Better Monitoring . . . Worse Productivity?

John Y. Zhu[1]

December 19, 2019

Abstract

Technological advances are enabling firms to collect more information 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 the provision of 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: abuse, non-contractible information, incentive power, statistical power, monitoring design, privacy, optimal contracts, principal-agent, moral hazard

[1]Department of Economics, Yale University (email: johnzhuyiran@gmail.com). I thank Dilip Abreu, Sylvain Chassang, Davide Cianciaruso, Tommaso Denti, Johannes Horner, David Pearce, Bill Sandholm, Eddie Schlee and seminar participants at University of Colorado Boulder Leeds, Cornell, University of Rochester, University of Minnesota, University of Washington at St Louis, INSEAD, HEC, University of Maryland Smith, UC Riverside, University of Wisconsin, Financial Research Association Conference, Federal Reserve Bank of Richmond, University of Kentucky Gatton, Michigan State Broad, Kansas Workshop on Economic Theory.

1 Introduction

A performance management revolution is reshaping the workplace driven in large part by a vast expansion in monitoring. 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 greatly increased the quantity of worker data that can be collected at any moment in time.[3] The potential benefits of all this extra information are legion – better worker development, more timely feedback, improved coordination amongst team members etc. However one drawback is that most of this information is not directly contractible and is, instead, filtered through a manager's discretionary feedback.[4] A natural concern then is that managers might somehow *abuse* the information, leading to distorted incentives.[5] I investigate and clarify this concern through a dynamic moral hazard model in which monitoring generates non-contractible information and incentive contracts depend

[2]CPM practices emphasize frequent performance appraisals, often in ratingless form and sometimes drawn from multiple sources (Ledford, Benson, and Lawler, 2016). They often depend on new technologies such as apps that can efficiently collect and aggregate information about performance. 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 the promise of 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]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.)

[5]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 continuous performance management systems – 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).

on the principal's reports about agent performance. Depending on the information structure, equilibria can involve principal reporting behavior resembling abuse. The anticipation of this abuse then distorts optimal contracting ex-ante. Generating extra information can trigger or intensify the principal's abusive tendency and lead to a worse contracting outcome. I then characterize what kinds of improvements to the information content of monitoring are counterproductive for the provision of incentives and discuss implications for when and how monitoring should be limited.

What is the abusive tendency? A typical contract will specify a "stick" such as termination that can be used as a threat to motivate the agent to put in effort. Since information is non-contractible, whether the stick gets used depends entirely on the principal's reports. This means even under a binding contract the principal still has considerable discretion in deciding how easily poor performance triggers the use of the stick. In this discretionary setting I show some improvements to the information content of monitoring tempt the principal to be abusive by giving her the ability to induce a lot of effort from the agent – even an inefficiently high amount of effort – but only if she makes heavy use of the stick. The principal is tempted because, ex-post, the agent bears the cost of the stick while the principal reaps the benefits from the effort being induced. Anticipating such abusive behavior, 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 monitoring. Conversely, reducing the amount of information generated by monitoring can be beneficial if it reduces how abusive the principal is tempted to be. One manifestation 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 (ab)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 (see the characterization in Section 4.2 of information that is weak in *statistical power* but strong in *incentive power*), and that maintaining formal, periodic performance reviews can – if done in the right way – be beneficial (Section 5).

Related Literature

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.[6]

The point I make about the potential benefits of reducing the amount of information generated by monitoring is more basic. Rather than adding another dimension to the baseline model, I observe that Holmstrom'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. I then tie this non-contractibility to the principal being abusive and show that more informative monitoring technologies can intensify the principal's abusive tendency and lead to a worse contracting outcome.

A growing literature is exploring, in a variety of settings, the impact of abusive behaviors of or related to the type considered in this paper and its relation to how much information players have. Zhu (2018) shows how such behaviors can simplify the optimal contract in a general dynamic moral hazard setting from something quite intricate to a memoryless efficiency wage contract. Baron and Guo (2019) consider a principal-agents model where the ability of agents to extort the principal can completely destroy cooperation. They then show how the release of certain types of contractible information can restore cooperation. Padró i Miquel and Chassang (2019) consider a principal-agent-monitor model where the agent can intimidate the monitor. The principal can reduce the agent's desire to intimidate and improve her own payoff by garbling the monitor's unverifiable information.

