Better Monitoring ... Worse Productivity?

John Y. Zhu[1]

August 1, 2019

Abstract

Changing workplace demands are driving managers to collect more information and provide more feedback about worker performance than ever before. Despite obvious benefits concerns persist about how all this information – much of which is non-contractible and must pass through discretionary feedback – might distort incentives. I highlight a better monitoring/worse outcome channel that speaks to these concerns. Some improvements to the informativeness of monitoring tempt managers to provide excessive negative feedback leading to overpunishment. Workers then refuse to accept contracts that do not severely constrain the size of the punishment threat. Without a serious punishment threat, effort and surplus decline.

JEL Codes: C73, D86, J41, M5
Keywords: incentive power, statistical power, monitoring, privacy, management style, information design, optimal contracts, principal agent, moral hazard

[1]I thank Dilip Abreu, Davide Cianciaruso, Tommaso Denti, Brett Green (discussant), David Pearce, Bill Sandholm, Eddie Schlee and seminar participants at UC Boulder, Cornell, University of Rochester, University of Minnesota, University of Washington at St Louis, INSEAD, HEC, University of Maryland, UC Riverside, University of Wisconsin, Financial Research Association Conference, Richmond Fed, University of Kentucky, Michigan State, Kansas Workshop on Economic Theory.

1 Introduction

A performance management revolution is reshaping the nature of monitoring in the workplace, emphasizing more information and more feedback. In just the last few years hundreds of companies ranging from small private firms to multinationals, non-profits and NGOs have introduced continuous performance management (CPM) practices as a complement to and sometimes as a replacement for the traditional end-of-year appraisal.[2] At the same time, technological advances such as those in biometrics have vastly increased the quantity of worker data that can be collected at any moment in time.[3] Despite obvious benefits of better monitoring (e.g. better worker development, quicker identification of problem areas, improved coordination amongst team members) there is widespread uncertainty about how all this extra information will affect the provision of incentives since much of it is not directly contractible and ends up being filtered through a manager's discretionary feedback. Given the rapidly expanding scope of monitoring, understanding if, when, and how monitoring should be limited is becoming increasingly important.

That monitoring should be limited at all may be surprising given that a fundamental result in principal-agent theory is about how better monitoring generically leads to a better outcome (Holmstrom, 1979). In light of this result, rationales for why reducing the informativeness of monitoring is beneficial have focused on introducing additional concerns into the baseline principal-agent model that can antagonize the otherwise positive relationship between principal monitoring and agent effort. One well-known example is career concerns.[4] Other possibilities include the fear of corruption or that "familiarity breeds contempt."[5]

The point I wish to make about the potential benefits of reducing the amount of information generated by monitoring is in some ways more universal. Rather than adding another dimension to the baseline model, I start by observing that Holm-

[2] CPM practices emphasize providing frequent feedback, often in ratingless form and sometimes drawn from multiple sources (Ledford, Benson, and Lawler, 2016). Numerous recent articles in *Harvard Business Review* – Buckingham and Goodall (2015), Cappelli and Tavis (2016), *Wall Street Journal* – Weber (2016), Hoffman (2017), and *Forbes* – Burkus (2016), Caprino (2016) have documented the shift toward CPM at companies across a broad range of industries, including *Google*, *Deloitte*, *Patagonia*, *Adobe*, *General Electric*, *Goldman Sachs*, *Kimberly-Clark*, and *Accenture*. John Doerr, venture capitalist at *Kleiner Perkins*, has also written about CPM. See Doerr (2018).

[3] While much attention has been paid to the kind of biometric technology used at Amazon warehouses that monitors discrete tasks, sophisticated monitoring technologies featuring machine learning and artificial intelligence are increasingly being deployed to evaluate performance in less structured environments. For example, *Humanyze* tracks vocal data including tone, speed, volume, and frequency. Algorithmic software then processes that data to help clients interpret office communication patterns and their impact on productivity.

[4] The canonical moral hazard model with career concerns is Holmstrom (1999). Crémer (1995), Dewatripont, Jewitt and Tirole (1999) and Prat (2005) explore ways in which better monitoring can lead to worse outcomes in various moral hazard models with career concerns.

[5] Outside the corporate world, the anti-fraternization rules between enlisted personnel and officers tasked with monitoring them are often justified based on these concerns.

strom's better monitoring/better outcome result assumes monitoring generates purely contractible information but for many agents monitoring generates information that is, at least partially, not contractible – think about the opinions a manager forms in her mind about how a worker is doing.[6] When information is not directly contractible, incentives must depend on the principal's reports of what she has observed. This means even if the principal and agent sign a contract, the principal still has quite a bit of discretion in deciding how much to discipline the agent. In this discretionary setting, I show some improvements to the information content of monitoring tempt the principal to be too tough by giving her the ability to induce a lot of effort from the agent but only if she makes heavy use of the "stick." Anticipating such a tough principal, the agent responds by refusing to accept any contract that gives the principal a big stick. The end result is that either the agent quits or the principal is forced to offer a new weaker contract with a stick so small that effort and surplus decline despite the better information. Conversely, reducing the amount of information generated by monitoring can be beneficial by reducing how tough the principal is tempted to be. A consequence of this result is that appraising an agent's overall performance every once in a while can dominate closely monitoring his day-to-day performance every single day.

Taken together these results highlight a *better monitoring/worse outcome channel* relevant across a broad range of principal-agent relationships where the principal has discretion in deciding how information generated by monitoring is used to provide incentives for the agent. In the coming sections I will show how the presence of such a channel means that when it comes to the design of monitoring technologies care should be taken to avoid generating information that is at once noisy but sensitive to effort (Section 4.2), that there is value to censoring raw performance data (Section 4.3), and that maintaining formal, periodic performance reviews can – if done in the right way – be beneficial even when organizational concerns beyond incentives necessitate a more frequent performance management component (Section 5.3).

The idea that more information can lead to a worse contracting outcome has also been pointed out in the insurance market setting: Limiting what counterparties know about hidden states can make everyone better off ex-ante.[7] My work can be viewed as exploring an analogous phenomenon for hidden actions.[8]

[6]Even if monitoring is in the form of a data generating technology, in practice that data is often still not contractible especially if the worker is not performing simple, repetitive tasks. The data is typically proprietary and observed only by the firm. It may be interpretable only in conjunction with other soft information. Directly conditioning outcomes on the data may be impractical if the dataset is huge, susceptible to manipulation, or evolving over time in response to changing market conditions that are hard to predict ex-ante. Worker privacy concerns also impede contractibility. Instead, what often happens is the manager uses the data to help make an unverifiable judgement call about worker performance. For example, one client used *Humanyze* data to help determine which teams seemed more crucial and which ones less so when deciding how to reposition personnel ahead of a major growth opportunity. (A Major Oil and Gas Company Faces Expansion, n.d.)

[7]See, for example, Hirshleifer (1971), Wilson (1975), and Schlee (2001).

[8]Gjesdal (1982) shows that when utility is nonseparable and contracts are deterministic better

2 A Motivating Example

A principal P (she) contracts an agent A (he) to work on a project. A exerts hidden effort $a \in [0, 1)$ with cost $h(a)$. Assume h is convex and differentiable with $h'(0) = 0$ and $h'(a) \to \infty$ as $a \to 1$. a determines the distribution of a signal $X \in \{x_B, x_G\}$ where $\mathbf{P}(X = x_B \mid a) = (1 - a)q$ for some constant $q \in (0, 1)$. x_B is a "bad" signal – the higher is a the less likely it occurs. Conversely, x_G is a "good" signal.

X is privately observed by P. However, nothing would change if instead X were mutually observable but unverifiable and A does not make verifiable reports about X that a contract could depend on. A's lack of reporting could be due to P's ability to falsify A's reports or perhaps it is too costly for A to take P to court if P violates a contract's dependence on A's reports. Given X, P receives a utility with $u(x_G) > u(x_B)$. Define $u(a) := \mathbf{E}_a u(X)$.

A *contract game* consists of an upfront payment $w \in \mathbb{R}$ and a punishment $p \geq 0$ that P can inflict on A after privately observing X. Punishing A does not affect P's utility but subtracts p from A's utility. P's punishment strategy is a mapping r from X to a possibly random decision to punish ($r = 1$) or not punish ($r = 0$). Given contract game (w, p), effort a, and punishment strategy r, P's payoff is $u(a) - w$. A's payoff is $w - h(a) - \mathbf{E}_a rp$.[9]

I assume the following contract negotiation protocol is in place: P offers a contract game and recommends an incentive-compatible way to play the game – that is, an equilibrium. A chooses whether or not to accept P's contract game offer. If he accepts he obeys P's recommendation about play. If he does not accept both parties exercise outside options normalized to 0.

Does it matter when P makes her recommendation about play? In theory, she can recommend an equilibrium at the time she offers a contract game and then, if A accepts the offer, recommend a different one. The ability to recommend again matters because after acceptance A's ex-ante participation constraint is no longer a concern. I assume if P does change her recommendation after A accepts then A obeys the latter recommendation. Equivalently, P only recommends an equilibrium after A accepts the contract game offer. For more on this contract negotiation protocol see the discussion below titled *Why Does the Principal Get to Dictate the Equilibrium?*

Let us now find the optimal contract game (w^*_X, p^*_X) under monitoring technology X. Given an offer (w, p), if A accepts P will subsequently recommend the equilibrium (a, r) with maximal a. Why? After A accepts, w is fixed and he can no longer exercise his outside option. This makes P's payoff $u(a) - w$ purely an increasing function of

monitoring in the sense of Blackwell (1950) can lead to a less efficient outcome. Essentially garbling allows deterministic contracts to mimic random contracts which can be beneficial under nonseparable utility. If contracts are allowed to be random better monitoring/better outcome is restored.

[9] My model is in the spirit of the subjective evaluation models found in MacLeod (2003) and Fuchs (2007). Those papers show that since monitoring is private incentives are provided through surplus destruction not performance sensitive pay. Thus, it is without loss of generality for there to be only an upfront payment.

a. The unique equilibrium with maximal a involves P punishing A if and only if the bad signal x_B is realized. When making his accept/reject decision, A anticipates that if he accepts P's final recommended equilibrium will be the one just described.

Taking the equilibrium as given, p_X^* is set to maximize surplus:

$$p_X^* \in \arg\max_{p \geq 0} u(h'^{-1}(qp)) - h(h'^{-1}(qp)) - (1 - h'^{-1}(qp))qp.$$

Let S_X^* denote the maximal surplus. If $S_X^* \geq 0$ an optimal contract game satisfying ex-ante participation constraints exists. A standard first-order condition now pins down A's effort as $a_X^* = h'^{-1}(qp_X^*)$. A's ex-ante participation constraint then pins down $w_X^* = h(a_X^*) + (1 - a_X^*)qp_X^*$.

Notice under X the optimal contract game and the way A and P play that game would have been the same even if P could not change her recommendation about play after A accepts her contract game offer. In general, however, the ability of P to change her recommendation after A accepts negatively impacts contracting:

Consider the following *better monitoring technology* X' derived from X by splitting x_G into two signals x_b and x_g where the probability of x_b decreases linearly from $1-q$ to $1-q-\varepsilon$ as a function of a over $[0,1)$. Here think of ε as being vanishingly small. Notice x_g is a good signal while x_b is an almost completely uninformative bad signal. Utility remains unchanged: $u(x_b) = u(x_g) = u(x_G)$.

Let us now find the optimal contract game $(w_{X'}^*, p_{X'}^*)$ under the better monitoring technology X'. Given an offer (w, p), A again rationally anticipates being punished whenever a bad signal is realized – except this time a bad signal is x_B or x_b. Since ε is vanishingly small, the effort induced by a punishment of size p under X' is arbitrarily close to that under X. Thus the argmax expression used to find an optimal punishment is approximately the one used before plus a $-(1-q)p$ term:

$$\arg\max_{p \geq 0} u(h'^{-1}(qp)) - h(h'^{-1}(qp)) - (1 - h'^{-1}(qp))qp - (1-q)p \tag{1}$$

This implies $S_{X'}^* < S_X^*$. If $S_{X'}^* < 0$ then both parties quit the relationship. Otherwise, an optimal contract game exists satisfying ex-ante participation constraints exists. Either way better monitoring has led to less surplus and a lower payoff for P. Moreover, if optimal punishments under X and X' are unique then optimal punishment and effort both decline: $p_{X'}^* < p_X^*$ and $a_{X'}^* < a_X^*$.

Why can't P just ignore the extra information contained in X', rendering better monitoring/worse outcome impossible? Because using the extra information of X' allows more effort to be induced, and maximizing effort is all P cares about after A accepts the contract game offer. But that extra effort comes at a great cost to efficiency because P now punishes the almost completely uninformative x_b in addition to x_B. Essentially P is being too tough on A. Allowing contract games to feature both a big and a small punishment – so that the punishment can fit the crime – would not help: After A accepts such a contract game offer, P will recommend a play that has A suffering the big punishment for both x_b and x_B. Ex-ante A anticipates that P cannot help but be too tough under the better monitoring technology X'. To prevent A from quitting, P is forced to offer a contract game with a small punishment, which hinders incentive provision and leads to a worse outcome.[10]

The motivating example suggests the following channel through which better monitoring can negatively affect the provision of incentives: *When an improvement to monitoring yields lots of low quality bad signals the principal wants to use the stick too often. The agent then refuses to accept a contract game that gives the principal a big stick. Without a big stick to discipline the agent effort and surplus decline.*

Is Better Monitoring/Worse Outcome Fragile?

