this direction.

Next, consider the "if" direction. Suppose x is càdlàg, monotone increasing and $[\phi, 1]$-valued; we show that there exists an information policy μ for which x is the disclosure path. We make use of Capasso and Bakstein (2005, Theorem 2.99).

Define $\mathcal{T} \equiv \{0-\} \cup [0, \infty)$ and $x_{0-} \equiv \phi$ and adopt the convention that $0- < t$ for all $t \geq 0$. We construct functions $q(t_1, t_2, a, Y)$ defined for $t_1, t_2 \in \mathcal{T}$ with $t_1 < t_2$, $a \in \mathbb{R}$ and $Y \in \mathcal{B}(\mathbb{R})$, which will be the desired transition probabilities for μ.

Consider any pair of times (t_1, t_2) with $0- \leq t_1 < t_2$. We begin by constructing q for cases where Y is a singleton. Define $q(t_1, t_2, a, \{y\})$ as follows. If $x_{t_1} = 1$, set $q(t_1, t_2, 0, \{0\}) = q(t_1, t_2, 1, \{1\}) = 1$ and set $q(t_1, t_2, a, \{y\}) = 0$ for all other (a, y) pairs. If $x_{t_2} = 0$, set $q(t_1, t_2, a, \{y\}) = 1$ if $a = y = \frac{1}{2}$ and set $q(t_1, t_2, a, \{y\}) = 0$ for all other (a, y) pairs. Next consider $x_{t_1} < 1$ and $x_{t_2} > 0$. For all $t \in \mathcal{T}$, define

$$m_t^{\pm} \equiv \frac{1 \pm x_t}{2}.$$

For $a \in \{m_{t_1}^+, m_{t_1}^-\}$, set $q(t_1, t_2, a, \{m_{t_2}^+\}) = \frac{a - m_{t_2}^-}{m_{t_2}^+ - m_{t_2}^-}$ and $q(t_1, t_2, a, \{m_{t_2}^-\}) = \frac{m_{t_2}^+ - a}{m_{t_2}^+ - m_{t_2}^-}$, making use of the fact that $m_{t_2}^+ - m_{t_2}^- > 0$ since $x_{t_2} > 0$. Set $q(t_1, t_2, a, \{y\}) = 0$ for all other (a, y) pairs.

Next, we extend this construction to all Borel sets Y. For any Borel set $Y \in \mathcal{B}(\mathbb{R})$, define

$$q(t_1, t_2, a, Y) \equiv \sum_{y \in Y \cap \{m_{t_2}^+, m_{t_2}^-\}} q(t_1, t_2, a, \{y\}). \tag{A.1}$$

Hence for all $a \in \mathbb{R}$, for all $0- \leq t_1 < t_2$, $q(t_1, t_2, a, \cdot)$ is a probability measure on $\mathcal{B}(\mathbb{R})$. Since $q(t_1, t_2, a, Y)$ is nonzero for at most two values of a, $q(t_1, t_2, \cdot, Y)$ is a Borel measurable function for all $\mathcal{B}(\mathbb{R})$ and all $0- \leq t_1 < t_2$. It remains to show that q satisfies the Chapman-Kolmogorov equation, which here specializes to

$$q(t_1, t_2, a, Y) = \sum_{z \in \{m^+_{t'}, m^-_{t'}\}} (q(t_1, t', a, \{z\}) q(t', t_2, z, Y)) \qquad \forall t_1 < t' < t_2. \tag{A.2}$$

By (A.1), and the fact that by construction the $q(t_1, t_2, a, \cdot)$ are probability measures for all t_1, t_2, a in the domain, it suffices to consider singletons $Y \in \{\{m^+_{t_2}\}, \{m^-_{t_2}\}\}$. Moreover, whenever $a \notin \{m^+_{t_1}, m^-_{t_1}\}$, by construction $q(t_1, t_2, a, Y) = 0$ for all t_1, t_2, a in the domain so both sides of (A.2) vanish and equality holds; we thus consider $a \in \{m^+_{t_1}, m^-_{t_1}\}$. Note that if $m^+_{t'} = m^-_{t'}$, then it must be that these both equal $a = 1/2$, and thus (A.2) reduces to $q(t_1, t_2, 1/2, \{m^+_{t_2}\}) = q(t', t_2, 1/2, \{m^+_{t_2}\})$. By construction, either (a) $x_{t_2} = 0$ so $m^+_{t_2} = 1/2$ and both sides of the equation are 1, or (b) $x_{t_2} > 0$ so $m^+_{t_2} > m^-_{t_2}$ and both sides are $\frac{1/2 - m^-_{t_2}}{m^+_{t_2} - m^-_{t_2}} = 1/2$. We conclude that (A.2) holds if $m^+_{t'} = m^-_{t'}$. Now whenever $m^+_{t'} \neq m^-_{t'}$, we have $m^+_{t_2} \neq m^-_{t_2}$ and

$$\begin{aligned}
\sum_{z \in \{m^+_{t'}, m^-_{t'}\}} \left(q(t_1, t', a, \{z\}) q(t', t_2, z, \{m^+_{t_2}\})\right) &= \frac{a - m^-_{t'}}{m^+_{t'} - m^-_{t'}} \frac{m^+_{t'} - m^-_{t_2}}{m^+_{t_2} - m^-_{t_2}} + \frac{m^+_{t'} - a}{m^+_{t'} - m^-_{t'}} \frac{m^-_{t'} - m^-_{t_2}}{m^+_{t_2} - m^-_{t_2}} \\
&= \frac{a - m^-_{t_2}}{m^+_{t_2} - m^-_{t_2}} \\
&= q(t_1, t_2, a, \{m^+_{t_2}\}),
\end{aligned}$$

as desired. By Capasso and Bakstein (2005, Theorem 2.99), there exists a unique Markov chain μ with q as its transition probabilities. By construction, μ is a martingale and x is its disclosure path.

A.4 Proof of Lemma 3

Let $x_t = \mathbb{E}|2\mu_t - 1|$ be the disclosure path induced by μ. The construction in the proof of Lemma 2 yields a deterministic information policy inducing x, so at least one such policy exists. Now fix a deterministic information policy μ' with disclosure path x. Note that the loyal agent's payoff under μ is

$$\mathbb{E} \int_0^\infty e^{-rt} |2\mu_t - 1| \, dt = \int_0^\infty e^{-rt} x_t \, dt = \int_0^\infty e^{-rt} |2\mu'_t - 1| \, dt,$$

so his payoff is the same under μ' and μ. As for the disloyal agent, his payoff under μ and any undermining policy b which is a deterministic function of time is

$$\begin{aligned}
&\mathbb{E}\int_0^\infty \exp\left(-rt-\gamma\int_0^t b_s\,ds\right)(Kb_t-1)|2\mu_t-1|\,dt\\
=&\int_0^\infty \exp\left(-rt-\gamma\int_0^t b_s\,ds\right)(Kb_t-1)x_t\,dt\\
=&\int_0^\infty \exp\left(-rt-\gamma\int_0^t b_s\,ds\right)(Kb_t-1)|2\mu'_t-1|\,dt,
\end{aligned}$$

hence is identical to his payoff choosing b under μ'. And as $|2\mu'_t-1|$ is a deterministic function, the disloyal agent's maximum possible payoff under μ' can be achieved by an undermining policy which is deterministic in time. It follows that the disloyal agent's maximum payoff under μ across *all* (not necessarily deterministic) undermining policies is at least as large as his maximum payoff under μ'. Therefore the principal's payoff under μ must be weakly lower than under μ'.

A.5 Proof of Proposition 1

Fix μ which is not a loyalty test. Without loss we take μ to be deterministic, as otherwise Lemma 3 ensures that we may pass to a payoff-improving deterministic policy. In this case it is also without loss to restrict the agent's strategy space to undermining policies which are deterministic functions of time. We will let $x_t=|2\mu_t-1|$ be the disclosure path associated with μ.

Fix an optimal action policy b^* for the disloyal agent under x, and let U^* be its associated continuation utility path. Let $T\equiv\inf\{t\ :\ Kx_t\leq\gamma U^*_t\}$. If $T=\infty$, then $dU^*/dt=(r+\gamma)U^*_t-(K-1)x_t$ a.e. by Lemma A.3, and hence $b^*=1$ a.e. The optimality of such a policy contradicts the assumption that μ is not a loyalty test. So $T<\infty$. And by right-continuity of x and continuity of U^*, it must be that $Kx_T\leq\gamma U^*_T$. Hence $U^*_T\geq\frac{Kx_T}{\gamma}\geq\frac{K\phi}{\gamma}$.

Now, define a new function $\widetilde{U}$ by

$$\widetilde{U}_t\equiv\begin{cases}U^*_t, & t<T\\ \min\{U^*_T\exp((r+\gamma/K)(t-T)),(K-1)/(r+\gamma)\}, & t\geq T\end{cases}.$$

Further define a function $\widetilde{x}$ by

$$\widetilde{x}_t\equiv\frac{1}{K-1}\left((r+\gamma)\widetilde{U}_t-\dot{\widetilde{U}}_t\right),$$

with $\dot{\widetilde{U}}$ interpreted as a right-derivative when the derivative is discontinuous.