Gigler and Hemmer (1998), Lizzeri, Meyer, and Persico (2002), and Fuchs (2007) show in dynamic settings how limiting feedback can be beneficial. While this less feedback/better outcome result sounds similar to my better monitoring/worse outcome result, the two results are distinct with different intuitions. The less feedback/better outcome result is about how it is optimal not to reveal anything to the agent until the end of the relationship. It is not about how much information is generated by monitoring. In those models more informative monitoring leads to sharper terminal incentives and, therefore, a better outcome. Moreover, the reason why interim feedback is suboptimal in those models is because such feedback generates additional incentive constraints on the agent's side making it more costly to sustain effort throughout the relationship. In contrast, more informative monitoring can be bad in my model because it can cause the principal to be more abusive.

The idea that more information can lead to a worse contracting outcome has also been noted 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]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 related models.

[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 monitoring in the sense of Blackwell (1950) can lead to a less efficient outcome. Garbling allows deterministic contracts to mimic random contracts which can be beneficial under nonseparability.

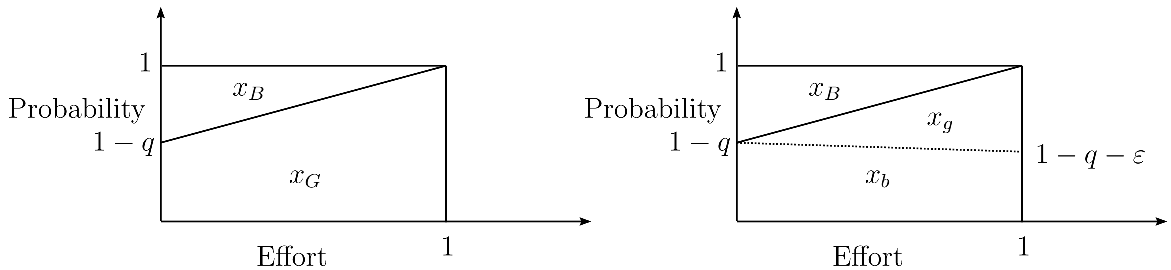

(a) Monitoring technology X. (b) Better monitoring technology X'.

Figure 1

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 $q \in (0, 1)$. See Figure 1a. 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 $w - h(a) - \mathbf{E}_a rp$. My definition of a contract game rules out performance sensitive pay $w(X)$. Since X is non-contractible, this is without loss: P would always choose to pay $\min_X w(X)$ regardless of X. In the next subsection I discuss how the model can be modified so that performance sensitive pay can play a meaningful role. The better monitoring/worse outcome result I am about to derive will continue to hold under that modification.

The timing of events is as follows: P offers a contract game. A chooses whether or not to accept P's offer. If he does not accept both parties exercise outside options normalized to 0. If he does accept then P announces a punishment strategy r, then A chooses an effort a, and finally payoffs are realized.

I now characterize the optimal contract game (w^*_X, p^*_X) under monitoring technology X along with the corresponding equilibrium. Given a contract game offer (w, p), if A accepts P will announce the punishment strategy that induces the most effort. Why? After a contract game is accepted, w is fixed and A can no longer exercise his outside option. This makes P's payoff $u(a) - w$ purely an increasing function of a. The unique punishment strategy that induces the most effort 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 punishment strategy 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 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 $w_X^* = h(a_X^*) + (1 - a_X^*)qp_X^*$.

Next, 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 $[0,1)$. See Figure 1b. 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'. By the same logic as before, 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 just slightly more than under X. Thus the argmax expression used to find the optimal punishment is approximately the one used before plus a $-(1-q)p$ term approximating the additional amount of punishment compared to before:

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

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

Why can't P ignore the extra information contained in X', making better monitoring/worse outcome impossible? 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. However 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 *abusive*: Ex-ante, P would like to commit to a restrained punishment strategy that punishes only x_B. But ex-interim she cannot help but announce the harsher punishment strategy since A is now locked into the contract. This dynamic has the flavor of a hold-up problem. See, for example, Williamson (1979). 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 announce the punishment strategy that has A suffering the big punishment for both x_b and x_B. Ex-ante A anticipates that P cannot help but be abusive under the better monitoring technology X'. To prevent A from quitting, P is forced to offer a contract game with a smaller punishment that then hinders incentive provision and leads to a worse outcome.[9]