In the motivating example P cannot stick to a recommendation about play that becomes even a tiny bit suboptimal after A accepts the contract game offer: Under X', P can induce slightly more effort and increase her utility a tiny bit after A accepts than under X by punishing the almost completely uninformative signal x_b. P's lack of commitment power then implies x_b will be punished, leading to a worse contracting outcome ex-ante. With ε-commitment power P could recommend punishing only x_B before A accepts a contract game offer and stick to that recommendation after A accepts. The worse outcome result then goes away.

This begs the question: Does better monitoring/worse outcome survive in general when P has ε-commitment power? Yes.

I will show that in general there are ways to improve monitoring that allow P to induce significantly higher effort but require an even more significant amount of punishment. Having ε-commitment power allows P to resist the temptation of inducing slightly higher effort but not significantly higher effort. By assumption, inducing significantly higher effort involves being too tough. Now the same logic as before leads to a worse outcome.

In the motivating example, P also lacks commitment power after observing X: The definition of a contract game only allows for an upfront payment because even if P were provided a menu of payments to choose from after observing X, she would only ever choose the minimum payment. With ε-commitment, P could commit to mappings from X to payments in the menu within ε of the minimum, allowing for a modicum of pay-to-performance sensitivity. Of course, this alternate source of effort incentives is not enough to render the better monitoring/worse outcome channel obsolete if ε is sufficiently small.

On the other hand, as ε tends to infinity, P will abandon the punishment tool p in favor of performance sensitive pay, which does not involve surplus destruction. The model becomes isomorphic to a standard moral hazard model with contractible

[10] Alternatively, P could keep the punishment size the same and just compensate A with a bigger upfront payment. However (1) implies this is never optimal – reducing punishment increases surplus which P can then capture ex-ante by lowering the upfront pay.

information and Holmstrom (1979) applies: Better monitoring not only never leads to a worse outcome but, in fact, generically leads to a strictly better outcome if A is risk-averse. In the motivating example both P and A are risk-neutral but clearly introducing risk-aversion or limited liability will not shut the better monitoring/worse outcome channel.

In the motivating example the information generated by monitoring is purely non-contractible but the better monitoring/worse outcome channel can remain open even if information contains a contractible component. Of course, as the contractible component becomes increasingly informative, the channel will weaken.

In the motivating example X has the property that each signal realization can be characterized as "good" or "bad" because its occurrence probability is a strictly monotonic function of effort. While in general signals need not be monotonic in this way, it is not clear that non-monotonicity would negatively interfere with the better monitoring/worse outcome channel. That being said, assuming monotonic signals simplifies the analysis of optimal contracting since it implies a very simple punishment strategy for maximizing effort. Throughout the rest of paper I continue to work with monotonic signals or signals that satisfy the monotone-likelihood-ratio-property (MLRP), another class of signals that implies a similarly simple punishment strategy for maximizing effort.

Why Does the Principal Get to Dictate the Equilibrium?

A's effort is hidden and P's action depends on non-contractible information. Thus, neither player's strategy can be court-enforced. Instead, the players must somehow settle outside of the courts on a way to play the contract game after it is accepted by A.

The contract negotiation protocol assumes that after A accepts a contract game offer, the players always settle on playing P's most preferred equilibrium. This protocol is important for better monitoring/worse outcome and I justify it by appealing to the credible threats idea of Tranæs (1998) and Zhu (2018) for selecting equilibria. See Dewatripont (1987) and Barron and Guo (2019) for additional applications of the idea.

Here is the gist of the credible threats idea: Consider a two player perfect information game of depth two with the property that player 2 is indifferent between all her actions at each of her decision nodes. Ex ante, player 2 ought to be able to (and will want to) commit to a strategy in order to induce player 1 to take an action most preferred by player 2. Tranaes (1998) generalizes this idea to select subgame-perfect equilibria in all finite extensive-form games. Zhu (2018) extends this idea to show that in repeated moral hazard models with private monitoring a unique (up to continuation payoff process) sequential equilibrium is selected in every contract game.

In the motivating example applying the credible threats idea leads to the equilibrium selected by the contract negotiation protocol: Notice P – player 2 – is indifferent between her two actions, punish and do not punish. The credible threats idea then says that, before the contract game is played, P can and will credibly commit to a punishment strategy that induces A – player 1 – to take the action most preferred by P. P's most preferred action is the highest effort A can be induced to exert.

In Tranæs (1998) and Zhu (2018), the credible threats idea only has bite in the knife edge case when a player is indifferent between actions. In the motivating example, the knife edge case is ensured because punishing A does not affect P's utility. But what if punish were, say, ε more costly to P than do not punish? As long as P has ε-commitment power, the credible threats idea still has bite under the contract negotiation protocol and the better monitoring/worse outcome channel remains open.

Looking Ahead

While it should be clear now that the better monitoring/worse outcome channel applies quite generally, a few important questions remain:

1. Are there ways to improve monitoring that are more natural than splitting a signal realization and still generate the better monitoring/worse outcome result?
2. How much worse can the outcome be when monitoring is improved?
3. What kinds of improvements to monitoring lead to a worse outcome?
4. Can the better monitoring/worse outcome channel provide practical guidance for how to beneficially limit monitoring?

In the remainder of the paper I address these questions. I begin by moving to a dynamic setting where the abstract utility destroying punishment tool p of the motivating example is replaced with the ability to terminate the relationship. This allows the scope for punishing the agent at any date to be endogenously bounded by the forgone surplus of the relationship going forward. I then:

1. Look at the universe of binary-valued monitoring technologies in the continuous-time limit and consider a situation where the principal begins with a single binary-valued technology X and then improves it by acquiring an additional conditionally independent binary-valued technology Y.
   - Acquiring an extra signal of effort represents a more intuitive way to improve monitoring than splitting a realization of a single signal into two signal realizations. The choice to restrict attention to all binary-valued monitoring technologies in the continuous-time limit represents a compromise between computational simplicity and generality.
2. Demonstrate that better monitoring can lead to a significantly worse outcome.

- In the leading example of the analysis I show that improving a bad news Poisson monitoring technology by bringing in an additional Brownian component causes the optimal contract game to collapse into a trivial arrangement that induces zero effort at all times and never terminates the agent.

3. Characterize the counterproductive improvements Y as those that are, relative to X, sufficiently strong in *incentive power* but sufficiently weak in *statistical power.*

   - Greater incentive power means being able to use a smaller punishment threat to induce any target effort level. Statistical power, to be defined later, is a measure of informativeness based on viewing the information generated by monitoring as a hypothesis test.

4. Show, as an application, that in a canonical setting in which the agent's efforts affect the drift of a Brownian process and the principal monitors by sampling that process, restricting the principal to sample every once in a while significantly improves the outcome.

   - In contrast, letting the principal sample continuously but only allowing her to view the results of those samplings every once in a while (i.e. batching) is strictly counterproductive. This contrasts with well-known results in the repeated games literature that highlight the benefits of batching. I then show how these findings shed light on ways to appropriately maintain a periodic performance appraisal for incentive provision purposes when other organizational concerns necessitate a continuous performance management component.

3 The Dynamic Model

The contracting horizon runs from date 0 to date T. Each date is of length $\Delta > 0$ and dates are denoted by $t = 0, \Delta, 2\Delta, \dots T$. Assume Δ divides T. The discount factor is $e^{-r\Delta}$ for some $r > 0$.

At the beginning of each date $t < T$ that A is still employed, P pays A some amount $w_t\Delta \in \mathbb{R}$. Next, A chooses effort $a_t \in [0, 1)$. a_t costs $h(a_t)\Delta$ with $h(0) = h'(0) = 0$, $h'' > 0$, and $\lim_{a_t \to 1} h(a_t) = \infty$. After A exerts effort P observes a private signal X_t smoothly controlled by a_t. Again, it is equivalent to assume X_t is only noncontractible but A cannot make reports about X_t. I assume X_t is strictly monotone in the sense that there exist two disjoint sets *Good* and *Bad* such that $Im(X_t) = Good \sqcup Bad$ and for any $x \in Good$ (Bad), $\mathbf{P}(X_t = x \mid a_t)$ is strictly increasing (decreasing) in a_t. X_t determines P's date t utility. Given a_t, I assume P's date t expected utility can be written as $u(a_t)\Delta$ where $u(\cdot)$ is a strictly increasing, weakly concave function and $u(0) > 0$. Next, P reports a public message m_t selected from a contractually pre-specified finite set of messages $\mathcal{M}_t$; then, a public randomizing device is realized; finally, A is possibly terminated at the beginning of date $t+\Delta$. If A is terminated A and P exercise outside options worth 0 at date $t+\Delta$ and P makes a final payment $w_{t+\Delta}$ to A. Otherwise, the same sequence of events as I just described is repeated for date $t+\Delta$. A is terminated at the beginning of date T.

A *contract game* $(\mathcal{M}, w, \tau)$ specifies a message book $\mathcal{M}$, a payment plan w, and a termination clause τ. Let h_t denote the public history of messages and public randomizing devices up through the end of date t. $\mathcal{M}$ consists of an $h_{t-\Delta}$-dependent finite message space $\mathcal{M}_t$ for each t. τ is a stopping time where $\tau = t$ is measurable with respect to $h_{t-\Delta}$. w consists of an $h_{t-\Delta}$-measurable payment $w_t\Delta$ (if $\tau(h_{t-\Delta}) > t$) or w_t (if $\tau(h_{t-\Delta}) = t$) to A for each t.

Given $(\mathcal{M}, w, \tau)$, an *assessment* (a, m) consists of an effort strategy a for A, a report strategy m for P, and a system of beliefs. a consists of an effort choice a_t for each t depending on $h_{t-\Delta}$ and A's private history $H^A_{t-\Delta}$ of prior effort choices. m consists of a message choice m_t for each t depending on $h_{t-\Delta}$ and P's private history H^P_t of observations $\{X_s\}_{s\leq t}$. The system of beliefs consists of a belief about $H^P_{t-\Delta}$ at each decision node $(H^A_{t-\Delta}, h_{t-\Delta})$ of A, and a belief about H^A_t at each decision node $(H^P_t, h_{t-\Delta})$ of P.

A *contract* $(\mathcal{M}, w, \tau, a, m)$ is a contract game plus an assessment. Given a contract, the date t continuation payoffs of A and P at the beginning of date t are

$$W_t(H^A_{t-\Delta}, h_{t-\Delta}) = \mathbf{E}^A_t \left[\sum_{t\leq s<\tau} e^{-r(s-t)}(w_s - h(a_s))\Delta + e^{-r(\tau-t)} w_\tau \right],$$

$$V_t(H^P_{t-\Delta}, h_{t-\Delta}) = \mathbf{E}^P_t \left[\sum_{t\leq s<\tau} e^{-r(s-t)}(-w_s + u(a_s))\Delta - e^{-r(\tau-t)} w_\tau \right].$$

3.1 The Optimal Contract

The optimal contracting problem is to find an incentive-compatible contract that maximizes V_0 subject to the agent's ex-ante participation constraint $W_0 \geq 0$ and an interim participation constraint $W_t + V_t \geq 0$ for all t. Intuitively, if the interim participation constraint were violated then both parties could be made strictly better off by separating under some severance pay.[11]

Incentive-compatibility requires that the principal's report strategy and the agent's effort strategy plus a system of beliefs comprise a sequential equilibrium. As we saw in the motivating example, the credible threats idea of Tranaes (1998) can be used to justify the assumption that P selects the equilibrium. In general, P's most preferred equilibrium changes over time. In the one shot setting this just means that P's most preferred equilibrium after A accepts the contract game offer is played. In the present

[11]It will be shown that for incentive compatible contracts W_t and V_t are both public, so violations of the interim participation constraint are common knowledge.

setting where P and A interact repeatedly, it is not a priori clear what equilibrium will be played. Zhu (2018) addresses this issue and argues that any reasonable application of the credible threats idea will inevitably lead to the same sequential equilibrium which at the end of each date t involves P reporting the message that leads to the highest (lowest) possible $W_{t+\Delta}$ subject to maximizing $V_{t+\Delta}$ if and only if $X_t \in Good$ ($\in Bad$).[12] This characterization of P's report strategy in any incentive-compatible contract then implies:

Theorem 1. *As $T \to \infty$, the optimal contract converges to a stationary efficiency wage contract with the following structure:*

- $\mathcal{M}_t = \{pass, fail\}$.
- $m_t = fail$ *iff* $X_t \in Bad$.
- w *consists of a pair of constants* $w_{salary}\Delta, w_{severance}$.
- *If* $m_t = pass$ *then* A *is retained for date* $t + \Delta$ *and paid* $w_{salary}\Delta$.
- *If* $m_t = fail$ *then* A *is terminated at date* $t + \Delta$ *with probability* p^*.
  - *If* A *is not terminated then it is as if* P *reported pass.*
  - *If* A *is terminated then he is paid* $w_{severance}$.

Proof. See appendix. □

The optimal contract for finite T has the same structure except w_{salary}, $w_{severance}$, and p^* will, in general, all depend on t. For computational simplicity, I will focus on the infinite horizon limit from now on.