We first claim that $\widetilde{x}$ is a valid disclosure path. Right-continuity is by construction, so it remains to show monotonicity and $[\phi, 1]$-valuedness. We proceed by explicitly computing the function. For $t < T$ it is simply

$$\widetilde{x}_t \equiv \frac{1}{K-1}\left((r+\gamma)U_t^* - \dot{U}_t^*\right),$$

and by definition $\dot{U}_t^* = (r+\gamma)U_t^* - (K-1)x_t$ for $t < T$, meaning $\widetilde{x}_t = x_t$. So $\widetilde{x}$ is monotone and $[\phi, 1]$-valued on $[0, T)$. Meanwhile for $t \geq T$ the function takes the form

$$\widetilde{x}_t = \begin{cases} \frac{\gamma U_T^*}{K} \exp\left((r+\gamma/K)(t-T)\right), & T \leq t < T' \\ 1, & t \geq T' \end{cases},$$

where T' solves $U_T^* \exp((r+\gamma/K)(t-T)) = (K-1)/(\gamma+r)$. This expression is monotone everywhere, including at the jump time T', where $\widetilde{x}_{T'-} = \frac{K-1}{K}\frac{\gamma}{\gamma+r} < 1$. And $\widetilde{x}_T = \frac{\gamma U_T^*}{K}$, which is at least ϕ given the earlier observation that $U_T^* \geq K\phi/\gamma$. Thus given that it attains a maximum of 1, the function is $[\phi, 1]$-valued on $[T, \infty)$. The last remaining check is that $\widetilde{x}$ is monotone at the jump time T. Recall that $\widetilde{x}_{T-} = x_T$ while $\widetilde{x}_T = \frac{\gamma U_T^*}{K} \geq x_T$, so monotonicity holds at T as well, proving $\widetilde{x}$ is a valid disclosure path.

Next we claim that $\widetilde{U}$ is the continuation value process of the policy $b = 1$ under $\widetilde{x}$, and that $b = 1$ is an optimal undermining policy. The first part follows from the definition of $\widetilde{x}$, which can be re-written as the total differential

$$\frac{d}{dt}\left(\exp(-(r+\gamma)t)\widetilde{U}_t\right) = -\exp(-(r+\gamma)t)(K-1)\widetilde{x}_t.$$

Integrating both sides from t to ∞ and using the fact that $\lim_{t\to\infty} \exp(-(r+\gamma)t)\widetilde{U}_t = 0$ then yields

$$\widetilde{U}_t = \int_t^\infty \exp(-(r+\gamma)(s-t))(K-1)\widetilde{x}_s\, ds,$$

as desired.

As for the optimality of $b = 1$, this follows from Lemma A.3 if we can show that $d\widetilde{U}/dt = f(\widetilde{U}_t, t)$ a.e., when f is defined using $\widetilde{x}$ as the underlying disclosure rule. Given that $\widetilde{U}_t = U_t^*$ and $\widetilde{x}_t = x_t^*$ for $t < T$, and U_t^* is an optimal continuation utility process under x^*, Lemma A.3 ensures the ODE is satisfied for $t < T$. Meanwhile for $t > T$, the definition of $\widetilde{x}$ ensures that

$$(r+\gamma)\widetilde{U}_t - (K-1)\widetilde{x}_t = \frac{d\widetilde{U}}{dt}.$$

So we're done if we can show that $(r+\gamma)\widetilde{U}_t - (K-1)\widetilde{x}_t \leq r\widetilde{U}_t + \widetilde{x}_t$, i.e. $\widetilde{U}_t \leq \frac{K\widetilde{x}_t}{\gamma}$. Using the definition of $\widetilde{U}$ and the explicit form for $\widetilde{x}$ constructed earlier, we can read off that $\widetilde{U}_t = \frac{K\widetilde{x}_t}{\gamma}$ for $t < T$, while $\widetilde{U}_t = \frac{K-1}{\gamma+r} < \frac{K}{\gamma} = \frac{K\widetilde{x}_t}{\gamma}$ for $t \geq T$. So indeed $d\widetilde{U}/dt = f(\widetilde{U}_t, t)$ a.e., and $b = 1$ is an optimal undermining policy under $\widetilde{x}$.

The proposition will be proven if we can show that the principal obtains strictly higher expected profits under $\widetilde{x}$ than under x. We just saw that $\widetilde{U}_0$ is the disloyal agent's payoff under $\widetilde{x}$, and by construction $\widetilde{U}_0 = U_0^*$, where U_0^* is the disloyal agent's payoff under x. So it suffices to show that the loyal agent's payoff is higher under $\widetilde{x}$ than under x.

The loyal agent's payoff under x is just

$$V = \int_0^\infty e^{-rt} x_t \, dt.$$

Let $U_t^1 \equiv \int_t^\infty e^{-(r+\gamma)(s-t)} x_s \, ds$ be the disloyal agent's continuation utility process when he undermines at all times. Then $\frac{d}{dt}(e^{-(r+\gamma)t} U_t^1) = -e^{-(r+\gamma)t} x_t$, and using integration by parts V may be rewritten

$$V = U_0^1 + \gamma \int_0^\infty e^{\gamma t} U_t^1 \, dt.$$

Similarly, the loyal agent's payoff under $\widetilde{x}$ may be written

$$\widetilde{V} = \widetilde{U}_0 + \gamma \int_0^\infty e^{\gamma t} \widetilde{U}_t \, dt,$$

where recall $\widetilde{U}$ is the disloyal agent's continuation utility process from undermining at all times. So we're done if we can show that $\widetilde{U} \geq U^1$ and $\widetilde{U}_0 > U_0^1$.

We first establish that $\widetilde{U} \geq U^*$. In light of Lemma A.3, whenever $Kx_t \geq \gamma U_t^*$ it must be that

$$\frac{dU^*}{dt} = (r+\gamma)U_t^* - (K-1)x_t \leq (r + \gamma/K)U_t^*,$$

and whenever $Kx_t < \gamma U_t^*$, it must be that

$$\frac{dU^*}{dt} = rU_t^* + x_t \leq (r + \gamma/K)U_t^*.$$

Hence $dU^*/dt \leq (r+\gamma/K)U_t^*$ a.e. Then by Grönwall's inequality,

$$U_s^* \leq U_t^* \exp((r+\gamma/K)(s-t))$$

for any t and $s > t$.

Now, by construction $\widetilde{U}_t = U_t^*$ for $t \leq T$, and for $t \in [T, T']$ we have

$$\widetilde{U}_t = U_T^* \exp((r + \gamma/K)(t - T)) \geq U_t^*.$$

Finally, for $t \geq T$ we have $\widetilde{U}_t = (K-1)/(r+\gamma) \geq \widetilde{U}_t$ given that $(K-1)/(r+\gamma)$ is an upper bound on the value of any continuation utility. So indeed $\widetilde{U} \geq U^*$.

Finally, note that trivially $U^* \geq U^1$ given that U^* is the continuation utility process of an optimal policy, bounding above the continuation utility of an arbitrary policy at any point in time. And by assumption $U_0^* > U_0^1$, as x is not a loyalty test. So indeed $\widetilde{U} \geq U^* \geq U^1$ and $\widetilde{U}_0 \geq U_0^* > U^1$, completing the proof.

A.6 Proof of Proposition 2

We begin by providing closed form expressions for the key solution parameters and a precise statement of Proposition 2. Define

$$\begin{aligned}
t^* &\equiv \frac{1}{\gamma} \ln \left((K-1) \frac{1-q}{q} \right) \\
\overline{x} &\equiv \frac{(K-1)\gamma}{K(\gamma + r)} \\
\underline{x} &\equiv \frac{(K-1)\gamma}{K(\gamma + r)} \left(\frac{Kq}{K-1} \right)^{1 + \frac{Kr}{\gamma}} \\
\overline{t}_L &\equiv \frac{K}{\gamma} \ln \left(\frac{K-1}{Kq} \right) \\
\underline{t}_M &\equiv t^* - \frac{1}{\gamma} \ln \left(K e^{\frac{\gamma}{K}\Delta} - (K-1) \right) \\
\Delta &\equiv \left(r + \frac{\gamma}{K} \right)^{-1} \ln \left(\frac{\overline{x}}{\phi} \right) \\
\overline{t}_M &\equiv \underline{t}_M + \Delta,
\end{aligned}$$

and for completeness, define $\underline{t}_L \equiv 0$ and $\underline{t}_H \equiv \overline{t}_H \equiv t^*$.