The motivating example suggests the following channel through which better monitoring can negatively affect the provision of incentives: *The principal wants to use the stick whenever a bad signal is generated. When an improvement to monitoring generates lots of low quality bad signals the principal's behavior becomes abusive. 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.*

The Scope of Better Monitoring/Worse Outcome

In the motivating example, after A accepts the contract game offer, P can induce slightly more effort and receive slightly more expected utility under X' than under X by announcing a strategy that punishes the almost completely uninformative signal x_b. P's inability to not take advantage of this small opportunity leads to a worse contracting outcome ex-ante.

This begs the question: Suppose before A accepts the contract game offer, P could commit to announce any punishment strategy that is, after the game is accepted, close to but not necessarily optimal. Does the better monitoring/worse outcome result go away? In the motivating example, adding such a modicum of commitment power would make the result go away but in general the answer is no.

I will show there are ways in general to improve monitoring that allow P to induce significantly higher effort but require an even more significant amount of excessive punishment. Having a little bit of commitment power may allow P to resist the temptation of inducing slightly higher effort but not significantly higher effort. Now the same logic as before leads to a worse contracting outcome.

In the motivating example, performance sensitive pay $w(X)$ is shut down without loss. The reason is P is both the monitor and the employer who pays A out of her own pocket. Therefore, she would never report a signal that did not lead to the smallest payment. Meaningful use of performance sensitive pay can be restored if the roles of monitor and employer are separated. I now sketch such a model and argue that the better monitoring/worse outcome phenomenon can continue to appear. The main assumptions I need are that the monitor's utility is an increasing function of effort and A is risk-averse. The original analysis assumed a risk-neutral A but it could

[9]Alternatively, P could keep the punishment size the same and just compensate A with a higher upfront payment. However (1) implies this is never optimal – reducing the punishment threat mitigates the decline in surplus that P can then recoup ex-ante with a lower upfront pay.

easily have been generalized to a risk-averse A (and a risk-averse P).

There are now three players, P, M, and A. A contract game specifies a pair of payments consisting of a guaranteed base pay $w \in \mathbb{R}$ and a discretionary bonus $b \geq 0$. Just as in the original example it was without loss to specify a single punishment threat, so it is without loss here to specify a single bonus. The timing is as follows: P offers a contract game (w, b). A accepts or rejects. If reject, all parties receive 0. If accept, M announces to A a bonus payment strategy $r : X \to \{0, 1\}$ where 0 is not paying the bonus and 1 is paying the bonus. A then chooses effort and finally payoffs are realized. P's payoff is $u(a) - w - \mathbf{E}_a rb$. M's payoff is some strictly increasing function $f(a)$. A's payoff is $\mathbf{E}_a v(w + rb) - h(a)$.

When v is the identity function any improvement to X has a neutral effect on P's payoff since the surplus maximizing effort level is induced by the optimal contract under X. In some of these cases, by adding a bit of concavity to v the improvement to X can now have a negative effect on P's payoff.

In the motivating example X has the property that each signal realization can be characterized as "good" or "bad" because its probability is a strictly monotonic function of effort. The advantage of assuming monotonic signals is that 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). MLRP signals imply a similarly simple punishment strategy for maximizing effort even though signal probabilities can now be non-monotonic in effort. More general signal structures can also be considered, but the better monitoring/worse outcome analysis becomes less tractable.

Looking Ahead

It should now be clear that the better monitoring/worse outcome channel applies quite generally. However, some important questions remain:

1. Signal splitting is abstract and can be hard to interpret. Can we restrict attention to an intuitive class of signal splittings 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 result 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 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 an intuitive way to improve monitoring. 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 that in the canonical setting in which a stochastic process tracks the agent's cumulative productivity and the principal monitors by sampling that process, restricting the principal to sample every once in a while can significantly improve the outcome of contracting.

   – Reducing sampling frequency reduces the informativeness of monitoring. The result provides a simple, broadly applicable demonstration of the better monitoring/worse outcome channel at work. 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 the release of information) is strictly counterproductive. This contrasts with well-known results in the repeated games literature that highlight the benefits of batching in other settings.