The optimal contract is a wage contract just like what was assumed in the motivating example.[13] At each date t, conditional on still being employed, the agent is paid the same amount regardless of performance history. There is a good reason for this. Suppose instead there was an additional message that leads to A receiving a big bonus which P is supposed to report if she observes some really positive information about A's performance (i.e. a *Good* signal whose probability increases sharply as a increases). The problem with this altered contract is that its strategy profile would not satisfy any reasonable notion of incentive-compatibility: Because monitoring is private, P can always claim she didn't see the really positive information even if she

[12] That such $W_{t+\Delta}$ are well-defined and do not depend on A's beliefs about P's private history depends on applications of the credible threats idea at dates $> t$. The use of backwards induction logic to refine sequential equilibria is why the model has a finite horizon.

[13] There are two payment levels here whereas in the motivating example there was only one. The difference arises only because I assumed in the motivating example that punishing A does not affect P's utility. In the current model, failing A changes P's continuation payoff. The difference between w_{salary} and $w_{severance}$ exists purely to make P indifferent ex-post between reporting $pass$ and $fail$.

did and thereby avoid having to pay A the big bonus. In general, P must be indifferent between reporting different messages that occur on the equilibrium path, which means in the optimal contract

$$V_{t+\Delta}(pass) = V_{t+\Delta}(fail).$$

By definition, $V_{t+\Delta}(fail) = -p^* w_{severance} + (1-p^*)V_{t+\Delta}(pass)$. Let S^* denote the Pareto-optimal surplus. By self-similarity and the fact that A's ex-ante participation constraint binds, $V_{t+\Delta}(pass) = V_0 = S^*$. Thus,

$$w_{severance} = -S^*.$$

Negative severance pay is just an artifact of how I normalized outside options.

Next, consider A's effort incentives. In my model, termination destroys surplus – by assumption, even zero effort generates positive surplus. Since P is completely insured against any surplus destruction, this means it is A who bears the cost of inefficient termination,

$$W_{t+\Delta}(pass) - W_{t+\Delta}(fail) = p^* S^*.$$

Consequently, A is willing to put in effort to reduce the chances of getting failed and terminated. The first-order condition that pins down A's effort level each date is,

$$h'(a^*)\Delta = -\frac{d\mathbf{P}(X_t \in Bad \mid a_t)}{da_t}|_{a_t=a^*} p^* S^*. \tag{2}$$

If there are multiple efforts that maximize A's utility, a^* is the highest one as this is the most preferred by P.

p^* and S^* are simultaneously determined by the following system of equations,

$$\begin{aligned} p^* &= \arg\max_{p\in[0,1]} (u(a^*(pS^*)) - h(a^*(pS^*)))\Delta + e^{-r\Delta}(1 - \mathbf{P}(Bad \mid a^*(pS^*))p)S^* \\ S^* &= (u(a^*(p^*S^*)) - h(a^*(p^*S^*)))\Delta + e^{-r\Delta}(1 - \mathbf{P}(Bad \mid a^*(p^*S^*))p^*)S^*. \end{aligned}$$

The solution can be recursively computed by setting $S_0^* = u(0)\Delta$ on the RHS of the two equations and then computing p_1^* and S_1^* and so on and so forth. S_i^* is strictly increasing in $i = 0, 1, 2 \ldots$ and $S^* = S_\infty^*$. Finally, w_{salary} is determined by A's binding ex-ante participation constraint $W_0 = 0$,

$$w_{salary} = h(a^*(p^*S^*)) + \frac{e^{-r\Delta}}{\Delta}\mathbf{P}(Bad \mid a^*(p^*S^*))p^*S^*.$$

Notice the sequence of computations that lead to the full characterization of the optimal dynamic contract mirrors the one in the motivating example. Not surprisingly then, a signal splitting construction like the one considered in the motivating example can generate a better monitoring/worse outcome result here as well. However, in the sections below I explore the better monitoring/worse outcome channel at a deeper level by addressing the issues listed at the end of Section 2.

4 Better Monitoring Worse Outcome

At each date t, P punishes A if and only if she observes a *Bad* signal at that date. This simple intuitive report strategy is a consequence of P wanting to maximize effort incentives at all times. But is this the efficient thing to do? Put another way, if P were a benevolent social planner instead of a utility maximizer would she still report in this way or something close given the contract game? The answer – it depends. If the monitoring technology generates "high quality" *Bad* signals that are suggestive of A not putting in the level of effort he is supposed to then intuitively the answer is yes (I will be precise about what "high quality" means shortly). Where this strategy becomes inefficient is when the monitoring technology generates mostly "low quality" *Bad* signals that are not very suggestive of A shirking. In this case, one would like to see P be a little more fair to A and fail him only if he generates a high quality *Bad* signal, or at least wait until he has generated low quality *Bad* signals across multiple dates. But P will not be fair: Sure, at the time of contracting, P would like to commit to be fair in the future – being fair increases surplus which P can then extract by offering a contract game with lower pay. The problem is, once A has accepted the contract game, P cannot help but change her report strategy to an overly tough one that maximizes effort incentives by punishing A any time any kind of a *Bad* signal occurs. Since changing a report strategy amounts to changing a function over non-contractible information, being too tough is not something that can be contracted away.

Now at the time of contracting A understands that in the future, if the monitoring technology is going to frequently generate low quality *Bad* signals, P will likely be too tough. To counteract this, the contracting parties then preemptively agree to an optimal contract that reduces the pain of punishment. That means setting p^* to be a low value. And in some cases when the typical *Bad* signal is very low quality it might even be optimal to lower p^* all the way to zero. Of course, once p^* hits zero there is no punishment threat and A will exert zero effort.

Is it possible to take a monitoring technology that generates mostly high quality *Bad* signals and *improve* it to the point where it generates lots of low quality *Bad* signals? The motivating example suggests some dilution in the quality of *Bad* signals is possible. I now show quality can be strongly diluted, to the point of triggering a complete collapse of the optimal contract. Moreover, such counterproductive improvements to monitoring need not involve signal splitting as in the motivating example and can instead take the more intuitive form of letting the principal observe an additional conditionally independent signal of effort. Of course not all additional signals are harmful. So lastly I classify the harmful additional signals as those that are, relative to the signal already in place, sufficiently strong in *incentive power* but sufficiently weak in *statistical power*. Before establishing the result at a reasonably general level, let us first work through an explicit example that demonstrates the basic idea.

4.1 Bad News Poisson and Brownian Monitoring

In this example I begin with a bad news Poisson monitoring technology where the Poisson event is the *Bad* signal. I show that this *Bad* signal is high quality in the sense that the ex-post efficiency loss from punishing A nontrivially whenever this signal occurs is more than outweighed by the ex-ante efficiency gain from the effort such a punishment induces. Consequently the optimal contract induces positive effort – depending on how the intensity of the Poisson process and the benefit function $u(\cdot)$ are parameterized the effort induced by the optimal contract can be made to be arbitrarily high (i.e. close to 1). I then improve the monitoring technology by letting P observe an additional, conditionally independent Brownian signal of effort where the drift is controlled by A's effort (actually a random walk where effort controls the degree of asymmetry). I show that the improved monitoring technology, where signals are vectors consisting of a Poisson and a Brownian component, generates a typical *Bad* vector that is very low quality. Consequently, the optimal contract collapses and P becomes worse off.

Under bad news Poisson monitoring, each date the incremental information X_t is

$$X_t = \begin{cases} \text{no event} & \text{with probability } 1-(1-a_t)\lambda\Delta \\ \text{event} & \text{with probability } (1-a_t)\lambda\Delta \end{cases}$$

for some $\lambda > 0$ and vanishingly small Δ. It is evident that the Poisson event itself is the *Bad* signal whereas no event is the *Good* signal. What is the quality of the event signal of bad news Poisson monitoring? The measure that is of interest to me is the *negative effort-elasticity of Bad signals*:

$$\left(-\frac{d\mathbf{P}(X_t \in Bad \mid a_t)}{da_t}\right) \cdot \frac{1}{\mathbf{P}(X_t \in Bad \mid a_t)}. \tag{3}$$

The first term of this elasticity measures the *incentive power* of information. It appears in the first-order condition that pins down A's best response effort – see equation (2). The larger is this term, the smaller is the punishment threat needed to induce a target effort level. What about the second term? In equilibrium, one knows the effort a_t that is being exerted by A. Thus, when a *Bad* signal is realized and punished, it is as if P is treating the signal as evidence that A did not in fact exert the effort he was supposed to exert. This constitutes a type II error. In hypothesis testing, the less likely a type II error occurs, the more statistically powerful is the test. In my analogy the probability of making a type II error is $\mathbf{P}(X_t \in Bad \mid a_t)$.

Thus, the second term of the elasticity can be thought of as measuring the *statistical power* of information.

When P is considering how much of a punishment threat her contract offer should feature the basic tradeoff she weighs is, for a given punishment threat, how much effort will she induce versus how much surplus will be destroyed. The first factor is captured by my measure of incentive power while the second factor is captured by my measure of statistical power. Holding one factor fixed, the other factor needs to be sufficiently attractive for it to be worth it to induce effort. How much is sufficient? It turns out the answer is given by the negative effort-elasticity of *Bad* signals, which multiples my measures of incentive and statistical power.

Formally I will show if the negative effort-elasticity of *Bad* signals goes to zero as Δ tends to zero, then in the continuous time limit the optimal contract does not induce positive effort. If on the other hand the measure stays bounded away from zero, then it is possible to parameterize the rest of the model in such a way so that the optimal contract induces arbitrarily high effort. See Theorem 2 below. Anticipating this result, I now use the negative effort-elasticity of *Bad* signals to classify the quality of the typical *Bad* signal generated by a monitoring technology in the continuous time limit.

A simple computation shows that the negative effort-elasticity of *Bad* signals under bad news Poisson monitoring is

$$\frac{1}{1-a_t}.$$

Notice it remains bounded away from zero as Δ becomes small no matter the effort level. Thus, bad news Poisson monitoring generates a high quality *Bad* signal and the optimal contract under bad news Poisson monitoring can induce positive effort.

Let us now see what happens when the bad news Poisson monitoring technology is improved by including a conditionally independent Brownian signal Y_t where effort controls the drift:

$$Y_t = \begin{cases} \sqrt{\Delta} & \text{with probability } \frac{1}{2} + \frac{a_t\sqrt{\Delta}}{2} \\ -\sqrt{\Delta} & \text{with probability } \frac{1}{2} - \frac{a_t\sqrt{\Delta}}{2} \end{cases}$$

Each date the Brownian signal is a single step of an extremely fine random walk. Whenever the random walk goes up it is a *Good* signal, whenever it goes down it is a *Bad* signal.

Under the improved monitoring technology, a signal is a vector (X_t, Y_t). Obviously when both components are *Bad* (*Good*) the vector is *Bad* (*Good*). But what about the other two cases? Notice, Brownian information has really strong incentive power, at least relative to bad new Poisson information:

$$-\frac{d\mathbf{P}(Y_t = -\sqrt{\Delta} \mid a_t)}{da_t} = \frac{\sqrt{\Delta}}{2} \gg \lambda\Delta = -\frac{d\mathbf{P}(X_t = \text{event} \mid a_t)}{da_t}.$$

For a fixed punishment threat, the ratio of the marginal benefits of effort induced by Brownian information versus bad news Poisson information is infinite in the limit. Thus, intuitively, under the improved monitoring technology P will punish A if the Brownian component is *Bad* no matter what is X_t since P cares about maximizing effort incentives. In fact, a simple application of the product rule shows that for this particular monitoring technology a vector is *Bad* if and only if at least one component signal is *Bad*.

But now we have a problem. Unfortunately Brownian information, despite its very strong incentive power, has even weaker statistical power with the Brownian *Bad* signal occurring about half the time no matter the effort A exerts. By punishing A whenever the Brownian component is *Bad*, P ensures that the extreme weakness of Brownian statistical power infects the statistical power of the information generated by the improved monitoring technology. Consequently, despite the considerably greater incentive power of the improved monitoring technology, the typical *Bad* vector has much lower quality than the original *Bad* signal. In fact, it is not hard to show that the negative effort-elasticity of *Bad* vectors is on the order of $\sqrt{\Delta}$ which goes to zero as Δ tends to zero. Thus, when bad news Poisson monitoring is improved by including a Brownian component, the optimal contract collapses into a trivial arrangement that always pays A a flat wage w_{salary} and never terminates A. A best responds by putting in zero effort, and P despite her better monitoring becomes worse off.

4.2 Incentive Power and Statistical Power

In the example above I showed that including a Brownian component with bad news Poisson monitoring leads to a worse outcome. I linked this result specifically to the fact that Brownian information has much greater incentive power but much weaker statistical power than bad news Poisson information.

I now generalize this result to the universe of binary-valued monitoring technologies in the continuous time infinite horizon limit: $\lim_{\Delta\to 0}\lim_{T\to\infty}$. Starting with a technology X_1, I show that introducing another technology X_2 causes the optimal contract to collapse into a trivial arrangement if, relative to X_1, X_2 has sufficiently strong incentive power but sufficiently weak statistical power. A new fact that emerges from the more general analysis: The condition that X_2 must have sufficiently strong incentive power relative to X_1 does not mean X_2 must have stronger incentive power than X_1 as in the bad news Poisson-Brownian example. In general, X_2's incentive power just needs to be above a certain threshold that is increasing in but could be strictly smaller than the incentive power of X_1. See Theorem 3 below.