We will prove that the unique optimal disclosure path x^* is as follows:

1. If $q \geq \frac{K-1}{K}$, $x_t^* = 1$ for all $t \geq 0$.

2. If $q < \frac{K-1}{K}$ and $\phi \geq \overline{x}$,

$$x_t^* = \begin{cases} \phi & \text{if } t \in [0, t^*) \\ 1 & \text{if } t \geq t^*. \end{cases}$$

3. If $q < \frac{K-1}{K}$ and $\phi \in (\underline{x}, \overline{x})$,

$$x_t^* = \begin{cases} \phi & \text{if } t \in [0, \underline{t}_M) \\ \phi e^{(r+\gamma/K)(t-\underline{t}_M)} & \text{if } t \in [\underline{t}_M, \overline{t}_M) \\ 1 & \text{if } t \geq \overline{t}_M. \end{cases}$$

4. If $q < \frac{K-1}{K}$ and $\phi \leq \underline{x}$,

$$x_t^* = \begin{cases} \underline{x} e^{(r+\gamma/K)t} & \text{if } t \in [0, \overline{t}_L) \\ 1 & \text{if } t \geq \overline{t}_L. \end{cases}$$

We solve the first two cases by solving a relaxed version of (PP), ignoring the IC constraint on $\dot{U}_t$, and checking afterward that it is satisfied. Note that the coefficient $q-(1-q)(K-1)e^{-\gamma t}$ in (1) is positive iff $q \geq (K-1)/K$ or both $q < \frac{K-1}{K}$ and $t > t^*$. Hence, the principal's relaxed problem is solved pointwise by setting $x_t = 1$ in those cases, and by setting $x_t = \phi$ otherwise. Now if $q \geq (K-1)/K$, we have $x_t^* = 1$ for all $t \geq 0$ and clearly the IC constraint is satisfied since $\dot{U}_t = 0$, solving the original problem and thus giving the first case of Proposition 2. If $q < (K-1)/K$, IC is clearly satisfied for the conjectured path at all times $t \geq t^*$. At times $t < t^*$, we have

$$\dot{U}_t = -(K-1)x_t + (r+\gamma)U_t,$$

and thus $\dot{U}_t \leq (r+\gamma/K)U_t$ if and only if $x_t \geq \frac{\gamma}{K}U_t$. Now $x_t = \phi$ and since U_t is increasing for $t < t^*$, this inequality is satisfied for all $t < t^*$ if and only if $\phi \geq \frac{\gamma}{K}U_{t^*-} = \frac{\gamma}{K}\frac{K-1}{r+\gamma} = \overline{x}$. Thus the solution to the relaxed problem solves (PP), giving the second case of Proposition 2.

For the remaining two cases of the proposition, where $q < \frac{K-1}{K}$ and $\phi < \overline{x}$, the IC constraint binds and we must solve the unrelaxed problem (PP''). We make use of notation and arguments from the proof of Lemma 6. The optimal disclosure path induces some expected utility U_0 for the disloyal agent, with U_0 lying in one of four possible regimes: $u = \underline{U}$, $u \in (\underline{U}, \widehat{U})$, $u \in [\widehat{U}, \overline{U})$ or $u = \overline{U}$, where $\widehat{U} = K\phi/\gamma$. We eliminate the first and fourth possibility as follows. If $U_0 = \underline{U}$, then $x_t^* = \phi$ for all time. But for sufficiently small $\varepsilon > 0$, an alternative contract with disclosure path $\tilde{x}_t$ defined by $\tilde{x}_t = \phi$ for $t \in [0, t^*)$ and $\tilde{x}_t = \phi + \varepsilon$ is a loyalty test in $\mathbb{X}$ and delivers a strictly higher payoff to the principal. Similarly, if $U_0 = \overline{U}$ then $x_t^* = 1$ for all time, but for sufficiently small $\varepsilon > 0$ an alternative contract with disclosure path $\tilde{x}_t$ defined by $\tilde{x}_t = 1 - \varepsilon$ for $t \in [0, t^*)$ and $\tilde{x}_t = 1$ for $t \geq t^*$ is a loyalty test in $\mathbb{X}$ and is a strict improvement over x^*. Hence, the optimal contract must satisfy either $U_0 \in (\underline{U}, \widehat{U})$ or $U_0 \in [\widehat{U}, \overline{U})$, with the associated optimal disclosure path x^* given by either equation (A.6) or equation (A.7), respectively.

Denote by $x^{*,M}$ and $x^{*,L}$ the paths defined in equations (A.6) and (A.7), respectively, with $u = U_0$. Let $\Pi^M[U_0]$ and $\Pi^L[U_0]$ be their respective payoffs to the principal. Also define the grand payoff function $\widehat{\Pi}[U_0]$ over $(\underline{U}, \overline{U})$ by

$$\widehat{\Pi}[U_0] = \begin{cases} \Pi^M[U_0], & U_0 \in (\underline{U}, \widehat{U}) \\ \Pi^L[U_0], & U_0 \in [\widehat{U}, \overline{U}). \end{cases}$$

The function $\widehat{\Pi}[U_0]$ captures the maximum profits achievable by the principal over all loyalty tests delivering an expected utility of exactly $U_0 \in (\underline{U}, \overline{U})$ to the disloyal agent. We will show that $\widehat{\Pi}$ is single-peaked with an interior maximizer, and we will derive a closed-form expression for the maximizer.

We begin with $x^{*,L}$, which has a single discontinuity at $\overline{t}_L(U_0) \equiv \frac{\ln\left(\frac{K-1}{(r+\gamma)U_0}\right)}{r+\frac{\gamma}{K}}$. The principal's profits under $x^{*,L}$ can therefore be written

$$\begin{aligned} \Pi^L[U_0] &= -\left(1 - \frac{K}{K-1}q\right)U_0 + \frac{q\gamma}{K-1}\int_0^\infty e^{-rt}U_t dt \\ &= -\left(1 - \frac{K}{K-1}q\right)U_0 + \frac{q\gamma}{K-1}\left[\frac{K}{\gamma}U_0\left(e^{\frac{\gamma}{K}\overline{t}_L(U_0)} - 1\right) + \frac{K-1}{r(r+\gamma)}e^{-r\overline{t}_L(U_0)}\right] \\ &= -U_0 + \frac{q\gamma}{K-1}\left(\frac{K}{\gamma} + \frac{1}{r}\right)\left(\frac{K-1}{r+\gamma}\right)^{\frac{\gamma/K}{r+\gamma/K}} U_0^{\frac{r}{r+\gamma/K}}. \end{aligned}$$

Taking a derivative, we obtain

$$\frac{d\Pi^L}{dU_0} = -1 + \frac{Kq}{K-1}\left(\frac{K-1}{r+\gamma}\right)^{\frac{\gamma/K}{r+\gamma/K}} U_0^{-\frac{\gamma/K}{r+\gamma/K}}.$$

Now $\Pi^L[\cdot]$ is strictly concave, and for $q < \frac{K-1}{K}$, it has a unique maximizer

$$U_0^{*,L} \equiv \frac{K-1}{\gamma+r}\left(\frac{Kq}{K-1}\right)^{1+\frac{Kr}{\gamma}}.$$

Given that $q < \frac{K-1}{K}$ and $\phi < \overline{x}$, it is straightforward to verify that $U_0^{*,L} \in [\widehat{U}, \overline{U})$ iff $\phi \leq \underline{x}$, and that in this case $\overline{t}_L(U_0^{*,L}) > t^*$. On the other hand, if $\phi > \underline{x}$, then $U_0^{*,L} < \widehat{U}$ and thus $\Pi^L[U_0]$ is decreasing in U_0 for $U_0 \in [\widehat{U}, \overline{U})$.

Next, we turn to $x^{*,M}$. To simplify the algebra that follows, define the function

$$\underline{t}(U_0) \equiv \frac{1}{r+\gamma} \log\left(\frac{\widehat{U} - \underline{U}}{U_0 - \underline{U}}\right)$$

to be the time at which $x^{*,M}$ begins increasing. This is a strictly decreasing function of U_0, and so may be inverted to write $\widehat{\Pi}^L$ as a function of $\underline{t}$ over $(0, \infty)$.

With this change of variables, the single discontinuity of $x^{*,M}$ occurs at time $\overline{t}(\underline{t}) = \underline{t} + \Delta$, where Δ is defined in the beginning of this proof. The principal's payoff function may then be written

$$\begin{aligned}
\Pi^M[\underline{t}] &\equiv \int_0^\infty e^{-rt}(q - (K-1)(1-q)e^{-\gamma t})x_t dt \\
&= \int_0^{\underline{t}} e^{-rt}(q - (K-1)(1-q)e^{-\gamma t})\phi dt + \int_{\underline{t}}^{\overline{t}(\underline{t})} e^{-rt}(q - (K-1)(1-q)e^{-\gamma t})\phi e^{(r+\gamma/K)(t-\underline{t})} dt \\
&\quad + \int_{\overline{t}(\underline{t})}^\infty e^{-rt}(q - (K-1)(1-q)e^{-\gamma t})dt.
\end{aligned}$$

Differentiating this expression wrt $\underline{t}$, and recalling that $\overline{t}'(\underline{t}) = 1$, we obtain

$$\begin{aligned}
\frac{d\Pi^M}{d\underline{t}} &= -\int_{\underline{t}}^{\overline{t}(\underline{t})} e^{-rt}(q - (K-1)(1-q)e^{-\gamma t})(r + \gamma/K)\phi e^{(r+\gamma/K)(t-\underline{t})} dt \\
&\quad - (1-\overline{x})e^{-r\overline{t}(\underline{t})}(q - (K-1)(1-q)e^{-\gamma \overline{t}(\underline{t})}) \\
&= -\left(r + \frac{\gamma}{K}\right)\phi e^{-r\underline{t}} \frac{K}{\gamma}\left[q\left(e^{\frac{\gamma}{K}\Delta} - 1\right) + (1-q)e^{-\gamma \underline{t}}\left(e^{-\gamma\left(\frac{K-1}{K}\right)\Delta} - 1\right)\right] \\
&\quad - (1-\overline{x})e^{-r\underline{t}}\left(qe^{-r\Delta} - (K-1)(1-q)e^{-r\Delta}e^{-\gamma(\underline{t}+\Delta)}\right) \\
&= \left(r + \frac{\gamma}{K}\right)e^{-r\underline{t}}\left[Ae^{-\gamma \underline{t}} + B\right],
\end{aligned}$$

where

$$\begin{aligned}
A &= -\phi \frac{K(1-q)}{\gamma}\left(e^{-\frac{\gamma(K-1)}{K}\Delta} - 1\right) + \frac{1-q}{r+\gamma}(K-1)e^{-(r+\gamma)\Delta} \\
&= \frac{1-q}{r+\gamma}(K-1)e^{-\left(r+\frac{\gamma}{K}\right)\Delta} > 0 \\
B &= -\phi \frac{Kq}{\gamma}\left(e^{\frac{\gamma}{K}\Delta} - 1\right) - \frac{1}{r+\gamma}qe^{-r\Delta} \\
&= -\frac{Kq}{\gamma(r+\gamma)}e^{-r\Delta}\left(\gamma - \phi(r+\gamma)e^{r\Delta}\right) < 0,
\end{aligned}$$

with the last inequality relying on the fact that $\phi < \overline{x}$.