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 $t = 0, \Delta, 2\Delta, \ldots T$. 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 non-contractible 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 $\{X_s\}_{s \le t}$. 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

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

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

Intuitively contracts will provide incentives in this setting by being structured in a way so that at the end of each date t, no matter what message P reports, $V_{t+\Delta}$ remains constant, but $W_{t+\Delta}$ can vary. This is just like in the motivating example where no matter if P punishes or not, P's utility is the same but A's utility can be different. By reporting messages that lead to different values of $W_{t+\Delta}$ depending on the observed signal, P can induce A to put in effort at date t.

Incentive-Compatibility

Notice the order of actions has switched moving from the motivating example to the dynamic model. In the motivating example, P announces a punishment strategy before A chooses effort. In the dynamic model, A first chooses effort, then X_t is realized, and finally P chooses a message.

To understand why I switched the order, suppose the motivating example were reformulated to match the order of the dynamic model: A chooses effort, X is realized, then P chooses $r(X) \in \{0, 1\}$. Under this reformulation, given a contract game, multiple equilibria exist: Since P is indifferent between 0 and 1 at every decision node, any punishment strategy combined with A's best-response effort comprise an equilibrium. This set can be quite large. It yields no prediction about P's behavior, and A's equilibrium effort can vary greatly depending on what P does. Which if any of these equilibria are plausible?

To answer the question, first notice that among all these equilibria, the one most preferred by P – call it (a^*, r^*) – is the one featuring the highest induced effort and the punishment strategy that punishes A if and only if a *Bad* signal is observed. This is the *unique* equilibrium under the original formulation of the motivating example. In the reformulation of the motivating example could P somehow ensure that (a^*, r^*) is played? Intuitively, yes. Suppose initially P and A agreed to play some other equilibrium (a, r). Then at the beginning of the game, before A has exerted effort, P could threaten A by announcing that she will unilaterally deviate from the agreed upon strategy r to strategy r^*. Unlike the non-credible threats that are ruled out by subgame perfection, this credible threat does not involve P promising to do anything that her future self would not want to do. Therefore, A should take it seriously and best-respond to it. Anticipating this, P will indeed threaten A with r^*, causing the equilibrium to change from (a, r) to (a^*, r^*). The original formulation of the motivating example micro-founds the credible-threats intuition for selecting (a^*, r^*).[10]

Moving to the dynamic model, I would again like to use the credible-threats intuition to winnow down the set of plausible equilibria, ideally down to a unique one. However, when P and A interact repeatedly a time-consistency issue pops up that is not relevant in the one-shot game. At the beginning of the game, P now needs to announce a complete report strategy for the entire game, not just a report strategy for date 0. But as is common in dynamic incentive models, an announcement that is optimal for P at the beginning of the game need not be optimal for P later on.

Of course, at the beginning of date $T - \Delta$, time-consistency is not an issue. So one could allow P to announce her final report strategy then. Then via backward

[10]Tranaes (1998) develops the credible-threats selection of subgame-perfect equilibria in all finite extensive-form games. For other applications of the credible-threats selection, see Dewatripont (1987), Zhu (2018), and Baron and Guo (2019).

induction, one could model P announcing her date t continuation strategy at the beginning of date t subject to the restriction that starting at date $t+\Delta$, this strategy is the time-consistent report strategy she would announce at $t+\Delta$.[11] Zhu (2018) explicitly uses this idea to refine sequential equilibria and shows that the surviving sequence of reports and efforts comprise a sequential equilibrium – the credible threats equilibrium – that is unique up to equivalence in the continuation payoff process. This equilibrium can be compactly characterized as follows: Given a contract game $(\mathcal{M}, w, \tau)$, a sequential equilibrium (a, m) is the credible threats equilibrium if the continuation payoff process (W_t, V_t) depends only on the public history and given any $h_{t-\Delta}$, $m_t(H_t^P, h_{t-\Delta}) = m^{max}(h_{t-\Delta})$ $(= m^{min}(h_{t-\Delta}))$ if $X_t \in Good$ $(\in Bad)$ where