I begin by restricting attention to all monitoring technologies in the continuous time infinite horizon limit satisfying the following regularity conditions: For all a_t,

$$\lim_{\Delta\to 0} -\frac{d}{da_t}\mathbf{P}(X_t \in \mathit{Bad} \mid a_t) = \Theta(\Delta^{\alpha}) \text{ for some } \alpha \geq 0$$
$$\lim_{\Delta\to 0} \mathbf{P}(X_t \in \mathit{Bad} \mid a_t) = \Theta(\Delta^{\gamma^b}) \text{ for some } \gamma^b \geq 0.$$
$$\lim_{\Delta\to 0} \mathbf{P}(X_t \in \mathit{Good} \mid a_t > 0) = \Theta(\Delta^{\gamma^g}) \text{ for some } \gamma^g \geq 0.$$

This class of monitoring technologies includes bad news Poisson monitoring, Brownian monitoring, good news Poisson monitoring:

$$X_t = \begin{cases} \text{event} & \text{with probability} \quad a_t\lambda\Delta \\ \text{no event} & \text{with probability} \quad 1 - a_t\lambda\Delta \end{cases}$$

as well as vector combinations of these technologies.

Here, α measures the incentive power of information – the lower is α the greater is the incentive power. γ^b measure the statistical power of information – the higher is γ^b the greater is the statistical power. Thus, $\alpha - \gamma^b$ measures the quality of *Bad* signals – the lower is $\alpha - \gamma^b$ the greater is the negative effort-elasticity and, therefore, the quality of *Bad* signals. It is always the case that $\alpha \geq \gamma^b$.

Theorem 2. *Given a monitoring technology, whether or not the optimal contract can induce positive effort is largely determined by the quality of Bad signals.*

Formally, assume $\alpha \leq 1$. If $\alpha - \gamma^b = 0$ then depending on how the rest of the model is parameterized the effort induced by the optimal contract can be made to be arbitrarily high (i.e. close to 1). Otherwise the optimal contract induces zero effort.

Proof. See appendix. □

Theorem 2 implies that in the continuous time infinite horizon limit if the negative effort-elasticity of *Bad* signals vanishes then the optimal contract collapses. My investigation of the better monitoring/worse outcome channel will be built around finding improvements to monitoring that cause this elasticity to vanish.

Corollary 1. *If X_t is Brownian or good news Poisson, the optimal contract induces zero effort. If X_t is bad news Poisson, there are parameterizations of the model under which the optimal contract induces nonzero effort.*

Corollary 1 matches classic results from the literature on repeated games with imperfect public monitoring. For example, Abreu, Milgrom, and Pearce (1991) shows that in a continuous time repeated prisoner's dilemma game with public monitoring cooperation can be supported as an equilibrium if monitoring is bad news Poisson but not good news Poisson. Sannikov and Skrzypacz (2007) shows that in a continuous time repeated Cournot oligopoly game with public monitoring collusion cannot be supported if monitoring is Brownian. This common baseline allows me to better highlight how my results on the relationship between monitoring and surplus, with their emphasis on incentive and statistical power, differ from related results in the repeated games literature. In particular, whereas better monitoring can lead to a worse outcome in my setting, improvements to the information content of signals at each date in the models described above always weakly improve the scope for cooperation.[14]

Armed with Theorem 2, I can now investigate how improvements to the monitoring technology affect optimality. I begin with a binary-valued monitoring technology $X_1 \in \{b_1, g_1\}$ with exponents $(\alpha_1, \gamma_1^b, \gamma_1^g)$. I then improve it by including a conditionally independent binary valued monitoring technology $X_2 \in \{b_2, g_2\}$ with exponents $(\alpha_2, \gamma_2^b, \gamma_2^g)$. I show that it is generically the case that effort has a strictly monotone effect on the vector valued information (X_1, X_2) generated by the improved monitoring technology. Thus, (X_1, X_2) also has some associated exponents $(\alpha, \gamma^b, \gamma^g)$. I derive the formulas for α, γ^b, and γ^g as a function of $(\alpha_1, \gamma_1^b, \gamma_1^g)$ and $(\alpha_2, \gamma_2^b, \gamma_2^g)$. Then, by inverting the formulas and using Theorem 2, I can show, given $(\alpha_1, \gamma_1^b, \gamma_1^g)$, what kinds of improvements $(\alpha_2, \gamma_2^b, \gamma_2^g)$ cause the optimal contract to collapse.

At each date t, (X_{1t}, X_{2t}) can take one of four values: $(g_1, g_2), (g_1, b_2), (b_1, g_2)$, or (b_1, b_2). Holding Δ fixed, $\mathbf{P}((X_{1t}, X_{2t}) = (g_1, g_2) \mid a_t, \Delta)$ is strictly increasing in a_t and $\mathbf{P}((X_{1t}, X_{2t}) = (b_1, b_2) \mid a_t, \Delta)$ is strictly decreasing in a_t. The probability that $(X_{1t}, X_{2t}) = (g_1, b_2)$ is $\mathbf{P}(X_{1t} = g_1 \mid a_t, \Delta) \cdot \mathbf{P}(X_{2t} = b_2 \mid a_t, \Delta)$. By the product rule, as $\Delta \to 0$, the derivative of $\mathbf{P}((X_{1t}, X_{2t}) = (g_1, b_2) \mid a_t, \Delta)$ with respect to a_t is $A(\Delta) - B(\Delta)$ where $A(\Delta) = \Theta(\Delta^{\alpha_1 + \gamma_2^b})$ and $B(\Delta) = \Theta(\Delta^{\gamma_1^g + \alpha_2})$. A sufficient condition for $\mathbf{P}((X_{1t}, X_{2t}) = (g_1, b_2) \mid a_t, \Delta)$ to be a strictly monotonic function of a_t in the continuous-time limit is $\alpha_1 + \gamma_2^b \neq \gamma_1^g + \alpha_2$. Similarly, a sufficient condition for $\mathbf{P}((X_{1t}, X_{2t}) = (b_1, g_2) \mid a_t, \Delta)$ to be a strictly monotonic function of a_t in the continuous-time limit is $\alpha_1 + \gamma_2^g \neq \gamma_1^b + \alpha_2$. Thus,

Lemma 1. *If $\alpha_1 - \alpha_2 \neq \gamma_1^g - \gamma_2^b$ or $\gamma_1^b - \gamma_2^g$ then effort has a strictly monotone effect on (X_{1t}, X_{2t}) as $\Delta \to 0$.*

Lemma 2. *Given X_{1t} and X_{2t} with exponents $(\alpha_1, \gamma_1^b, \gamma_1^g)$ and $(\alpha_2, \gamma_2^b, \gamma_2^g)$, if $\alpha_1 \geq \alpha_2$ then the exponents (a, γ^b, γ^g) associated with the vector-valued (X_{1t}, X_{2t}) are*

$$(\alpha = \alpha_2, \gamma^b = \min\{\gamma_1^b, \gamma_2^b\}, \gamma^g = \gamma_2^g) \quad \text{if } \gamma_1^g - \gamma_2^b < \alpha_1 - \alpha_2 < \gamma_1^b - \gamma_2^g$$
$$(\alpha = \alpha_2, \gamma^b = \gamma_2^b, \gamma^g = \min\{\gamma_1^g, \gamma_2^g\}) \quad \text{if } \gamma_1^b - \gamma_2^g < \alpha_1 - \alpha_2 < \gamma_1^g - \gamma_2^b$$
$$(\alpha = \alpha_2, \gamma^b = \gamma_2^b, \gamma^g = \gamma_2^g) \quad \text{if } \gamma_1^g - \gamma_2^b,\ \gamma_1^b - \gamma_2^g < \alpha_1 - \alpha_2$$

[14] In a repeated games setting with public monitoring, Kandori (1992) shows that making monitoring more informative in the sense of Blackwell (1950) causes the pure-strategy sequential equilibrium payoff set to expand in the sense of set inclusion.

Lemma 2 only considers $\alpha_1 \geq \alpha_2$. The other case, $\alpha_2 \geq \alpha_1$, is implied by symmetry.

Proof. See appendix. □

Lemma 2 yields an explicit characterization of counterproductive improvements to the monitoring system.

Theorem 3. *Improving monitoring by introducing new information that is, relative to the original information, sufficiently strong in incentive power but sufficiently weak in statistical power causes the optimal contract to collapse.*

Formally, suppose $\alpha_1 = \gamma_1^b$*. If* $\alpha_2 < \alpha_1 + \gamma_2^b$ *and* $\gamma_2^b < \min\{\gamma_1^b, \alpha_2\}$*, then* $\alpha > \gamma^b$*. The result is tight in the sense that if either of the inequalities is reversed then* $\alpha = \gamma^b$*.*

The inequality $\alpha_2 < \alpha_1 + \gamma_2^b$ of Theorem 3 is the formal expression of what it means for incentive power to be sufficiently strong. Notice if γ_2^b is positive then Y does not need to have stronger incentive power than X to cause the optimal contract to collapse. Similarly, sufficiently weak in statistical power means $\gamma_2^b < \min\{\gamma_1^b, \alpha_2\}$. This latter inequality admits a natural interpretation: The statistical power of the new information needs to be weak enough so that its *Bad* signal is low quality and much more common than the *Bad* signal of the original information. As a corollary, the better monitoring/worse outcome result described in the previous subsection is recovered.

Corollary 2. *Improving a bad news Poisson monitoring technology by including a conditionally independent Brownian signal of effort causes the optimal contract to collapse.*

Proof. Let X_1 denote bad news Poisson monitoring and X_2 denote Brownian monitoring. Notice, $\alpha_2 = 0.5 < 1 + 0 = \alpha_1 + \gamma_2^b$ and $\gamma_2^b = 0 < \min\{1, 0.5\} = \min\{\gamma_1^b, \alpha_2\}$. The corollary now follows from Theorem 3. □

4.3 Noisy Information Does Not Mean Weak Incentives

When it comes to the provision of incentives there is a common misconception that if information is very noisy then incentives cannot be very strong. This confounding of "informativeness" and incentives likely arises due to the fact that one often works with parametric families of information structures within which the rankings based on incentive power and any reasonable measure of informativeness are coincident.

For example, imagine a moral hazard model where an agent can either exert effort $a = 1$ or shirk $a = 0$. a determines the mean of a normally distributed payoff X with variance σ. In this setting incentive power can be measured by how small of a punishment threat one can use and still induce effort. The set of payoffs whose likelihood of occurring decreases when the agent exerts effort is $\{X \leq 0.5\}$. Thus, statistical power can be measured by $\mathbf{P}(X \leq 0.5 | a = 1)$ and the quality of *Bad* signals can be measured by the likelihood ratio $\mathbf{P}(X \leq 0.5|a = 0)/\mathbf{P}(X \leq 0.5|a = 1)$. As σ decreases, incentive power, statistical power, and the quality of *Bad* signals all go up. These improvements can all be traced to the fact that Blackwell informativeness goes up.

Greater Blackwell informativeness always means greater incentive power. However Blackwell informativeness is not a complete order and, depending on the situation, there may be other reasonable ways to compare informativeness across information structures that either do not agree with Blackwell or apply when Blackwell does not. For example, in the setting of this paper where, in equilibrium the effort is known but nevertheless *Bad* signals are punished as if they indicate a downward deviation, the quality of *Bad* signals as measured by negative effort-elasticity is a natural way to compare informativeness. And it is hard to argue a bad news Poisson increment that generates a high quality *Bad* signal is not more informative than a Brownian increment that generates a very low quality *Bad* signal even though the two increments cannot be compared via Blackwell. In this case, the parametric intuition gets in the way, leading us to be perhaps surprised that Brownian information still strongly dominates bad news Poisson information in terms of incentive power.

The exploration of the better monitoring/worse outcome channel in this section helps clarify the distinction between measuring informativeness and measuring incentive power. Theorem 3 and Corollary 2 highlight how information that is at once not very informative but still very sensitive to effort exists, is not unusual, and can have an outsized effect on optimal contracting.

That outsized effect is of course negative and it points to a possible danger as companies increasingly incorporate monitoring technologies that can generate vast amounts of worker data. Much of this raw data is quite noisy and if this noisy data also happens to be sensitive to effort, then the better monitoring/worse outcome channel implies incentives can be compromised.

One possible way to preserve incentives is to have contracts directly condition worker outcomes (e.g. bonus pay) on the raw data, thereby eliminating the managerial discretion upon which better monitoring/worse outcome rests. There are, however, a number of issues that make this difficult to implement in practice. The data may reveal information that the company is reluctant to make public. The worker may also be concerned about maintaining his privacy. In industries that are constantly in flux, it may not be clear ex-ante how to optimally condition contracts on the raw data. And lastly, as long as the manager still makes reports that affect worker outcomes – perhaps there is another source of non-contractible information that needs reporting on – there is no way to prevent the manager from conditioning her reports (that are supposed to be about that other source of information) on the contractible data generated by the monitoring technology, in which case, the better monitoring/worse outcome channel remains open.

A more plausible way to overcome better monitoring/worse outcome may be to have the technology itself censor the raw data. From this point of view an important aspect of the continued development of monitoring technologies is the development of algorithms that can process the noisy input and turn it into information that represent better quality signals of performance. As a toy example, imagine the input is the Brownian random walk considered earlier. One easy way to increase the quality of *Bad* signals through data censoring is to let P observe a *Bad* signal only when the random walk has gone down multiple times, on the order of $\frac{1}{\Delta}$.