This expression vanishes at the unique time $\underline{t}$ such that $Ae^{-\gamma \underline{t}} + B = 0$, in particular at time $\underline{t}_M$ as defined at the start of the proof. To the left of this time Π^M is stricty increasing, while to the right Π^M is strictly decreasing. It is straightforward to verify that $\underline{t}_M > 0$ iff $\phi > \underline{x}$, in which case $\Pi^M[U_0]$ is single-peaked with an interior maximizer

$$U_0^{*,M} \equiv \underline{t}^{-1}(\underline{t}_M) \in (\underline{U}, \widehat{U}),$$

and $0 < \underline{t}_M < t^* < \bar{t}_M$. Otherwise if $\phi \leq \underline{x}$, then Π^M is monotonically decreasing in $\underline{t}$, or equivalently, monotonically increasing in U_0 for $U_0 \in (\underline{U}, \widehat{U})$.

Since $\Pi^M[\widehat{U}] = \Pi^L[\widehat{U}]$, the above arguments imply that $\widehat{\Pi}[U_0]$ is single-peaked, and its maximizer is (i) $U^{*,L} \in [\widehat{U}, \overline{U})$ if $\phi \leq \underline{x}$ or (ii) $U^{*,M} \in (\underline{U}, \overline{U})$ if $\phi \in (\underline{x}, \overline{x})$. In case (i), $x^{*,L}$ with $u = U_0^{*,L}$ is the optimal disclosure path, and in case (ii), $x^{*,M}$ with $u = U_0^{*,M}$ is the optimum.

To summarize, in all parametric cases, the optimal disclosure path satisfies the optimal growth property and is uniquely detemined in closed form, characterized by the parameters stated in the beginning of this section.

A.7 Proof of Proposition 3

By Lemma 2, there exists an underlying information policy μ for the optimal disclosure path x^* in each of the parametric cases. We now consider the three items of the proposition. Suppose that $x_t > x_{t-}$. Recall that μ is a martingale and has left limits, so $\mu_{t-} = \mathbb{E}_{t-}[\mu_t] \equiv \lim_{t' \uparrow t} \mathbb{E}_{t'}[\mu_t]$. Since μ_t takes values $m_t^{\pm}$ defined in the proof of Lemma 2, it suffices to consider binary signals $S_t \in \{L, R\}$ and without loss label these signals so that, given any μ_{t-}, $\mu_t = m_t^+$ if and only if $S_t = R$. For $\omega' \in \{L, R\}$, let $p_t(\omega')$ be the probability that $S_t = R$ conditional on information up to time t and conditional on $\omega = \omega'$. By Bayes' rule, we have two equations in two unknowns:

$$\begin{aligned} m_t^+ &= \frac{\mu_{t-} p_t(R)}{\mu_{t-} p_t(R) + (1 - \mu_{t-}) p_t(L)} \\ m_t^- &= \frac{\mu_{t-}(1 - p_t(R))}{\mu_{t-}(1 - p_t(R)) + (1 - \mu_{t-})(1 - p_t(L))}. \end{aligned}$$

The solution is

$$\begin{aligned} p_t(R) &= \frac{m_t^+(\mu_{t-} - m_t^-)}{\mu_{t-}(m_t^+ - m_t^-)} = \frac{(1 + x_t)(2\mu_{t-} - 1 + x_t)}{4\mu_{t-} x_t} \\ p_t(L) &= \frac{(1 - m_t^+)(\mu_{t-} - m_t^-)}{(1 - \mu_{t-})(m_t^+ - m_t^-)} = \frac{(1 - x_t)(2\mu_{t-} - 1 + x_t)}{4(1 - \mu_{t-}) x_t}. \end{aligned}$$

The agent's state conjecture is correct if $\mu_{t-} > \frac{1}{2}$ and $\omega = R$ or if $\mu_{t-} < \frac{1}{2}$ and $\omega = L$. By symmetry, suppose $\mu_{t-} < \frac{1}{2}$. If the agent's state conjecture is correct, then $\omega = L$ and he gets a contradictory signal with probability $p_t(L) = \frac{(1-x_t)(x_t - x_{t-})}{2x_t(1+x_{t-})}$. If instead the agent's state conjecture is incorrect, then $\omega = R$ and he gets a contradictory signal with probability $p_t(R) = \frac{(1+x_t)(x_t - x_{t-})}{2x_t(1-x_{t-})}$.

Next, suppose that $\dot{x}_t > 0$. Since beliefs move deterministically conditional on no jump occuring, we consider Poisson signals. Since beliefs move across $\frac{1}{2}$ when a jump occurs, we consider contradictory Poisson learning, that is, Poisson processes with higher arrival rate when the agent's current action being wrong. Let $\overline{\lambda}_t, \underline{\lambda}_t$ be the arrival rates conditional on the current action being right or wrong, respectively. Consider $\mu_t > 1/2$. Then by Bayes' rule, absent an arrival,

$$\dot{x}_t = 2\dot{\mu}_t = 2\mu_t(1-\mu_t)(\overline{\lambda}_t - \underline{\lambda}_t) \tag{A.3}$$

$$= (1+x_t)\frac{1-x_t}{2}(\overline{\lambda}_t - \underline{\lambda}_t). \tag{A.4}$$

In addition, conditional on an arrival, the belief updates from $\mu_t = \frac{1+x_t}{2}$ to $m_t^- = \frac{1-x_t}{2}$, so by Bayes' rule,

$$\frac{1-x_t}{2} = \frac{\frac{1+x_t}{2}\underline{\lambda}_t}{\frac{1+x_t}{2}\underline{\lambda}_t + \frac{1-x_t}{2}\overline{\lambda}_t}. \tag{A.5}$$

The unique solution to the linear system in (A.4) and (A.5) is $\overline{\lambda}_t$, $\underline{\lambda}_t$ as stated in the proposition. A symmetric argument shows that the same arrival rates apply when $\mu_t < 1/2$. To obtain the expected arrival rate, again assume wlog that $\mu_t > 1/2$ and recall that $\dot{x}_t = x_t(r + \gamma/K)$; by direct computation,

$$\begin{aligned}\mu_t\underline{\lambda}_t + (1-\mu_t)\overline{\lambda}_t &= \frac{1+x_t}{2}\frac{\dot{x}_t}{2x_t}\frac{1-x_t}{1+x_t} + \frac{1-x_t}{2}\frac{\dot{x}_t}{2x_t}\frac{1+x_t}{1-x_t} \\ &= \frac{1}{2}(r + \gamma/K).\end{aligned}$$

The case where $\dot{x}_t = 0$ and $x_{t-} = x_t$ is trivial.

A.8 Proof of Lemma 4

Fix a disclosure path x corresponding to some information policy. Suppose there existed a time t at which $\dot{U}_t$ is defined and (2) failed. Consider any undermining policy which unconditionally sets $b_t = 0$ from time t to time $t + \varepsilon$ for $\varepsilon > 0$, and then unconditionally sets $b_t = 1$ after time $t + \varepsilon$. Let $V(\varepsilon)$ be the (time-zero) expected continuation payoff to the disloyal agent of such a strategy beginning at time t, conditional on not having been terminated prior to time t. Then

$$V(\varepsilon) = -\int_t^{t+\varepsilon} e^{-r(s-t)} x_s \, ds + e^{-r\varepsilon} U_{t+\varepsilon}.$$

Differentiating wrt ε and evaluating at $\varepsilon = 0$ yields

$$V'(0) = -x_t - rU_t + \dot{U}_t.$$

If $V'(0) > 0$, then for sufficiently small ε the expected payoff of unconditionally undermining forever is strictly improved on by deviating to not undermining on the interval $[t, t+\varepsilon]$ for sufficiently small $\varepsilon > 0$. (Note that undermining up to time t yields a strictly positive probability of remaining employed by time t and achieving the continuation payoff $V(\varepsilon)$.) So

$$-x_t - rU_t + \dot{U}_t \leq 0$$

is a necessary condition for unconditional undermining to be an optimal strategy for the disloyal agent.

Now, by the fundamental theorem of calculus $\dot{U}_t$ is defined whenever x_t is continuous, and is equal to

$$\dot{U}_t = (r+\gamma)U_t - (K-1)x_t.$$

As x is a monotone function it can have at most countably many discontinuities. So this identity holds a.e. Solving for x_t and inserting this identity into the inequality just derived yields (2), as desired.

In the other direction, suppose μ is a deterministic information policy with disclosure path x satisfying (2) a.e. To prove the converse result we invoke Lemma A.3, which requires establishing that

$$\dot{U}_t = \min\{(r+\gamma)U_t - (K-1)x_t, rU_t + x_t\}$$

a.e. By the fundamental theorem of calculus, $\dot{U}_t$ exists a.e. and is equal to

$$\dot{U}_t = (r+\gamma)U_t - (K-1)x_t$$

wherever it exists. So it is sufficient to show that $(r+\gamma)U_t - (K-1)x_t \leq rU_t + x_t$ a.e., i.e. $\gamma U_t \leq K x_t$.