$$\begin{aligned}
\mathcal{M}_t^*(h_{t-\Delta}) &= \operatorname*{arg\,max}_{m_t \in \mathcal{M}_t(h_{t-\Delta})} V_{t+\Delta}(h_{t-\Delta}m_t), \\
m^{max}(h_{t-\Delta}) &\in \operatorname*{arg\,max}_{m_t \in \mathcal{M}_t^*(h_{t-\Delta})} W_{t+\Delta}(h_{t-\Delta}m_t), \\
m^{min}(h_{t-\Delta}) &\in \operatorname*{arg\,min}_{m_t \in \mathcal{M}_t^*(h_{t-\Delta})} W_{t+\Delta}(h_{t-\Delta}m_t).
\end{aligned}$$

In keeping with the motivating example, a contract in the dynamic model is incentive-compatible if the assessment is the credible threats equilibrium of the contract game.

3.1 The Optimal Contract

The optimal contracting problem is to find an incentive-compatible contract that maximizes V_0 subject to A'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.[12]

Theorem 1. *As $T \to \infty$, the optimal contract converges to a stationary contract with the following structure: There exist a fixed effort level a^* and a fixed probability p^* such that at each date $t < \tau$,*

- *A exerts effort a^*,*
- *$\mathcal{M}_t = \{pass, fail\}$,*
- *$m_t = fail$ iff $X_t \in Bad$,*

[11] The date $T-\Delta$ report strategy will actually be trivial since the contract terminates for sure at date T. As a result, A exerts 0 effort at date $T-\Delta$. However, by assumption, even zero effort generates positive surplus. This then creates the opportunity for P to announce a non-trivial report strategy starting at date $T-2\Delta$.

[12] Since W_t and V_t are both public in an incentive-compatible contract, violations of the interim participation constraint are common knowledge.

- *If $m_t = pass$ then A is retained for date $t + \Delta$,*
- *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.*

Proof. See appendix. □

The optimal contract for finite T has the same structure except a^* and p^* will, in general, both depend on t but nothing else, and termination occurs at date T regardless of the report at date $T - \Delta$.

Theorem 1 is silent about w because it is partially indeterminate. Without loss all payments can be delayed until termination provided they are compounded at the discount rate. w plays an auxiliary role in the optimal contract, ensuring at the end of each date t, P is indifferent between reporting *pass* and *fail*,

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

Incentives are instead driven by the threat of termination, parameterized by p^*.

The optimal contract and, by extension, all Pareto-optimal contracts can be computed as follows. Let S^* denote the Pareto-optimal surplus. At each date $t < \tau$, the continuation contract starting from date t is a Pareto-optimal contract. P's indifference between *pass* and *fail* implies that A bears the cost of inefficient termination,

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

The first-order condition that pins down A's effort level each date is

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

If there are multiple efforts that maximize A's utility, a^* is the highest one.

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

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

The solution can be recursively computed by setting $S_0^* = u(0)\Delta$ on the right hand side 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 bounded above. Its limit is S^*. p^* then falls out of the first equation.

With the optimal contract characterized, I now explore the better monitoring/worse outcome channel further by addressing the questions posed at the end of Section 2.

4 Better Monitoring Worse Outcome

In the optimal contract, at the end of each date t, P punishes A iff she observes a *Bad* signal at that date just like in the motivating example. Given what we have learned, it is clear how abusive reporting might pop up. Suppose the monitoring technology generates a lot of "low quality" *Bad* signals – I will formalize the notion of quality shortly – then clearly punishing the agent for every *Bad* signal that is generated is abusive. To counteract this behavior, the parties will then preemptively agree to an optimal contract that reduces the size of the punishment threat. That means setting p^* to be a low value. Indeed when the typical *Bad* signal is of 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 mostly 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. 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. I show that the Poisson event is a "high quality" *Bad* signal in some formal sense. Consequently the optimal contract induces positive effort – depending on how the model is 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 random walk approximating a Brownian motion with drift controlled by A's effort. 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 of 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 *Bad* signal that is the Poisson event? The measure that is of interest to me is the following product:

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

The first term of this product 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 shirked. 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 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 incentive power while the second factor is captured by 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 looking at the product of incentive power and statistical power.