5 Periodic Performance Appraisal

The better monitoring/worse outcome channel explored in the previous section implies that limiting the information content of monitoring can be beneficial. How should information content be limited? One obvious way, if applicable, is to remove a signal that is strong in incentive power but weak in statistical power. Another possibility that was discussed is to have a monitoring technology self-censor. In this section I explore a third way to beneficially limit the information content of monitoring by limiting when monitoring is allowed to occur.

I consider a canonical setting in which there is a stochastic process tracking *cumulative* productivity and P monitors by *sampling* that process. I show how reducing the frequency of sampling can significantly improve surplus and productivity. When the frequency of sampling is reduced, so is the information content of monitoring. P does not observe the kind of detailed day-to-day performance data she would observe if she sampled frequently. Instead, all she observes are signals of how A has performed overall since the previous sampling date a while ago. Thus the fact that infrequent sampling can be beneficial means that *appraising a worker's overall performance every once in a while can dominate closely monitoring his day-to-day performance every single day.*

5.1 Information Content or Observation Frequency?

Reducing the frequency of sampling actually changes two things about monitoring at once: It reduces the information content of monitoring but it also reduces the frequency at which P observes new information. In the discussion above, I suggest that it is the first effect that is the source of potential benefits from monitoring infrequently: Reducing observation frequency is beneficial only because it reduces information content. But a seminal result of Abreu, Milgrom, and Pearce (1991) concerns how the second effect *by itself* can lead to greater efficiency in repeated games with imperfect monitoring. In that paper the second effect is isolated by batching the release of information so that, say, every 10 units of time, P observes the information generated from the past 10 units of time all at once. Batching allows the information content of monitoring to stay exactly the same while still reducing the frequency of observation.

I now show that batching is counterproductive in my setting where equilibrium play is determined by the prospect of credible threats. Moreover, as long as the optimal contract induces nonzero effort, batching is strictly counterproductive.

Releasing information in batches every once in a while is equivalent to releasing information as it is generated but restricting the players to respond to the flow of information only every once in a while. Unlike the repeated games literature where the game is taken as given, P and A in my model are doing optimal contracting and can choose the structure of the contract game. In particular, they can choose to use a contract game that only allows P to react to the flow of information every once in a while: For example, the contract game could be structured so that the message space is a singleton between t_1 and $t_2 - \Delta$. In this case, P does not have a rich enough message space to react to new information between t_1 and $t_2 - \Delta$ and it is equivalent to batching the information generated between t_1 and t_2 and releasing it all at once at date t_2. Since Theorem 2 is a result about optimal contracting, contract games that allow P to react to new information only every once in a while are already folded into the analysis. Thus, my optimality result indirectly implies that batching cannot increase surplus.

In fact, choosing a contract game that allows P to react to new information only every once in a while is not only not helpful, it is usually strictly hurtful. Suppose the contract game does not allow P to react to new information between t_1 and $t_2 - \Delta$. On date t_2 when P finally has the opportunity to affect A's continuation payoff through her reports, all of A's efforts before date t_2 have been sunk. The credible threats refinement implies that P's goal standing at the beginning of date t_2 is to choose a date t_2 report strategy that maximizes date t_2 effort incentives. This means P will report in a way so that A's date $t_2 + \Delta$ continuation payoff is maximized (minimized) depending only on if $X_{t_2} \in \mathit{Good}$ ($\in \mathit{Bad}$). In particular, P ignores all signals generated before date t_2. Anticipating this, A best responds by exerting zero effort from t_1 to $t_2 - \Delta$.

5.2 Optimal Sampling Frequency: Efficiency versus Capacity

To explore the costs and benefits of infrequent sampling, I now consider a canonical setting where effort controls the drift of a Brownian motion and P monitors by sampling the Brownian motion. Such a model can be adapted from the original model as follows:

The timing is the same as before: $t = 0, \Delta, \dots T$. The sequence of events within each date t is the same as before except P may or may not monitor A. If P does monitor A at date t, then she samples the current value Z_t of the Brownian process:

$$Z_t = \sum_{i=0}^{t} a_i \Delta + B_t$$

where B_t is standard Brownian motion. In addition, P reports a public message $m_t \in \mathcal{M}_t$ and a public randomizing device is realized just like before. The model then moves to the next date $t + \Delta$. If P does not monitor, then the model immediately moves to $t + \Delta$ after A exerts date t effort.

A contract game in addition to specifying $\mathcal{M}$, w, and τ also specifies a predictable sequence of *sampling times* $e_1 < e_2 < \ldots < T$. Continuation payoffs can be defined exactly as before at every date, including non-sampling dates. The credible threats refinement admits a natural generalization. Contracts and the optimal contracting problem are defined exactly as before. I am interested in studying the properties of the optimal contract in the continuous time infinite horizon limit.

Lemma 3. *The optimal contract features a deterministic sequence of sampling times. There exist a D^* such that as $\Delta \to 0$ and $T \to \infty$, the optimal contract's sequence of sampling times converges to the sequence $\{D^*, 2D^*, 3D^* \ldots\}$.*

Given Lemma 3 characterizing the optimal contract in the continuous time infinite horizon limit can be broken down into two steps: First, characterize the optimal contract in the continuous time infinite horizon limit given a fixed *sampling frequency* $\frac{1}{D}$. Then, find the optimal sampling frequency $\frac{1}{D^*}$.

Theorem 4. *Fix a sampling frequency $\frac{1}{D}$. There exist $\rho^*(D)$ and $p^*(D)$ such that the optimal contract given sampling frequency $\frac{1}{D}$ has the following structure:*

- *For each $k \in \mathbb{Z}^+$, $\mathcal{M}_{kD} = \{pass, fail\}$.*
- *$m_{kD} = fail$ iff $Z_{kD} - Z_{(k-1)D} \leq \rho^*(D)$.*
- *w consists of a pair of constants $w_{salary}(D)dt, w_{severance}(D)$.*
- *If $m_{kD} = pass$ then A is retained for the sampling period $(kD, (k+1)D]$ and is paid a stream $w_{salary}(D)dt$.*
- *If $m_{kD} = fail$ then A is terminated with probability $p^*(D)$.*
  - *If A is not terminated then it is as if P reported pass.*
  - *If A is terminated then he is paid a lump sum $w_{severance}(D)$.*

The credible threats refinement implies that at each date kD the report depends only on $Z_{kD} - Z_{(k-1)D}$. $Z_{kD} - Z_{(k-1)D}$ is not monotone with respect to effort but it does satisfy MLRP with respect to effort. In the proof I show that as a consequence of MLRP P's effort maximizing report strategy involves setting a threshold $\rho^*(D)$ and endogenously splitting the range of $Z_{kD} - Z_{(k-1)D}$ into the "*Bad*" signals below $\rho^*(D)$ and the "*Good*" signals above $\rho^*(D)$. Once this result is established, the rest of the proof closely follows that of Theorem 1. The resulting continuous time optimal contract can be viewed as equivalent to the original discrete time optimal contract except the exogenously fixed date length Δ is now replaced with a to-be-endogenously-determined optimal sampling period length D^*:

Theorem 5. *There exists an optimal sampling frequency* $\frac{1}{D^*} \in (0, \infty)$*. As sampling frequency converges to infinity or zero effort induced by the optimal contract given the sampling frequency decreases to zero.*

Given that P's report strategy identifies endogenous *Good* and *Bad* signals, the statistical power of $Z_{kD} - Z_{(k-1)D}$ can be defined similar to before based on the probability of *Bad* signals occurring. This then allows me to use the negative effort-elasticity of *Bad* signals to measure the quality of *Bad* signals. As sampling frequency increases, the quality of *Bad* signals worsens. Theorem 2 can be easily adapted to then show that in the continuous sampling limit the optimal contract must induce zero effort. As sampling frequency decreases, the quality of *Bad* signals continues to improve. If there were no discounting, P's flow payoff would increase to the first-best level as sampling frequency decreases to zero. However, since there is discounting, as the length of a sampling period increases discounting begins eroding the capacity to provide incentives: In the beginning of a long sampling period, the threat of termination – which is bounded by the size of the continuation surplus – in the distant future when the period concludes has little effect on the continuation payoff of A today. Thus, as sampling frequency goes to zero, the optimal contract again induces zero effort. Since the quality of *Bad* signals determines how efficient it is to provide incentives, the optimal sampling frequency can be viewed as balancing the efficiency versus the capacity of incentive provision.

One concern about maintaining an infrequent sampling of performance is whether or not the act of sampling itself is verifiable. My model assumes implicitly that the principal cannot sample outside of the specified sampling dates. If the principal could sample outside of the specified sampling dates then he would want to sample constantly. Anticipating this deviation, the contracting parties would then optimally agree to a trivial contract. Thus, my analysis suggests that it is important to force the principal to go through formal channels in order to appraise worker performance. Another interpretation is that my analysis suggests a benefit to making appraising performance sufficiently costly to P. This way while it can still be verified if P has conducted a formal performance appraisal when the contract calls for it, there is less fear that P will conduct an unwarranted informal performance appraisal at some other date.

5.3 Development versus Accountability

"Performance management's purpose is shifting, structurally and dramatically. With blurring lines between performance management and talent development, executives will have to consider how to balance the assessment of past performance with the ongoing need to develop employee skills. The professional development function — emphasizing performance improvement, coaching, and feedback — often receives short shrift." (Schrage et al, 2019)

"The tension between the traditional and newer approaches stems from a long-running dispute about managing people: Do you "get what you get" when you hire your employees? Should you focus mainly on motivating the strong ones with money and getting rid of the weak ones? Or are employees malleable? Can you change the way they perform through effective coaching and management and intrinsic rewards such as personal growth and a sense of progress on the job? With traditional appraisals, the pendulum had swung too far toward the former." (Cappelli and Tavis, 2016)

"Ultimately, we need to accomplish three things: review contributions, reward accomplishments, and give and receive feedback. Do they need to be conflated into a cumbersome process? Under our new system, our contributors get highly specific performance feedback at least once every six weeks. But in practice it happens every week. Instead of lagging, the performance management process is leading. The true mechanisms for success are the ones that build capabilities." (Donna Morris, Chief Human Resources Officer at Adobe, an early adopter of continuous performance management practices. See Doerr, 2018.)

The rise of continuous performance management (CPM) practices has been driven by the need for greater worker development in an increasingly dynamic world where, to quote Schrage et al (2019), "the half-life of skills is going to get shorter." Making performance appraisals more frequent obviously makes it easier for managers to help workers continually identify and develop new skills needed to adapt to a changing environment. What is less obvious is how such a change might impact incentive provision. Cappelli and Tavis (2016) notes that "companies changing their systems are trying to figure out how their new practices will affect the pay-for-performance model." Steffen Maier, cofounder of Impraise – a company that helps clients implement CPM – has pointed to a general apprehension about how these new practices, which often involve nuanced, ratingless feedback, might lead to less transparency and more bias in compensation decisions (Caprino, 2016).

What the better monitoring/worse outcome channel does is clarify an explicit downside for incentive provision when performance appraisals become more frequent and tie it specifically to the principal's making discretionary reports about agent performance. At the same time, my exploration of the better monitoring/worse outcome channel also yields strategies for potentially overcoming that downside.

Since the discretionary nature of principal reports is a key reason why better monitoring can lead to a worse outcome, one strategy for preventing CPM from negatively affecting incentive provision is to make the feedback less discretionary. Companies have attempted to do this through eliciting feedback from multiples sources (e.g. calibration meetings between managers, peer-based feedback) and by making feedback be about team performance which can be measured more objectively than individual performance (Caprino, 2016). Ledford, Benson, and Lawler (2016), in their study of ongoing, ratingless, and crowdsourced feedback, finds that "the most effective patterns ... for meeting rewards system objectives are the use of all three practices together and/or ongoing feedback plus crowdsourced feedback."

Another potential strategy for overcoming the downside of frequent performance appraisals is to complement those frequent appraisals with a separate infrequent appraisal of performance reserved for incentive provision. While some companies have eliminated annual performance reviews entirely, most practitioners of CPM still keep one annual review specifically for making compensation decisions (Caprino, 2016). For example, *Patagonia* still maintains a formal review that provides "an annual reckoning" to help determine compensation and bonuses (Ramirez, 2018).

One concern about such a hybrid arrangement is whether or not the two modes of performance appraisals interfere with each other. As Cappelli and Tavis (2016) point out, "It will be interesting to see whether most supervisors end up reviewing the feedback they've given each employee over the year before determining merit increases. And could that subtly undermine development by shifting managers' focus back to accountability?" My analysis of batching in Section 5.1 highlights a potential pitfall of trying to hold workers accountable periodically while monitoring continuously. If the periodic appraisal of past performance that determines worker incentives is simply a review of past continuous appraisals then introducing the periodic appraisal is counterproductive.

Instead, continuous and periodic appraisal should be focused on gathering different types of information to serve the competing demands of development versus accountability. The idea of splitting appraisals into ones focused on development and on accountability has been around for decades. "After running a well-publicized experiment in 1964, General Electric concluded it was best to split the appraisal process into separate discussions about accountability and development, given the conflicts between them. Other companies followed suit." (Cappelli and Tavis, 2016). The better monitoring/worse outcome channel provides guidance for how to successfully implement such a split through continuous and periodic performance appraisals. The development oriented continuous appraisal could focus on corrective feedback by gathering information about what actions the worker should and should not take rather than how much effort workers are putting into a certain action. The accountability oriented periodic appraisal could then focus on gathering information about how successful were the worker's actions overall across the entirety of the appraisal period.