Combining the hypothesized inequality $\dot{U}_t \leq \left(r + \frac{\gamma}{K}\right) U_t$ and the identity $\dot{U}_t = (r+\gamma)U_t -$

$(K-1)x_t$ yields the inequality

$$(r+\gamma)U_t - (K-1)x_t \le \left(r + \frac{\gamma}{K}\right) U_t.$$

Re-arrangement shows that this inequality is equivalent to $x_t \ge \gamma U_t / K$, as desired.

A.9 Proof of Lemma 5

Let

$$\Pi = \int_0^\infty e^{-rt} x_t (q - (1-q)(K-1)e^{-\gamma t})\, dt.$$

Use the fact that $\dot{U}_t = -(K-1)x_t + (r+\gamma)U_t$ as an identity to eliminate x_t from the objective, yielding

$$\Pi = \int_0^\infty \left(\frac{q}{K-1}e^{\gamma t} - (1-q)\right) e^{-(r+\gamma)t}((r+\gamma)U_t - \dot{U}_t)\, dt.$$

This expression can be further rewritten as a function of U only by integrating by parts. The result is

$$\Pi = -\left(1 - \frac{K}{K-1}q\right) U_0 - \lim_{t\to\infty} \left(\frac{q}{K-1}e^{\gamma t} - (1-q)\right) e^{-(r+\gamma)t} U_t + \frac{q\gamma}{K-1} \int_0^\infty e^{-rt} U_t\, dt.$$

Since U is bounded in the interval $[0, (K-1)/(r+\gamma)]$, the surface term at $t = \infty$ must vanish, yielding the expression in the lemma statement.

A.10 Proof of Lemma 6

Fix $u \in [\underline{U}, \overline{U}]$. Ignoring the fixed term involving U_0 and all multiplicative constants, the principal's objective reduces to

$$\int_0^\infty e^{-rt} U_t\, dt,$$

which is increasing in the path of U_t pointwise. Let $g(U) \equiv \min\{(r+\gamma/K)U, (r+\gamma)U - (K-1)\phi\}$. We begin by constructing the unique pointwise maximizer of U subject to $U_0 = u$, $U \in \mathbb{U}$, and $\dot{U}_t \le g(U_t)$ a.e.

Define a function $U^\dagger$ as follows. If $u = \underline{U}$, then set $U_t^\dagger = u$ for all time. If $u \in (\underline{U}, \widehat{U})$ where $\widehat{U} \equiv K\phi/\gamma$, then set

$$U_t^\dagger = \begin{cases} e^{(r+\gamma)t}\left(u - \underline{U}\left(1 - e^{-(r+\gamma)t}\right)\right), & t \le \underline{t} \\ \widehat{U} \exp\left(\left(r + \frac{\gamma}{K}\right)(t - \underline{t})\right), & t > \underline{t} \end{cases},$$

where $\underline{t} \equiv \min\left\{0, \frac{1}{r+\gamma}\log\left(\frac{\widehat{U}-\underline{U}}{u-\underline{U}}\right)\right\}$. Finally, if $u \in [\widehat{U}, \overline{U}]$, then set $U_t^\dagger = u\exp\left(\left(r+\frac{\gamma}{K}\right)t\right)$ for all time. It is readily checked by direct computation that $U^\dagger$ satisfies $U_0^\dagger = u$ and $\dot{U}_0^\dagger = g(U_t^\dagger)$ a.e.

Now, note that g is strictly increasing in U. By a slight variant of the arguments used to prove Lemma A.1, one can show that for any absolutely continuous function U satisfying $\dot{U}_t \leq g(U_t)$ a.e. and $U_0 = u$, it must be that $U^\dagger \geq U$. Further, if $\dot{U}_t = g(U_t)$ a.e., then $U = U^\dagger$, and if on the other hand $\dot{U}_t < g(U_t)$ on a positive-measure set of times on any interval $[0, T]$, then $U_T^\dagger > U_T$.

Define a function U^* by $U_t^* = \min\{U_t^\dagger, \overline{U}\}$. If $u = \underline{U}$, then the previous facts imply that the only element of $\mathbb{U}$ satisfying the control constraint and $U_0 = u$ is the constant function $U^* = \underline{U}$, which is then trivially the pointwise maximizer of the objective. Otherwise, note that $U^\dagger$ is a strictly increasing function, and there exists a unique time $\bar{t}$ at which $U_{\bar{t}}^\dagger = \overline{U}$. Then by the facts of the previous paragraph, U^* is the unique pointwise maximizer of U on the interval $[0, \bar{t}]$ subject to the control constraint and $U_0 = u$; and it is trivially the unique pointwise maximizer of U subject to $U \leq \overline{U}$ on the interval $(\bar{t}, \infty)$. Thus U^* is the unique pointwise maximizer of U in $\mathbb{U}$ subject to the control constraint and $U_0 = u$.

The final step in the proof is to show that the disclosure path implied by U^* lies in $\mathbb{X}$. Let

$$x_t^* \equiv \frac{1}{K-1}\left((r+\gamma)U_t^* - \dot{U}_t^*\right).$$

If $u = \underline{U}$, then $x^* = \phi$, which is in $\mathbb{X}$. Next, if $u \in (\underline{U}, \widehat{U})$ then

$$x_t^* = \begin{cases} \phi, & t \leq \underline{t} \\ \phi\exp\left(\left(r+\frac{\gamma}{K}\right)(t-\underline{t})\right), & \underline{t} < t < \bar{t} \\ 1, & t \geq \bar{t}. \end{cases} \tag{A.6}$$

Note that $x_{\bar{t}}^* = 1$ while $x_{\bar{t}-}^* = \frac{\gamma}{K}\frac{K-1}{r+\gamma} < 1$. So x^* is monotone increasing, bounded below by $x_0^* = \phi$ and above by 1. So $x \in \mathbb{X}$. If $u \in [\widehat{U}, \overline{U})$, then

$$x_t^* = \begin{cases} \frac{u\gamma}{K}\exp\left(\left(r+\frac{\gamma}{K}\right)t\right), & t < \bar{t} \\ 1, & t \geq \bar{t}. \end{cases} \tag{A.7}$$

Given $u \geq \widehat{U}$, $u\gamma/K \geq \phi$, so x^* is monotone increasing at $t = 0$. And at $t = \bar{t}$ it satisfies $x_{\bar{t}}^* = 1$ and $x_{\bar{t}-}^* = \frac{\gamma}{K}\frac{K-1}{r+\gamma} < 1$. Thus x^* is monotone increasing everywhere, bounded below by $x_0^* = u\gamma/K \geq \phi$ and above by 1. So $x^* \in \mathbb{X}$.

Finally, if $u = \overline{U}$, then $x^* = 1$ which is clearly in $\mathbb{X}$.

A.11 Proof of Proposition 4

More precisely, we prove the following:

- If $q \geq \frac{K-1}{K}$ or $\phi \geq \overline{x}$, $\beta_t = 1$ for all $t \geq 0$.
- If $q < \frac{K-1}{K}$ and $\phi \leq \underline{x}(q)$, $\beta_t = \frac{1}{K(1-qe^{\gamma/Kt})}$ for $t \in [0, \overline{t}_L)$ and $\beta_t = 1$ for $t \geq \overline{t}_L$.
- If $q < \frac{K-1}{K}$ and $\phi \in (\underline{x}(q), \overline{x})$,

$$\beta_t = \begin{cases} 1 & \text{if } t \in [0, \underline{t}_M) \\ \frac{1}{K(1-Ce^{(\gamma/K)(t-\underline{t}_M)})} & \text{if } t \in [\underline{t}_M, \overline{t}_M) \\ 1 & \text{if } t \geq \overline{t}_M, \end{cases}$$

  where $C \equiv \left(\frac{\phi}{\overline{x}}\right)^{\frac{\gamma}{\gamma+Kr}} \frac{K-1}{K}$.

We also show that β_t is increasing whenever $\beta_t < 1$, and that β_t is continuous at time $\overline{t}_L$ (in the low knowledge case) or $\overline{t}_M$ (in the moderate knowledge case).