Of course the boundary between development-relevant and accountability-relevant information may not be clear cut. Indeed, figuring out which new actions need to be taken probably involves monitoring the success of previous actions. As companies continue to rethink their monitoring systems to better serve worker development, an important open question is how to do this in a way that minimizes any negative impact on holding workers accountable. My work contributes to a better understanding of this fundamental problem of monitoring design by highlighting a potentially significant cost to incentive provision that the use of CPM can entail.

6 Conclusion

In this paper I considered a principal-agent model with moral hazard and highlighted a better monitoring/worse outcome channel when information generated by monitoring is non-contractible and must pass through the principal's discretionary reports in order to affect the agent's payoff. Allowing the principal to observe additional information that is, relative to the original information, sufficiently strong in incentive power but sufficiently weak in statistical power can lead to an optimal contract that induces lower effort and generates less surplus. Such improvements to monitoring tempt the principal to make too many negative reports leading to overpunishment of the agent. The agent then responds by refusing to accept contracts with strong punishment threats. Without a strong punishment threat, incentive provision is hindered, leading to lower effort and less surplus. The better monitoring/worse outcome result implies that limiting the information content of monitoring can be beneficial. I explored a number of ways to do this including data-censoring and reducing the frequency of monitoring.

7 Appendix

Proof of Theorem 1. Based on Zhu (2018), I take as given that for any contract game $(\mathcal{M}, w, \tau)$ the credible threats refinement selects the set of sequential equilibria $\mathcal{E}^*(\mathcal{M}, w, \tau)$ – the credible threats equilibria – constructed as follows:

Let ξ_t denote the public randomizing device realized at the end of date t. For every public history of the form $h_{T-\Delta}$, define $(W_T^*(h_{T-\Delta}), V_T^*(h_{T-\Delta})) = (w_T(h_{T-\Delta}), -w_T(h_{T-\Delta}))$. Fix a $t < T$ and suppose by backwards induction a unique continuation payoff process $(W_{s+\Delta}^*(h_s), V_{s+\Delta}^*(h_s))$ has been constructed for all h_s where $t \leq s < T$. Given a public history $h_{t-\Delta}$, if $\tau(h_{t-\Delta}) = t$ then define $(W_t^*(h_{t-\Delta}), V_t^*(h_{t-\Delta})) = (w_t(h_{t-\Delta}), -w_t(h_{t-\Delta}))$. Otherwise, define

$$\mathcal{M}_t^*(h_{t-\Delta}) = \underset{m_t \in \mathcal{M}_t(h_{t-\Delta})}{\arg\max} \mathbf{E}_{\xi_t} V_{t+\Delta}^*(h_{t-\Delta} m_t \xi_t)$$

$$V_{t+\Delta}^{max}(h_{t-\Delta}) = \max_{m_t \in \mathcal{M}_t(h_{t-\Delta})} \mathbf{E}_{\xi_t} V_{t+\Delta}^*(h_{t-\Delta} m_t \xi_t)$$

$$W_{t+\Delta}^{pass}(h_{t-\Delta}) = \max_{m_t \in \mathcal{M}_t^*(h_{t-\Delta})} \mathbf{E}_{\xi_t} W_{t+\Delta}^*(h_{t-\Delta} m_t \xi_t)$$

$$W_{t+\Delta}^{fail}(h_{t-\Delta}) = \min_{m_t \in \mathcal{M}_t^*(h_{t-\Delta})} \mathbf{E}_{\xi_t} W_{t+\Delta}^*(h_{t-\Delta} m_t \xi_t)$$

$$a_t^*(h_{t-\Delta}) = \max\left\{\underset{a_t\in[0,1)}{\arg\max} -h(a_t)\Delta - e^{-r\Delta}\mathbf{P}(X_t \in Bad \mid a_t)(W_{t+\Delta}^{pass}(h_{t-\Delta}) - W_{t+\Delta}^{fail}(h_{t-\Delta}))\right\}$$

$$\begin{aligned}W_t^*(h_{t-\Delta}) &= (w_t(h_{t-\Delta}) - h(a_t^*(h_{t-\Delta})))\Delta \\ &\quad + e^{-r\Delta}\left[W_{t+\Delta}^{pass}(h_{t-\Delta}) - \mathbf{P}(X_t \in Bad \mid a_t^*(h_{t-\Delta}))(W_{t+\Delta}^{pass}(h_{t-\Delta}) - W_{t+\Delta}^{fail}(h_{t-\Delta}))\right]\end{aligned}$$

$$V_t^*(h_{t-\Delta}) = (u(a_t^*(h_{t-\Delta})) - w_t(h_{t-\Delta}))\Delta + e^{-r\Delta}V_{t+\Delta}^{max}(h_{t-\Delta})$$

$a_t^*(h_{t-\Delta})$ is well-defined because the expression inside the arg max is continuous and has a maximum. By induction, I have now constructed a unique continuation payoff process $(W_t^*(h_{t-\Delta}), V_t^*(h_{t-\Delta}))$ for all t and $h_{t-\Delta}$. I refer to this process as the continuation payoff process of $(\mathcal{M}, w, \tau)$ and (W_0^*, V_0^*) as the ex-ante payoff of $(\mathcal{M}, w, \tau)$

Fix any strategy profile (a, m) satisfying

$$\begin{aligned}a_t(H_{t-\Delta}^A, h_{t-\Delta}) &= a_t^*(h_{t-\Delta}) \\ m_t(H_t^P, h_{t-\Delta}) &\in \mathcal{M}_t^*(h_{t-\Delta}) \\ \mathbf{E}_{\xi_t}W_{t+\Delta}^*(h_{t-\Delta}m_t(H_t^P, h_{t-\Delta})\xi_t) &= W_{t+\Delta}^{pass}(h_{t-\Delta}) \text{ if } X_t \in Good \\ \mathbf{E}_{\xi_t}W_{t+\Delta}^*(h_{t-\Delta}m_t(H_t^P, h_{t-\Delta})\xi_t) &= W_{t+\Delta}^{fail}(h_{t-\Delta}) \text{ if } X_t \in Bad\end{aligned}$$

for all t and $h_{t-\Delta}$. There exists an assessment with such a strategy profile that is a sequential equilibrium. All such sequential equilibria generate the continuation payoff process of $(\mathcal{M}, w, \tau)$. $\mathcal{E}^*(\mathcal{M}, w, \tau)$ is defined to be this set of sequential equilibria. An element of this set is called a credible-threats equilibrium.

I claim given a contract game $(\mathcal{M}, w, \tau)$, there exists another contract game with the same ex-ante payoff and with the property that the message space is $\{pass, fail\}$ at all times. To construct this other contract game, first note that there exists an element $(a, m) \in \mathcal{E}^*(\mathcal{M}, w, \tau)$ with the property that for all $h_{t-\Delta}$, $H_{t-\Delta}^P$, and $x_1, x_2 \in Im(X_t)$, $m_t(H_{t-\Delta}^P x_1, h_{t-\Delta}) = m_t(H_{t-\Delta}^P x_2, h_{t-\Delta})$ if $x_1, x_2 \in Good$ or $x_1, x_2 \in Bad$. By definition, (a, m) generates ex-ante payoff (W_0^*, V_0^*). Now remove all the edges and vertices of $(\mathcal{M}, w, \tau)$ that are not reached by (a, m). Call the resulting game $(\mathcal{M}'', w'', \tau'')$. (a, m) also $\in \mathcal{E}^*(\mathcal{M}'', w'', \tau'')$ and clearly still generates ex-ante payoff (W_0^*, V_0^*). $(\mathcal{M}'', w'', \tau'')$ has the property that the message space has at most two elements for each t and $h_{t-\Delta}$. Pick any $h_{t-\Delta}$ in $(\mathcal{M}'', w'', \tau'')$. If $|\mathcal{M}_t''(h_{t-\Delta})| = 2$ then by definition of (a, m) one of the messages is the one that is always reported under (a, m) if $X_t \in Good$ and the other is the message that is always reported under (a, m) if $X_t \in Bad$. Relabel the *Good* message *pass* and the *Bad* message $fail$. If $|\mathcal{M}_t''(h_{t-\Delta})| = 1$ then split the message into two messages, *pass* and $fail$, and attach copies of the continuation game following the original message after both the *pass* and the $fail$ messages. This altered game, call it $(\mathcal{M}', w', \tau')$, has the property that the message space is $\{pass, fail\}$ after all histories. Moreover, it is clear the ex-ante payoff of $(\mathcal{M}', w', \tau')$ is (W_0^*, V_0^*). Thus, the claim is proved.

From now on, I assume without loss of generality that in a contract game the message space at all times is $\{pass, fail\}$ and there is a credible threats equilibrium with the property that $m_t = pass$ iff $X_t \in Good$.

Suppose by induction there exists an optimal contract game in the T-model with the following properties: For all $t < T$, there exist numbers $w_{salary,t}\Delta, w_{severance,t+\Delta}, p^*_{t+\Delta}$ such that

- If $m_t = pass$ then A is retained for date $t + \Delta$ and paid $w_{salary,t}\Delta$.
- If $m_t = fail$ then A is terminated at date $t + \Delta$ with probability $p^*_{t+\Delta}$.
  - If A is not terminated then it is as if P reported $pass$.
  - If A is terminated then he is paid $w_{severance,t+\Delta}$.
- $W^*_{t+\Delta}(h_{t-\Delta} pass \xi_t) = 0$ for all $h_{t-\Delta}$ and ξ_t, and $W_0^* = 0$.

The existence of such an optimal contract game is trivially true when $T = 0$. For general T, the ex-ante payoff of this contract game is $(W_0^* = 0, V_0^* = S_T^*)$ where S_T^* is the Pareto-optimal surplus in the T-model.

Now fix an optimal contract game $(\mathcal{M}, w, \tau)$ in the $T + \Delta$-model. Consider the subgame $(\mathcal{M}, w, \tau)|pass\xi_0$ following the date 0 $pass$ report and a realization of ξ_0. It can be identified with a contract game in the T-model. Replace every $(\mathcal{M}, w, \tau)|pass\xi_0$ with the same optimal contract game $\mathscr{C}$ in the T-model with the properties described above. Next, change the portion of $(\mathcal{M}, w, \tau)$ following the date 0 $fail$ report to a randomization between $\mathscr{C}$ and termination with severance pay $w_{severance,\Delta} = -S_T^*$ where the randomization is based on ξ_0 and is structured so that A's expected date Δ payoff equals $-(W_\Delta^{pass} - W_\Delta^{fail})$. Finally shift w_0 so that A's ex-ante payoff remains 0. The modified contract game's continuation payoff process continues to satisfy ex-ante and ex-interim participation constraints and delivers a weakly larger ex-ante payoff to P compared to $(\mathcal{M}, w, \tau)$. Thus, it is also an optimal contract game. Moreover, it has the same structure that I am trying to prove by induction. This completes the induction.

Fix an optimal contract game $(\mathcal{M}, w, \tau)$ in the $T + \Delta$-model with the structure described above and consider the subgame $(\mathcal{M}, w, \tau)|pass\xi_0$ which does not depend on ξ_0. This subgame, by construction, is an optimal contract game in the T-model. Now take $T \to \infty$. By self-similarity $(\mathcal{M}, w, \tau)|pass\xi_0$ and $(\mathcal{M}, w, \tau)$ generate the same ex-ante payoffs. Thus, one can replace $(\mathcal{M}, w, \tau)|pass\xi_0$ with $(\mathcal{M}, w, \tau)$ itself. By doing this repeatedly across all dates, the resulting optimal contract game in the infinite horizon limit has the stationary structure described in the Theorem.

$\square$

Proof of Theorem 2. Let $a_t^*(\Delta)$ denote the limiting effort induced by the optimal contract at date t as $T \to \infty$ holding fixed $\Delta > 0$. Suppose $\lim_{\Delta\to 0} a_t^*(\Delta) > 0$. Since A is exerting an interior effort, the first-order condition equating marginal cost, $h'(a_t^*(\Delta))\Delta$, to marginal benefit,

$$\left(-\frac{d}{da_t}\mathbf{P}(X_t \in Bad \mid a_t, \Delta)|_{a_t=a_t^*(\Delta))}\right) \cdot p^*(\Delta) \cdot e^{-r\Delta} S^*(\Delta),$$

must hold. Since marginal cost $= \Theta(\Delta)$, therefore marginal benefit $= \Theta(\Delta)$. Since $e^{-r\Delta}S^*(\Delta) = \Theta(\Delta^0)$ and, by assumption, $-\frac{d}{da}\mathbf{P}(X_t \in Bad \mid a_t, \Delta)|_{a_1=a_t^*(\Delta))} = \Theta(\Delta^\alpha)$, therefore $p^*(\Delta) = \Theta(\Delta^{1-\alpha})$.

The contribution to surplus of $a_t^*(\Delta)$ relative to zero effort is $= \Theta(\Delta)$. The cost to surplus of $p^*(\Delta)$ relative to zero probability of termination is $\mathbf{P}(X_t \in Bad \mid a_t^*(\Delta), \Delta) \cdot p^*(\Delta) = \Theta(\Delta^{\gamma^b+(1-\alpha)})$. For the contributions to exceed the costs it must be that $\gamma^b + 1 - \alpha \geq 1 \Rightarrow \alpha - \gamma^b \leq 0 \Rightarrow \alpha - \gamma^b = 0$. Feasibility of $p^*(\Delta) = \Theta(\Delta^{1-\alpha})$ implies $\alpha \leq 1$. □

Proof of Lemma 2. Case 1a: $\gamma_1^g - \gamma_2^b < \alpha_1 - \alpha_2 = 0 < \gamma_1^b - \gamma_2^g$.

It is easy to show $\gamma_1^g = \gamma_2^g = 0$. By the product rule, as $\Delta \to 0$, the derivative of $\mathbf{P}(X_t = (g_1, b_2) \mid a_t, \Delta)$ with respect to a_t is $A(\Delta) - B(\Delta)$ where $A(\Delta) = \Theta(\Delta^{\alpha_1+\gamma_2^b})$ and $B(\Delta) = \Theta(\Delta^{\gamma_1^g+\alpha_2})$. Since $\alpha_1 + \gamma_2^b > \gamma_1^g + \alpha_2$, $B(\Delta) \gg A(\Delta)$ and therefore $(g_1, b_2) \in Bad$. By the product rule, as $\Delta \to 0$, the derivative of $\mathbf{P}(X_t = (b_1, g_2) \mid a_t, \Delta)$ with respect to a_t is $-A(\Delta) + B(\Delta)$ where $A(\Delta) = \Theta(\Delta^{\alpha_1+\gamma_2^g})$ and $B(\Delta) = \Theta(\Delta^{\gamma_1^b+\alpha_2})$. Since $\alpha_1 + \gamma_2^g < \gamma_1^b + \alpha_2$, $A(\Delta) \gg B(\Delta)$ and therefore $(b_1, g_2) \in Bad$.

Given the results above, $\gamma^b = \min\{\gamma_1^b + \gamma_2^b, \gamma_1^g + \gamma_2^b, \gamma_1^b + \gamma_2^g\} = \min\{\gamma_1^b, \gamma_2^b\}$. $\gamma^g = \gamma_1^g + \gamma_1^g = 0$. $\alpha = \min\{\alpha_1 + \gamma_2^g, \gamma_1^g + \alpha_2\} = \alpha_1 = \alpha_2$.

Case 1b: $\gamma_1^g - \gamma_2^b \leq 0 < \alpha_1 - \alpha_2 < \gamma_1^b - \gamma_2^g$.

$\gamma_1^g = 0$, $\gamma_1^b > 0$. $(g_1, b_2) \in Bad$, $(b_1, g_2) \in Bad$. $\gamma^b = \min\{\gamma_1^b + \gamma_2^b, \gamma_1^g + \gamma_2^b, \gamma_1^b + \gamma_2^g\} = \min\{\gamma_1^b + \gamma_2^g, \gamma_2^b\} = \min\{\gamma_1^b, \gamma_2^b\}$. $\gamma^g = \gamma_2^g$. $\alpha = \min\{\alpha_1 + \gamma_2^g, \gamma_1^g + \alpha_2\} = \alpha_2$.

Case 2: $\gamma_1^b - \gamma_2^g \leq 0 < \alpha_1 - \alpha_2 < \gamma_1^g - \gamma_2^b$.

$\gamma_1^b = 0$, $\gamma_1^g > 0$. $(g_1, b_2) \in Good$, $(b_1, g_2) \in Good$. $\gamma^b = \gamma_2^b$. $\gamma^g = \min\{\gamma_1^g + \gamma_2^g, \gamma_1^g + \gamma_2^b, \gamma_1^b + \gamma_2^g\} = \min\{\gamma_1^g + \gamma_2^b, \gamma_2^g\} = \min\{\gamma_1^g, \gamma_2^g\}$. $\alpha = \min\{\alpha_1 + \gamma_2^b, \gamma_1^b + \alpha_2\} = \alpha_2$.

Case 3a: $\gamma_1^g - \gamma_2^b \leq 0 \leq \gamma_1^b - \gamma_2^g < \alpha_1 - \alpha_2$.

$\gamma_1^g = 0$ or $\gamma_2^g = \gamma_1^b = 0$. $(g_1, b_2) \in Bad$, $(b_1, g_2) \in Good$. $\gamma^b = \min\{\gamma_1^b + \gamma_2^b, \gamma_1^g + \gamma_2^b\} = \gamma_2^b$. $\gamma^g = \min\{\gamma_1^g + \gamma_2^g, \gamma_1^b + \gamma_2^g\} = \gamma_2^g$. $\alpha = \min\{\alpha_1 + \gamma_2^b, \gamma_1^b + \alpha_2, \gamma_1^g + \alpha_2\} = \alpha_2$.

Case 3b: $\gamma_1^b - \gamma_2^g \leq 0 \leq \gamma_1^g - \gamma_2^b < \alpha_1 - \alpha_2$.

$\gamma_1^b = 0$ or $\gamma_1^g = \gamma_2^b = 0$. $(g_1, b_2) \in Bad$, $(b_1, g_2) \in Good$. $\gamma^b = \min\{\gamma_1^b + \gamma_2^b, \gamma_1^g + \gamma_2^b\} = \gamma_2^b$. $\gamma^g = \min\{\gamma_1^g + \gamma_2^g, \gamma_1^b + \gamma_2^g\} = \gamma_2^g$. $\alpha = \min\{\alpha_1 + \gamma_2^b, \gamma_1^b + \alpha_2, \gamma_1^g + \alpha_2\} = \alpha_2$. □

Definition. *Fix a function* $H : [0,1) \to [0,\infty)$ *with* $H(0) = H'(0) = 0$, $H'' > 0$, *and* $lim_{a\to 1} H(a) = \infty$. *Let* $\mathcal{F}_W$ *denote the set of measurable functions* $f : (-\infty,\infty) \to [0,W]$. *Given* $\rho \in [-\infty,\infty]$, *define* 1_ρ *to be the function that is 0 on* $(-\infty,\rho]$ *and 1 on* (ρ,∞). *Let* $N(a,\sigma^2)$ *denote a normal random variable with mean a and variance* σ^2. *Given* $f \in F_W$ *and a differentiable strictly increasing function* $g : [0,1) \to (-\infty,\infty)$ *define* $\mathcal{A}(H,f,g,\sigma^2)$ *to be the set*

$$\arg\max_{a\in[0,1)} \mathbf{E}f(N(g(a),\sigma^2)) - H(a).$$

For each $x \in \mathbb{R}$, *define* $\Phi_{\rho,\sigma^2}(x) := \mathbf{P}(N(x,\sigma^2) \geq \rho)$.

Lemma 4. *There exists a pair* $\rho(H,W,g,\sigma^2)$ *and* $a(H,W,g,\sigma^2) \in a(H, W1_{\rho(H,W,g,\sigma^2)}, g, \sigma^2)$ *such that there does not exist an* $f \in \mathcal{F}_W$ *with nonempty* $\mathcal{A}(H,f,g,\sigma^2)$ *and an element* $a' \in \mathcal{A}(H,f,g,\sigma^2)$ *such that* $a' > a(H,W,g,\sigma^2)$. *If* $W > 0$, $\rho(H,W,g,\sigma^2)$ *and* $a(H,W,g,\sigma^2)$ *are unique and* $\rho(H,W,g,\sigma^2) \in [g(0), g(a(H,W,g,\sigma^2))]$.

Proof. Fix H, f, g, and σ^2 with non-empty $\mathcal{A}(H,f,g,\sigma^2)$. Since $\mathcal{A}(H,f,g,\sigma^2)$ is compact, let a' be the maximum element. Let $\rho' \in \mathbb{R}$ satisfy

$$\mathbf{E}W1_{\rho'}(N(g(a'),\sigma^2) = \mathbf{E}f(N(g(a'),\sigma^2)).$$

By construction, we have $a < a' \Rightarrow \mathbf{E}W1_{\rho'}(N(g(a),\sigma^2) \leq \mathbf{E}f(N(g(a),\sigma^2))$ and $a > a' \Rightarrow \mathbf{E}W1_{\rho'}(N(g(a),\sigma^2) \geq \mathbf{E}f(N(g(a),\sigma^2))$. Thus, $\mathcal{A}(H,W1_{\rho'},g,\sigma^2)$, which is non-empty, has an element that is weakly larger than a'.

Now consider the set of a such that there exists a ρ satisfying $a \in \mathcal{A}(H,W1_\rho,g,\sigma^2)$. This is a closed set: Let a_∞ be the limit of a sequence $\{a_n\}$ in this set. To each a_n, I can associate a threshold ρ_n such that $a_n \in \mathcal{A}(H,W1_{\rho_n},g,\sigma^2)$. There is a convergent subsequence in the compact set $[-\infty,\infty]$ with limit ρ_∞. $a_\infty \in \mathcal{A}(H,W1_{\rho_\infty},g,\alpha^2)$. Thus, there exists a maximum element a^* with associated threshold ρ^*. a^* is the largest element of $\mathcal{A}(H,W1_{\rho^*},g,\sigma^2)$. This proves existence.

To prove uniqueness, assume $W > 0$. Suppose $\rho^* < g(0)$. Then $\Phi'_{g(0),\sigma^2}(x) > \Phi'_{\rho^*,\sigma^2}(x)$ for all $x \geq g(0)$. Thus, $\mathcal{A}(H,W1_{g(0)},g,\sigma^2)$ has an element $> a^*$. Contradiction.

Suppose $\rho^* > g(a^*)$. Since $W > 0$, it must be that $a^* > 0$ and $W\Phi'_{\rho^*,\sigma^2}(g(a^*))g'(a^*) = H'(a^*)$. Now consider the function $\Phi_{g(a^*),\sigma^2}$. By construction, $W\Phi'_{g(a^*),\sigma^2}(g(a^*))g'(a^*) > H'(a^*)$. Moreover, $\Phi'_{g(a^*),\sigma^2}(x) > \Phi'_{\rho^*,\sigma^2}(x)$ for all $x \leq g(a^*)$. Thus, $\mathcal{A}(H,W1_{g(a^*)},g,\sigma^2)$ has an element $x > a^*$. Contradiction. Thus, $\rho^* \in [g(0), g(a^*)]$. Suppose there were two such thresholds $\rho_1^* \neq \rho_2^*$. We then have $\Phi'_{\rho_1^*,\sigma^2}(g(a^*)) = \Phi'_{\rho_2^*,\sigma^2}(g(a^*))$. Since $\rho_1^* \neq \rho_2^*$ and Φ_{ρ,σ^2} is logistic shaped with inflection point equal to ρ, it must be that $\max\{\rho_1^*,\rho_2^*\} > g(a^*)$. Contradiction. □

Definition. *Fix a* $D > 0$ *divisible by* Δ *and an* $a \in [0,1)$. *For* $t = 0, \Delta, 2\Delta \ldots$ *let* $a(t)$ *denote the unique effort level satisfying* $e^{-rt}h'(a(t)) = h'(a)$. *Define* $H^D(a) = \sum_{t=0}^{D-\Delta} e^{-rt}h(a(t))\Delta$ *and* $a^D = \sum_{t=0}^{D-\Delta} a(t)\Delta$.

Proof of Lemma 3 and Theorem 4. The proof relies on Lemma 4 and the two definitions above.

Based on Zhu (2018) and Lemma 4, I take as given that for any contract game $(\mathcal{M}, w, \tau, e)$ the credible threats refinement selects the set of sequential equilibria $\mathcal{E}^*(\mathcal{M}, w, \tau, e)$ – the credible threats equilibria – constructed as follows:

Since sampling dates are predictable, the end of each sampling period is known at the beginning of that period. In particular, given a sampling date t it is known at the end of date t if this was the final sampling or not. Fix a final sampling date e_k for some k. The termination date is then measurable with respect to h_{e_k} since there are no more messages or public randomizing devices after date e_k. Define

$$W^*_{e_k+\Delta}(h_{e_k}) = \left(\sum_{t=e_k+\Delta}^{\tau(h_{e_k})-\Delta} e^{-r(t-e_k-\Delta)} w_t(h_{e_k})\Delta\right) + e^{-r(\tau(h_{e_k})-e_k-\Delta)} w_{\tau(h_{e_k})}.$$

and $V^*_{e_k+\Delta}(h_{e_k}) := \left(\sum_{t=e_k+\Delta}^{\tau(h_{e_k})-\Delta} e^{-r(t-e_k-\Delta)} u(0)\Delta\right) - W^*_{e_k+\Delta}(h_{e_k})$. Next, fix a non-final sampling date e_k for some k. Define

$$\mathcal{M}^*_{e_{k+1}}(h_{e_k}) = \operatorname*{arg\,max}_{m_{e_{k+1}} \in \mathcal{M}_{e_{k+1}}(h_{e_k})} \mathbf{E}_{\xi_{e_{k+1}}} V^*_{e_{k+1}+\Delta}(h_{e_k} m_{e_{k+1}} \xi_{e_{k+1}})$$

$$V^{max}_{e_{k+1}+\Delta}(h_{e_k}) = \max_{m_{e_{k+1}} \in \mathcal{M}_{e_{k+1}}(h_{e_k})} \mathbf{E}_{\xi_{e_{k+1}}} V^*_{e_{k+1}+\Delta}(h_{e_k} m_{e_{k+1}} \xi_{e_{k+1}})$$

$$W^{pass}_{e_{k+1}+\Delta}(h_{e_k}) = \max_{m_{e_{k+1}} \in \mathcal{M}^*_{e_{k+1}}(h_{e_k})} \mathbf{E}_{\xi_{e_{k+1}}} W^*_{e_{k+1}+\Delta}(h_{e_k} m_{e_{k+1}} \xi_{e_{k+1}})$$

$$W^{fail}_{e_{k+1}+\Delta}(h_{e_k}) = \min_{m_{e_{k+1}} \in \mathcal{M}^*_{e_{k+1}}(h_{e_k})} \mathbf{E}_{\xi_{e_{k+1}}} W^*_{e_{k+1}+\Delta}(h_{e_k} m_{e_{k+1}} \xi_{e_{k+1}})$$

$$a^*(h_{e_k}) = \max\Bigg\{ \operatorname*{arg\,max}_{a \in [0,1)} -H^{e_{k+1}-e_k}(a) - e^{-r(e_{k+1}-e_k)} \cdot \mathbf{P}\Big(N(a^{e_{k+1}-e_k}, e_{k+1}-e_k) \leq \rho(H^{e_{k+1}-e_k}, e^{-r(e_{k+1}-e_k)}(W^{pass}_{e_{k+1}+\Delta}(h_{e_k}) - W^{fail}_{e_{k+1}+\Delta}(h_{e_k})), a^{e_{k+1}-e_k}, e_{k+1}-e_k)\Big) \cdot (W^{pass}_{e_{k+1}+\Delta}(h_{e_k}) - W^{fail}_{e_{k+1}+\Delta}(h_{e_k})) \Bigg\}$$

$$W^*_{e_k+\Delta}(h_{e_k}) = \sum_{t=e_k+\Delta}^{e_{k+1}} (w_t(h_{e_k}) - h(a^*(h_{e_k})(t-e_k-\Delta)))\Delta +$$

$$e^{-r(e_{k+1}-e_k)}\Bigg[W^{pass}_{e_{k+1}+\Delta}(h_{e_k}) - \mathbf{P}\Bigg(N(a^{e_{k+1}-e_k}, e_{k+1}-e_k) \leq$$

$$\rho(H^{e_{k+1}-e_k}, e^{-r(e_{k+1}-e_k)}(W^{pass}_{e_{k+1}+\Delta}(h_{e_k}) - W^{fail}_{e_{k+1}+\Delta}(h_{e_k})), a^{e_{k+1}-e_k}, e_{k+1}-e_k)\Bigg)\cdot$$

$$(W^{pass}_{e_{k+1}+\Delta}(h_{e_k}) - W^{fail}_{e_{k+1}+\Delta}(h_{e_k}))\Bigg]$$

$$V^*_{e_k+\Delta}(h_{e_k}) = \sum_{t=e_k+\Delta}^{e_{k+1}} (u(a^*(h_{e_k})(t-e_k-\Delta)) - w_t(h_{e_k}))\Delta + e^{-r(e_{k+1}-e_k)}V^{max}_{e_{k+1}+\Delta}(h_{e_k})$$

By defining $e_0 = -\Delta$, one can also define (W^*_0, V^*_0) in the same way as above.

By induction, I have now constructed a unique continuation payoff process

$$(W^*_{e_k+\Delta}(h_{e_k}), V^*_{e_k+\Delta}(h_{e_k}))$$

across all sampling dates e_k and corresponding histories h_{e_k}. I refer to this process as the continuation payoff process of $(\mathcal{M}, w, \tau, e)$ and (W^*_0, V^*_0) as the ex-ante payoff of $(\mathcal{M}, w, \tau, e)$.

Fix any strategy profile (a, m) satisfying

$$\begin{aligned} a_t(H^A_{t-\Delta}, h_{t-\Delta}) &= a^*(h_{e_k})(t-e_k-\Delta) \text{ for all } e_k < t \leq e_{k+1} \\ m_{e_{k+1}}(H^P_{e_{k+1}}, h_{e_k}) &\in \mathcal{M}^*_{e_{k+1}}(h_{e_k}) \\ \mathbf{E}_{\xi_{e_{k+1}}} W^*_{e_{k+1}+\Delta}(h_{e_k} m_{e_{k+1}}(H^P_{e_{k+1}}, h_{e_k})\xi_{e_{k+1}}) &= W^{pass}_{e_{k+1}+\Delta}(h_{e_k}) \text{ if } X_t \in Good \\ \mathbf{E}_{\xi_{e_{k+1}}} W^*_{e_{k+1}+\Delta}(h_{e_k} m_{e_{k+1}}(H^P_{e_{k+1}}, h_{e_k})\xi_{e_{k+1}}) &= W^{fail}_{e_{k+1}+\Delta}(h_{e_k}) \text{ if } X_t \in Bad \end{aligned}$$

for all t, $h_{t-\Delta}$ and sampling dates e_{k+1} with corresponding histories $h_{e_{k+1}}$. There exists an assessment with such a strategy profile that is a sequential equilibrium. All such sequential equilibria generate the continuation payoff process of $(\mathcal{M}, w, \tau, e)$. $\mathcal{E}^*(\mathcal{M}, w, \tau, e)$ is defined to be this set of sequential equilibria. An element of this set is called a credible-threats equilibrium.

Recycling the proof of Theorem 1, it is without loss of generality to assume in a contract game each message space $\mathcal{M}_{e_{k+1}}(h_{e_k}) = \{pass,\ fail\}$ and there is a credible threats equilibrium with the property that $m_{e_{k+1}}(H^P_{e_{k+1}}, h_{e_k}) = pass$ iff $Z_{e_{k+1}} - Z_{e_k} > \rho(H^{e_{k+1}-e_k}, e^{-r(e_{k+1}-e_k)}(W^{pass}_{e_{k+1}+\Delta}(h_{e_k}) - W^{fail}_{e_{k+1}+\Delta}(h_{e_k})), a^{e_{k+1}-e_k}, e_{k+1}-e_k)$.

Suppose by induction Lemma 3 is true for all models of length $\leq T$. Consider the

$T+\Delta$-model. Fix a contract game $(\mathcal{M}, w, \tau, e)$ and define

$$z \in \underset{z' \in Im(\xi_{e_1})}{\arg\max} W^*_{e_1+\Delta}(pass\ z') + V^*_{e_1+\Delta}(pass\ z').$$

Change the contract game as follows: Replace the subgame following $pass\ z'$ with the one following $pass\ z$ for all $z' \in Im(\xi_{e_1})$. Replace the subgame following $fail$ with the randomization between the one following $pass\ z$ and termination with $w_{e_1+\Delta} = -V^*_{e_1+\Delta}(pass\ z)$ calibrated so that A's expected payoff is $W^*_{e_1+\Delta}(pass\ z) - (W^{pass}_{e_1+\Delta} - W^{fail}_{e_1+\Delta})$. Under this sequence of changes, the difference between A's expected payoff following $pass$ and $fail$ does not change. Thus the ex-ante surplus weakly increases. Finally shift w between dates 0 and e_1 so that A's ex-ante payoff remains W^*_0. This contract game's ex-ante payoff weakly Pareto-dominates the original game's ex-ante payoff. By construction, it has a deterministic e_1 and all subsequent sampling dates do not depend on m_{e_1} and ξ_{e_1}. By induction, all subsequent sampling dates do not depend on messages or public randomizations after e_1. Thus, the contract has a deterministic sequence of sampling dates. The self-similarity of the infinite horizon limit means that there is a sequence of optimal contract games such that the deterministic sequence of sampling times converges to an evenly spaced sequence as $T \to \infty$. It is straightforward to show that this even spacing must have a convergent subsequence as $\Delta \to 0$.

□

Proof of Theorem 5. By Lemma 4 and the characterization of credible threats equilibria, the optimal contract game's threshold $\rho(D)$ is always between 0 and D. As $D \to 0$, $\mathbf{P}(N(a^D, D) < \rho) \to 0.5$ for any $a \in [0,1)$ and $\rho \in [0, D)$. Thus, $\mathbf{P}(Z_{kD} - Z_{(k-1)D} \leq \rho^*(\Delta) \mid a^*(\Delta)) = \Theta(1)$ as $\Delta \to 0$ where $a^*(\Delta)$ is the effort sequence induced within a sampling period of the optimal contract given sampling frequency $\frac{1}{\Delta}$. This implies that the negative effort-elasticity of *Bad* signals goes to zero as $\Delta \to 0$. Now by Theorem 2 the effort induced by the optimal contract goes to zero as $\Delta \to 0$. □

References

[1] "A Major Oil and Gas Company Faces Expansion," Retrieved from `https://www.humanyze.com/case-studies-oil-and-gas/`

[2] Abreu, D., P. Milgrom and D. Pearce (1991) "Information and Timing in Repeated Partnerships," *Econometrica* Vol. 59, pp. 1713-1733

[3] Barron, D. and Y. Guo (2019) "The Use and Misuse of Coordinated Punishments," working paper

[4] Blackwell, D. (1951) "Comparison of Experiments," *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability* pp. 93-102

[5] Buckingham, M. and A. Goodall (2015, October) "Reinventing Performance Management," Retrieved from `https://hbr.org/2015/04/reinventing-performance-management`

[6] Burkus, D. (2016, June 1) "How Adobe Scrapped Its Performance Review System And Why It Worked," Retrieved from `https://www.forbes.com/sites/davidburkus/2016/06/01/how-adobe-scrapped-its-performance-review-system-and-why-it-worked/#76fd1a0155e8`

[7] Cappelli, P. and A. Tavis (2016, October) "The Performance Management Revolution," Retrieved from https://hbr.org/2016/10/the-performance-management-revolution

[8] Caprino, K. (2016, December 13) "Separating Performance Management From Compensation: New Trend For Thriving Organizations," Retrieved from `https://www.forbes.com/sites/kathycaprino/2016/12/13/separating-performance-management-from-compensation-new-trend-for-thriving-organ` `#6b0e119877d4`

[9] Crémer, J. (1995) "Arm's Length Relationships," *Quarterly Journal of Economics* Vol. 110, pp. 275-295

[10] Dewatripont, M. (1987) "The Role of Indifference in Sequential Models of Spatial Competition: An Example," *Economics Letters* Vol. 23, pp. 323-328

[11] Dewatripont, M., I. Jewitt, and J. Tirole (1999) "Economics of Career Concerns, Part I: Comparing Information Structures," *Review of Economic Studies* Vol. 66, pp. 183-198

[12] Doerr, J. (2018, November 26) "Why Adobe Killed Performance Reviews," Retrieved from `https://www.whatmatters.com/stories/why-adobe-killed-performance-reviews`

[13] Fuchs, W. (2007) "Contracting with Repeated Moral Hazard and Private Evaluations," *American Economic Review* Vol. 97, pp. 1432-1448

[14] Gjesdal, F. (1982) "Information and Incentives: The Agency Information Problem," *Review of Economic Studies* Vol. 49, pp. 373-390

[15] Hirshleifer, J. (1971) "The Private and Social Value of Information and the Reward to Inventive Activity," *American Economic Review* Vol. 61, pp. 561-574

[16] Hoffman, L. (2017) "Goldman Goes Beyond Annual Review With Real-Time Employee Feedback," Retrieved from https://www.wsj.com/articles/goldman-goes-beyond-annual-review-with-real-time-employee-ratings-1492786653

[17] Holmstrom, B. (1979) "Moral Hazard and Observability," *Bell Journal of Economics* Vol. 10, pp. 74-91

[18] Holmstrom, B. (1999) "Managerial Incentive Problems: A Dynamic Perspective," *Review of Economic Studies* Voi. 66, pp. 169-182

[19] Kandori, M. (1992) "The Use of Information in Repeated Games with Private Monitoring," *Review of Economic Studies* Vol. 59, pp. 581-593

[20] Ledford, G. E., G. Benson, and E. E. Lawler (2016) "Cutting Edge Performance Management," *WorldatWork Research*

[21] MacLeod, W. B. (2003) "Optimal Contracting with Subjective Evaluation," *American Economic Review* Vol. 93, pp. 216-240

[22] Prat, A. (2005) "The Wrong Kind of Transparency," *American Economic Review* Vol. 95, pp. 862-877

[23] Ramirez, J. C. (2018, March 15) "The Future of Feedback," Retrieved from http://hrexecutive.com/future-of-feedback/

[24] Sannikov, Y. and A. Skrzypacz (2007) "Impossibility of Collusion under Imperfect Monitoring with Flexible Production," *American Economic Review* Vol. 97, pp. 1794-1823

[25] Schlee, E. E. (2001) "The Value of Information in Efficient Risk-Sharing Arrangements," *American Economic Review* Vol. 91, pp. 509-524

[26] Schrage, M., D. Kiron, B. Hancock, and R. Breschi (2019, February 26) "Performance Management's Digital Shift," Retrieved from https://sloanreview.mit.edu/projects/performance-managements-digital-shift/

[27] Tranæs, T. (1998) "Tie-Breaking in Games of Perfect Information," *Games and Economic Behavior* Vol. 22, pp. 148-161

[28] Weber, L. (2016, August 21) "At Kimberly-Clark, 'Dead Wood' Workers Have Nowhere to Hide," Retrieved from https://www.wsj.com/articles/focus-on-performance-shakes-up-stolid-kimberly-clark-1471798944

[29] Wilson, R. (1975) "Information Economies of Scale," *Bell Journal of Economics* Vol. 6, pp. 184-195

[30] Zhu, J. Y. (2018) "A Foundation for Efficiency Wage Contracts," *American Economic Journal: Microeconomics* Vol. 10, pp. 248-288