The discussion preceding the proposition covers the first item. For the second item, we begin by deriving a unique candidate for β_t and then show that it supports time-consistency. Note that a necessary condition for time-consistency is

$$\underline{x}(\pi_t) = x_t^* = \underline{x}(q)e^{(r+\gamma/K)t}$$
$$\iff \frac{(K-1)\gamma}{K(\gamma+r)}\left(\frac{K\pi_t}{K-1}\right)^{1+\frac{Kr}{\gamma}} = \frac{(K-1)\gamma}{K(\gamma+r)}\left(\frac{Kq}{K-1}\right)^{1+\frac{Kr}{\gamma}} e^{(r+\gamma/K)t}$$
$$\iff \pi_t = qe^{\gamma/Kt}.$$

Differentiating the above yields $\dot{q}_t = \frac{\gamma}{K}qe^{\gamma/Kt} = \frac{\gamma}{K}\pi_t$. To derive β_t, differentiate the Bayes' rule equation:

$$\pi_t = \frac{q}{q + (1-q)\exp\left(-\gamma \int_0^t \beta_s ds\right)}$$
$$\implies \dot{q}_t = \gamma\beta_t\pi_t(1-\pi_t).$$

Equating the two expressions for $\dot{\pi}_t$, we obtain the unique candidate

$$\beta_t = \frac{1}{K(1-\pi_t)} = \frac{1}{K(1-qe^{\gamma/Kt})}.$$

Now β_t is strictly increasing, nonnegative, and satisfies $\lim_{t\uparrow\bar{t}_L}\beta_t = \frac{1}{K(1-qe^{\gamma/K\bar{t}_L})} = 1$. We now check that the rest of the contract is time-consistent. Denote by $(x_s^{t_0})_{s\geq 0}$ the disclosure path of an optimal contract starting (with time reset to 0) with the time-t_0 agent reputation π_{t_0} and disclosure level $x^*_{t_0}$ from the path of the original contract; we aim to show that $x_s^{t_0} = x_{t_0+s}$ for all $s \geq 0$, $t_0 > 0$. First consider $t_0 \in (0, \bar{t}_L(q))$. Observe first that in the original contract, information grows at the constant rate $r + \gamma/K$ until time $\bar{t}_L(q)$; hence for all $t_0 \in (0, \bar{t}_L(q))$, $x^*_{t_0+s} = x_s^{t_0}$ for all $s \in [0, \bar{t}_L(q) - t_0)$. Moreover, $\bar{t}_L(q) = t_0 + \bar{t}_L(\pi_{t_0})$, so the final jump in the new contract also occurs at the originally scheduled time $\bar{t}_L(q)$, and for all $s \geq \bar{t}_L(q) - t_0$, we have $x_s^{t_0} = 1 = x^*_{t_0+s}$. This establishes time-consistency for $t_0 \in (0, \bar{t}_L(q))$. For $t_0 \geq \bar{t}_L(q)$, note that $\pi_{t_0} \geq \frac{K-1}{K}$, so $x_s^{t_0} = 1 = x^*_{t_0+s}$ for all $s \geq 0$.

For the moderate knowledge case, we use a similar process. The agent strictly prefers to undermine at all times before $\underline{t}_M(q)$ and after time $\bar{t}_M(q)$. Hence undermining with certainty at those times must be part of the disloyal agent's strategy.

For times $t \in (\underline{t}_M(q), \bar{t}_M(q))$, the agent is indifferent and undermines to satisfy the same equality above:

$$
\begin{aligned}
& \underline{x}(\pi_t) = x^*_t \\
\iff & \frac{(K-1)\gamma}{K(\gamma+r)}\left(\frac{K\pi_t}{K-1}\right)^{1+\frac{Kr}{\gamma}} = \phi e^{(r+\gamma/K)(t-\underline{t}_M)} \\
\iff & \pi_t = Ce^{(\gamma/K)(t-\underline{t}_M)},
\end{aligned}
$$

where $C \equiv \left(\frac{\phi K(\gamma+r)}{(K-1)\gamma}\right)^{\frac{\gamma}{\gamma+Kr}}\frac{K-1}{K} = \left(\frac{\phi}{\bar{x}}\right)^{\frac{\gamma}{\gamma+Kr}}\frac{K-1}{K}$. Again, taking derivatives of this equation and the Bayes' rule equation and equating the two, we get

$$
\begin{aligned}
& \dot{\pi}_t = \gamma\beta_t\pi_t(1-\pi_t) = C\frac{\gamma}{K}e^{(\gamma/K)(t-\underline{t}_M)} = \frac{\gamma}{K}\pi_t \\
\iff & \beta_t = \frac{1}{K(1-\pi_t)} = \frac{1}{K(1-Ce^{(\gamma/K)(t-\underline{t}_M)})}.
\end{aligned}
$$

By inspection, $\beta_t \in [0,1]$ is increasing; it can also be verified that $\pi_{\bar{t}_M(q)} = \frac{K-1}{K}$ and hence $\lim_{t\uparrow\bar{t}_M}\beta_t = 1$. To conclude, we check time-consistency. Time-consistency for $t_0 \geq \bar{t}_M(q)$ is immediate from the fact that $\pi_{t_0} \geq \frac{K-1}{K}$, so $x_s^{t_0} = 1 = x^*_{t_0+s}$ for all $s \geq 0$. Next consider $t_0 \in [0, \underline{t}_M(q))$ and note that $\underline{x}(\pi_{\underline{t}_M(q)}) = \phi$. Since $\underline{x}$ is strictly increasing in q and q is strictly increasing in t, we have $\underline{x}(\pi_{t_0}) < \phi < \bar{x}$ for all $t_0 \in [0, \underline{t}_M(q))$. It follows that for such t_0, the state $(x^*_{t_0}, \pi_{t_0}) = (\phi, \pi_{t_0})$ is in the moderate knowledge case. Using that $\pi_t = \frac{q}{q+(1-q)\exp(-\gamma t)}$ for $t \in [0, \underline{t}_M(q))$, it is straightforward to verify that for all $t_0 \in [0, \underline{t}_M(q))$, $\underline{t}_M(\pi_{t_0}) = \underline{t}_M(q) - t_0$ and $\bar{t}_M(\pi_{t_0}) = \bar{t}_M(q) - t_0$, which shows that $x_s^{t_0} = x^*_{t_0+s}$ for all $s \geq 0$, establishing time-consistency for such t_0. For $t_0 \in [\underline{t}_M(q), \overline{t}_M(q))$, we have $\pi_{t_0} < \frac{K-1}{K}$ and (by construction) $x^*_{t_0} = \underline{x}(\pi_{t_0})$, hence the low knowledge case applies for the new contract. Information then grows at rate $r + \gamma/K$, so $x^{t_0}_s = x^*_{t_0} + s$ for all $s \in [0, \overline{t}_M(q) - t_0)$. It is straightforward to verify that $\overline{t}_L(\pi_{t_0}) = \overline{t}_M(q) - t_0$, hence $x^{t_0}_s = x^*_{t_0+s} = 1$ for all $s \geq \overline{t}_M(q) - t_0$, establishing time-consistency for $t_0 \in [\underline{t}_M(q), \overline{t}_M(q))$. This concludes the proof.

A.12 Comparative statics

In this section we provide comparative statics results for the low knowledge and high knowledge cases and prove a proposition which covers all three cases.

Table 2: Comparative statics (high knowledge)

	$dt^*/$	$d\Pi/$
$d\gamma$	$-$	$+$
dq	$-$	$+$
dK	$+$	$-$
dr	0	$-$
$d\phi$	0	$-$

Table 3: Comparative statics (low knowledge)

	$d\overline{x}/$	$d\underline{x}/$	$d\overline{t}_L/$	$d\Pi/$
$d\gamma$	$+$	$+$	$-$	$+$
dq	0	$+$	$-$	$+$
dK	$+$	$-$	$+$	$-$
dr	$-$	$-$	0	$-$
$d\phi$	0	0	0	$-$

Proposition A.1. *In the low, moderate and high knowledge cases, the relationships between model inputs and outputs are as given in Tables 1, 2 and 3.*

Proof of Proposition A.1. We proceed one output at a time. The analysis of $\overline{x}$ applies to both the low and moderate knowledge cases.[24] The analysis of the principal's profit — with the exception of comparative statics with respect to r, which we treat separately in Lemma A.4 — applies to all three cases, and is provided last.

[24]This value has a single definition, and the arguments used in signing its derivatives do not use any additional facts about parameter values specific to the low or moderate knowledge cases.

- t^* (Table 2) and $\overline{x}$ (Tables 1 and 3): The signs of these derivatives are immediate from inspection of the respective formulas given in Section A.6.

- $\underline{x}$ (Tables 1 and 3): Recall that $\underline{x} = \frac{(K-1)\gamma}{K(\gamma+r)} \left(\frac{Kq}{K-1}\right)^{1+\frac{Kr}{\gamma}}$, which is clearly increasing in q and independent of ϕ. The first factor is decreasing in r and increasing in γ, as is the second factor since $q < \frac{K-1}{K}$, which implies $d\underline{x}/dr < 0 < d\underline{x}/d\gamma$. By direct calculation, $d\underline{x}/dK$ has the same sign as $-1 + (K-1)\ln\left(\frac{Kq}{K-1}\right) < -1$.

- Δ (Table 1): From $\Delta = \left(r + \frac{\gamma}{K}\right)^{-1} \ln\left(\frac{\overline{x}}{\phi}\right)$, it is immediate that $d\Delta/d\phi < 0 = d\Delta/dq$. Now both factors are increasing in K, so unambiguously $d\Delta/dK > 0$. By direct calculation, $d\Delta/dr = -\frac{K}{(r+\gamma)(Kr+\gamma)^2}\left[Kr + \gamma + K(r+\gamma)\ln\left(\frac{\overline{x}}{\phi}\right)\right]$. For $\phi < \overline{x}$, all terms in square brackets are positive, so $d\Delta/dr < 0$. We have $d\Delta/d\gamma = \frac{K}{(Kr+\gamma)^2}\left[\frac{r(Kr+\gamma)}{\gamma(r+\gamma)} - \ln\left(\frac{\overline{x}}{\phi}\right)\right]$; for ϕ sufficiently close to $\overline{x}$, this is strictly positive. However, evaluating at $\phi = \underline{x}$ yields $\frac{K\left[r+(r+\gamma)\ln\left(\frac{Kq}{K-1}\right)\right]}{\gamma(r+\gamma)(Kr+\gamma)}$, which is strictly negative for r sufficiently small. We conclude that $d\Delta/d\gamma$ can be positive or negative.

- $\underline{t}_M$ (Table 1): First, we have $d\underline{t}_M/dq = -\frac{1}{q(1-q)\gamma} < 0$. Next, define locally $\alpha \equiv \left(\frac{\overline{x}}{\phi}\right)^{\frac{\gamma}{Kr+\gamma}}$, where $\alpha > 1$ when the moderate knowledge case applies (since $\phi > \overline{x}$). By direct calculation,

$$
\begin{aligned}
d\underline{t}_M/dK &= \frac{Kr\alpha[1+(K-1)\ln\alpha]}{(K-1)\gamma(Kr+\gamma)[1-K+K\alpha]} > 0 \\
d\underline{t}_M/dr &= \frac{K\alpha[1+K(r/\gamma+1)\ln\alpha]}{(r+\gamma)(Kr+\gamma)[1-K+K\alpha]} > 0 \\
d\underline{t}_M/d\phi &= \frac{K\alpha}{\phi(Kr+\gamma)[1-K+K\alpha]} > 0 \\
d\underline{t}_M/d\gamma &= \gamma^{-2}[\ln(X) - Y] < 0, \qquad \text{where} \\
X &\equiv \frac{q(1-K+K\alpha)}{(1-q)(K-1)} \\
Y &\equiv \frac{Kr\gamma\alpha[Kr+\gamma+K(r/\gamma+1)\ln\alpha]}{(r+\gamma)(Kr+\gamma)[1-K+K\alpha]}.
\end{aligned}
$$

Now $\ln(X) \leq \ln\left[\frac{q}{(1-q)(K-1)}\left(1 - K + K\left[\frac{\overline{x}}{\underline{x}}\right]^{\frac{\gamma}{Kr+\gamma}}\right)\right] = 0$, and by inspection $Y > 0$ so $d\underline{t}_M/d\gamma < 0$, as desired.

- $\overline{t}_M$ (Table 1): Since $d\Delta/dq = 0$, $d\overline{t}_M/dq = d\underline{t}_M/dq < 0$. Also, $d\overline{t}_M/dK = d\Delta/dK +$

$d\underline{t}_M dK > 0$ unambiguously. By direct calculation, using the definition of α from above,

$$d\bar{t}_M/dr = -\frac{(K-1)K(\alpha-1)[1+K(1+r/\gamma)\ln\alpha]}{(r+\gamma)(Kr+\gamma)[1-K+K\alpha]} < 0$$
$$d\bar{t}_M/d\phi = -\frac{(K-1)K(\alpha-1)}{\phi(Kr+\gamma)[1-K+K\alpha]} < 0.$$

Finally, by tedious calculation, we have $\frac{d^2\bar{t}_M}{dqd\gamma} = \frac{1}{\gamma^2 q(1-q)} > 0$. For an upper bound on $d\bar{t}_M/d\gamma$, we increase q until $\underline{x} = \phi$. We obtain

$$\begin{aligned} d\bar{t}_M/d\gamma|_{\phi=\underline{x}} &= \frac{(K-1)\left[-\left(\frac{Kq}{K-1}-1\right)r\gamma + \frac{K}{K-1}(r+\gamma)[r+\gamma(1-q)]\ln\left(\frac{Kq}{K-1}\right)\right]}{(1-q)\gamma^2(r+\gamma)(Kr+\gamma)} \\ &< \frac{(K-1)\left(\frac{Kq}{K-1}-1\right)}{(1-q)\gamma^2(r+\gamma)(Kr+\gamma)}\left[-r\gamma + \frac{K}{K-1}(r+\gamma)[r+\gamma(1-q)]\right] \\ &< 0, \end{aligned}$$

where we have used $q < \frac{K-1}{K}$ and the inequality $\ln x < x - 1$ for $x > 0$.

- $\bar{t}_L$ (Table 3): By inspection, $\bar{t}_L$ is independent of r and ϕ and is decreasing in γ and q. For K, Calculate $d\bar{t}_L/dK = \frac{1}{\gamma}\left[\frac{1}{(K-1)} - \ln\left(\frac{Kq}{K-1}\right)\right] > \frac{1}{(K-1)\gamma} > 0$.

- Π (Tables 1, 2 and 3): Recall that $\mathbb{X}$ is the set of monotone, càdlàg $[\phi, 1]$-valued functions x and let $\mathbb{B}$ denote the set of pure strategies b for the disloyal agent. The principal's payoff is the maximum of 0 and

$$\max_{x\in\mathbb{X}}\min_{b\in\mathbb{B}} \int_0^\infty x_t e^{-rt}\left[q - (1-q)(b_t K - 1)\exp\left(-\gamma\int_0^t b_s ds\right)\right] dt. \qquad (A.8)$$

Suppose that the expression in (A.8) is strictly positive (the other case is trivial). Now $\mathbb{X}$ and $\mathbb{B}$ are independent of q, K and γ, while the integrand is increasing in q and decreasing in K and γ, so the value in (A.8) inherits these properties. On the other hand, ϕ only enters via the principal's choice set $\mathbb{X}$; if the constraint $x_0 \geq \phi$ does not bind, then the principal is unaffected, but if it does bind, she is made worse off by an increase in ϕ. Finally, the result with respect to r is proved separately in Lemma A.4 below.

□

Lemma A.4. *In the low, moderate, and high knowledge cases, the principal's normalized payoff is decreasing in the discount rate.*

Proof. We begin with the high knowledge case. Define $z \equiv \frac{q}{(K-1)(1-q)}$, and note that $z \in (0,1)$ as $q \in \left(0, \frac{K-1}{K}\right)$. The principal's normalized payoff is

$$\begin{aligned}\Pi_H &= \int_0^{t^*} \phi\left[q-(1-q)(K-1)e^{-\gamma t}\right] re^{-rt}dt + \int_{t^*}^{\infty} \left[q-(1-q)(K-1)e^{-\gamma t}\right] re^{-rt}dt \\ &= \frac{r\phi[Kq-(K-1)] + q\gamma[z^{r/\gamma}(1-\phi)+\phi]}{r+\gamma}.\end{aligned}$$

Differentiating w.r.t. r yields

$$\frac{d\Pi_H}{dr} = \frac{-\gamma[q(1-\phi)z^{r/\gamma} + (k-1)\phi(1-q)] + qz^{r/\gamma}(r+\gamma)(1-\phi)\ln(z)}{(r+\gamma)^2}.$$

The expression in square brackets is strictly positive and the last term of the numerator is strictly negative, so the full expression is unambiguously strictly negative.

For the low knowledge case, the principal's normalized payoff is

$$\begin{aligned}\Pi_L &= \int_0^{\bar{t}_L} \underline{x}e^{(r+\gamma/k)t}\left[q-(1-q)(K-1)e^{-\gamma t}\right] re^{-rt}dt \\ &\quad + \int_{\bar{t}_L}^{\infty} \left[q-(1-q)(K-1)e^{-\gamma t}\right] re^{-rt}dt \\ &= \frac{q\gamma\left(\frac{Kq}{K-1}\right)^{Kr/\gamma}}{r+\gamma}.\end{aligned}$$

Now the numerator is decreasing in r as $\frac{Kq}{K-1} < 1$ and the denominator is increasing in r, so we conclude that $d\Pi_L/dr < 0$.

For the moderate knowledge case, we have

$$\begin{aligned}\Pi_M &= \int_0^{\underline{t}_M} \phi\left[q-(1-q)(K-1)e^{-\gamma t}\right] re^{-rt}dt \\ &\quad + \int_{\underline{t}_M}^{\bar{t}_M} \phi e^{(r+\gamma/K)(t-\underline{t}_M)}\left[q-(1-q)(K-1)e^{-\gamma t}\right] re^{-rt}dt \\ &\quad + \int_{\bar{t}_M}^{\infty} \left[q-(1-q)(K-1)e^{-\gamma t}\right] re^{-rt}dt \\ &= \frac{-\phi r(r+\gamma)(K-1) + q(Kr+\gamma)\left[\gamma\alpha^{-Kr/\gamma}X^{r/\gamma} - (r+\gamma)\left(X^{r/\gamma}-1\right)\right]}{(r+\gamma)^2},\end{aligned} \tag{A.9}$$

where $\alpha \equiv \left(\frac{\overline{x}}{\phi}\right)^{\frac{\gamma}{Kr+\gamma}}$ and $X \equiv \frac{q(1-K+K\alpha)}{(1-q)(K-1)}$. Note that in the moderate knowledge case, $\alpha \in \left(1, \frac{K-1}{Kq}\right]$ and $X \in \left(\frac{q}{(1-q)(K-1)}, 1\right]$.

Differentiating (A.9), we obtain

$$
\begin{aligned}
d\Pi_M/dr &= -\frac{g(X;r,\gamma,q,K,\phi)}{(r+\gamma)^3}, \qquad \text{where} \\
g(\hat{X};r,\gamma,q,K,\phi) &\equiv \phi\gamma(r+\gamma)(K-1)\left(1-q+q\hat{X}^{r/\gamma}\right) \\
&\quad + q(r+\gamma)(Kr+\gamma)\hat{X}^{r/\gamma}\ln(\hat{X})\left(\phi[1+r/\gamma]-\alpha^{-Kr/\gamma}\right) \\
&\quad + q\hat{X}^{r/\gamma}\alpha^{-Kr/\gamma}\left(\gamma[r+\gamma(2-K)]+K[r+\gamma]^2\ln[\alpha]\right).
\end{aligned}
$$

For later purposes, we establish that $\phi[1+r/\gamma]-\alpha^{-Kr/\gamma}<0$. Using the definitions of α and $\overline{x}$,

$$
\begin{aligned}
\phi[1+r/\gamma]-\alpha^{-Kr/\gamma} &= \phi\left[\frac{r+\gamma}{\gamma}-\phi^{-\frac{\gamma}{Kr+\gamma}}\overline{x}^{-\frac{Kr}{Kr+\gamma}}\right] \\
&< \phi\left[\frac{r+\gamma}{\gamma}-\overline{x}^{-\frac{\gamma}{Kr+\gamma}}\overline{x}^{-\frac{Kr}{Kr+\gamma}}\right] \\
&= \phi\frac{r+\gamma}{\gamma}\left[1-\frac{K}{K-1}\right]<0.
\end{aligned}
$$

We now show that for all $\hat{X}\in(0,1]$ (and in particular for $\hat{X}=X$), $g(\hat{X};r,\gamma,q,K,\phi)>0$, establishing that $d\Pi_M/dr<0$. To this end, we show that $g:\hat{X}\mapsto g(\hat{X};r,\gamma,q,K,\phi)$ is quasiconcave and that $\lim_{\hat{X}\downarrow 0} g(\hat{X};r,\gamma,q,K,\phi)>0$ and $g(1;r,\gamma,q,K,\phi)>0$:

- g is quasiconcave in $\hat{X}$: we have

$$
\begin{aligned}
\frac{\partial}{\partial\hat{X}}g(\hat{X};r,\gamma,q,K,\phi) = q\hat{X}^{\frac{r}{\gamma}-1}\{&\phi r(r+\gamma)(K-1) \\
&+q(r+\gamma)(Kr+\gamma)\left(\phi[1+r/\gamma]-\alpha^{-Kr/\gamma}\right)\left(\ln[\hat{X}]r/\gamma+1\right) \\
&+\alpha^{-Kr/\gamma}(r/\gamma)\left(\gamma[r+\gamma(2-K)]+K[r+\gamma]^2\ln[\alpha]\right)\}.
\end{aligned} \tag{A.10}
$$

  Observe that in (A.10), the factor outside the braces is positive, and let $g_2(\hat{X};r,\gamma,q,K,\phi)$ denote the expression inside the braces. As argued above, the coefficient on $\ln[\hat{X}]$ in $g_2(\hat{X};r,\gamma,q,K,\phi)$ is negative; hence, as $\hat{X}\downarrow 0$, the $g_2(\hat{X};r,\gamma,q,K,\phi)\uparrow\infty$, and thus for sufficiently small $\hat{X}$, $\frac{\partial}{\partial\hat{X}}g(\hat{X};r,\gamma,q,K,\phi)>0$. Moreover, we have that $g_2(\hat{X};r,\gamma,q,K,\phi)$ is monotonically decreasing in $\hat{X}$, giving us the quasiconcavity of $g(\hat{X};r,\gamma,q,K,\phi)$ w.r.t. $\hat{X}$.

- $\lim_{\hat{X}\downarrow 0} g(\hat{X};r,\gamma,q,K,\phi)>0$: By taking limits directly as $\hat{X}\downarrow 0$ and applying L'Hôpital's rule to obtain $\hat{X}^{r/\gamma}\ln[\hat{X}] \uparrow 0$, we have

$$\lim_{\hat{X}\downarrow 0} g(\hat{X}; r, \gamma, q, K, \phi) = \phi\gamma(r+\gamma)(K-1)(1-q) > 0.$$

- $g(1; r, \gamma, q, K, \phi) > 0$: By plugging in $\hat{X} = 1$ and simplifying, we get

$$g(1; r, \gamma, q, K, \phi) = (K-1)\gamma(r+\gamma)\phi + q\alpha^{-Kr/\gamma}\left(\gamma[r+(2-K)\gamma] + K[r+\gamma]^2\ln[\alpha]\right). \tag{A.11}$$

  Observe that α is independent of q, and thus the right hand side of (A.11) is linear in q. To show that it is positive, it suffices to show positivity for extreme values of q. For the moderate knowledge case, $\overline{x}$ is independent of q so $\phi < \overline{x}$ places no tighter lower bound on q besides $q \geq 0$. But $\underline{x} \leq \phi$ implies $q \leq \overline{q} \equiv \frac{K-1}{K}\left(\frac{\phi}{\overline{x}}\right)^{\frac{\gamma}{\gamma+Kr}}$. Evaluating at these extremes, we have

$$g(1; r, \gamma, 0, K, \phi) = (K-1)\gamma(r+\gamma)\phi > 0$$
$$g(1; r, \gamma, \overline{q}, K, \phi) = (r+\gamma)^2\phi\left[1 + K(1+r/\gamma)\ln(\alpha)\right] > 0.$$

  Hence, $g(1; r, \gamma, q, K, \phi) > 0$ for all q for which the moderate knowledge case applies.

Together, these claims imply that $g(\hat{X}; r, \gamma, q, K, \phi) > 0$ for all $\hat{X} \in (0, 1]$, and thus $d\Pi_M/dr < 0$, concluding the proof. □

A.13 Proof of Lemma 7

Fix any information policy, and let μ and σ^2 be the induced mean and variance processes for the agent's posterior beliefs. By the law of iterated expectations, μ must be a martingale. Let $\mu_t^{(2)} = \mathbb{E}_t[\omega^2]$ be the uncentered second moment of the agent's time-t beliefs about ω. Then the time-s posterior variance of the agent's beliefs may be written

$$\sigma_t^2 = \mu_t^{(2)} - \mu_t^2.$$

Taking time-s expectations of both sides for any $s < t$ and using the law of iterated expectations and Jensen's inequality yields

$$\mathbb{E}_s[\sigma_t^2] = \mu_s^{(2)} - \mathbb{E}_s\left[(\mathbb{E}_t[\omega])^2\right] \leq \mu_s^{(2)} - \mu_s^2 = \sigma_s^2.$$

So σ^2 must be a supermartingale. Taking time-zero expectations of both sides establishes that x must be monotone increasing, with in particular $x_t \geq x_{0-} = \phi$. The upper bound $x_t \leq 1$ follows from the fact that σ^2 is a non-negative process. Finally, as both μ and $\mu^{(2)}$ are martingales on a filtration satisfying the usual conditions, σ^2 is càdlàg (possibly by passing to an appropriate version). Meanwhile x is a monotone function, thus has well-defined limits from both sides everywhere. And by Fatou's lemma, $\mathbb{E}[\sigma_t^2] \leq \liminf_{s\downarrow t} \mathbb{E}[\sigma_s^2]$ for every t. Hence $x_t \geq x_{t+}$, and since x is monotone this implies $x_t = x_{t+}$. So x is càdlàg.

A.14 Proof of Lemma 8

In the high-knowledge case, x^* can be implemented via disclosure of ω at time t^*.

In the moderate-knowledge case, x^* is implemented by a combination of two signals. First, over the time inteval $[\underline{t}_M, \overline{t}_M]$ the agent observes a signal process ξ distributed as

$$d\xi_t = \omega\, dt + (1 - x_t^*)\sqrt{\frac{C}{(r + \gamma/K)x_t^*}}\, dB_t,$$

where B is a standard Brownian motion independent of ω and the exogenous signal. Second, at time $\overline{t}_M$ the agent observes ω. Under standard results on Brownian filtering (for example, see Liptser and Shiryaev (2013)), observing the signal process ξ induces a sequence of normal posteriors with variance process σ^2 evolving as

$$\frac{d\sigma^2}{dt} = -\sigma_t^4(1 - x_t^*)^{-2}\frac{(r + \gamma/K)x_t^*}{C},$$

with initial condition $\sigma^2_{\underline{t}_M} = \eta^2$. Note that $\sigma_t^2 = C(1 - x_t^*)$ satisfies this ODE with the correct initial condition, given that $\dot{x}_t^* = (r + \gamma/K)x_t^*$ on $[\underline{t}_M, \overline{t}_M)$. So the agent's posterior variance evolves as $\sigma_2^2 = C(1 - x_t^*)$ on this interval. Meanwhile on $[0, \underline{t}_M]$ the agent's posterior variance is fixed at η^2, while on $[\overline{t}_M, \infty)$ the posterior variance is 0. Thus the disclosure path induced by this information policy is indeed x^*, and the policy is deterministic.

Finally, in the low-knowledge case x^* is implemented by a combination of three signals. First, at time zero the agent observes a signal distributed as $\mathcal{N}\left(\omega, ((C(1 - \underline{x}))^{-1} - \eta^{-2})^{-1}\right)$, independent of the exogenous signal conditional on ω. Second, over the time interval $[0, \overline{t}_L]$ the agent observes a signal process ξ distributed as in the moderate-knowledge case, with the Brownian motion independent of the time-zero signal. Third, at time $\overline{t}_L$ the agent observes ω. Bayesian updating implies that after observing the time-zero signal, the agent's posterior beliefs are normally distributed with variance $\sigma_0^2 = C(1 - \underline{x})$. Meanwhile on $[0, \overline{t}_L]$ the work for the moderate-knowledge case shows that the agent's posterior beliefs have variance

$\sigma_t^2 = C(1 - x_t^*)$. Finally, the agent's posterior variance on $[\bar{t}_L, \infty)$ is 0. Thus the disclosure path induced by this information policy is x^*, and the policy is deterministic.