In the more general analysis of the next section, I will formally show if the product of incentive power and statistical power 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. Anticipating this result, I now use the product of incentive power and statistical power 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 quality of the *Bad* signal 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 motion Y_t where effort controls the drift:

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

Whenever Y_t 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 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 costs 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 the bad news Poisson component 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 what 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 is of much lower quality than the original *Bad* signal. In fact, it is not hard to show that the quality 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 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 adding a Brownian component to bad news Poisson monitoring leads to a worse contracting 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 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,

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

Here, $\lim_{\Delta\to 0} f(\Delta) = \Theta(g(\Delta))$ means there exist $c_1, c_2, \varepsilon > 0$ such that $c_1 g(\Delta) \leq f(\Delta) \leq c_2 g(\Delta)$ for all $\Delta < \varepsilon$.

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 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 determined by the quality of Bad signals.*

Formally, assume $\alpha \leq 1$. If $\alpha - \gamma^b = 0$ then for any $a < 1$, there exist parameterizations of the rest of the model such that the effort a^* induced by the optimal contract is greater than a. Otherwise the optimal contract induces zero effort.

Proof. See appendix. □

Theorem 2 implies that in the continuous time limit if the quality of *Bad* signals as measured by the product of incentive power and statistical power 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 the quality of *Bad* signals 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 about the relationship between monitoring and surplus, with their emphasis on incentive power 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.[13]

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,

[13] In a repeated games setting with public monitoring, Kandori (1992) shows that making monitoring more informative in the sense of Blackwell (1951) causes the pure-strategy sequential equilibrium payoff set to expand in the sense of set inclusion.

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*

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

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 $\alpha = \gamma^b$.

Proof. See appendix. □

The inequality $\alpha_2 < \alpha_1 + \gamma_2^b$ of Theorem 3 is the formal expression of what it means for X_2 to have sufficiently strong incentive power relative to X_1. Notice if γ_2^b is positive then X_2 does not need to have stronger incentive power than X_1 to cause the optimal contract to collapse. Similarly, $\gamma_2^b < \min\{\gamma_1^b, \alpha_2\}$ is the formal expression of what it means for X_2 to have sufficiently weak statistical power relative to X_1. 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 of low quality ($\alpha_2 - \gamma_2^b > 0$) and much more common than the *Bad* signal of the original information ($\gamma_2^b < \gamma_1^b$). The better monitoring/worse outcome result described in the previous section now falls out as a corollary.

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 $\gamma_2^b = 0 < \min\{1, 0.5\} = \min\{\gamma_1^b, \alpha_2\}$. □

4.3 Noisy Information Does Not Mean Weak Incentives

When it comes to the provision of incentives it is tempting to think that if information is very noisy then incentives cannot be very strong. This confounding of noisy information and weak incentives likely arises due to the fact that one often works with parametric families of information structures within which the rankings based on any reasonable measures of noisiness of information and weakness of incentives 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, the set of payoffs whose likelihood of occurring decreases when the agent exerts effort is $\{X \leq 0.5\}$. Thus, incentive power can be measured by the difference $\mathbf{P}(X \leq 0.5|a = 0) - \mathbf{P}(X \leq 0.5|a = 1)$, and statistical power can be measured by $\mathbf{P}(X \leq 0.5|a = 1)$. The quality of *Bad* signals can then be measured by the likelihood ratio $\mathbf{P}(X \leq 0.5|a = 0)/\mathbf{P}(X \leq 0.5|a = 1)$. As σ increases, incentive power, statistical power, and the quality of *Bad* signals all go down. These weakenings of informativeness and incentives can all be traced to the fact that Blackwell informativeness goes down.

However Blackwell informativeness is not a complete order and, depending on the situation, there may be other reasonable ways to compare noisiness 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 shirking, statistical power is a natural way to compare informativeness. And it is natural to think of a bad news Poisson increment that rarely generates a type II error as being less noisy than a Brownian increment that generates a type II error about half of the time. 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 emphasizes the distinction between noisy information and weak incentives. Theorem 3 and Corollary 2 highlight how signals that are at once very noisy but still very sensitive to effort exist, are not unusual, and can have an outsized effect on optimal contracting.

This 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 possibility of managerial abuse upon which the better monitoring/worse outcome result 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. For jobs 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 data generated by the monitoring technology, in which case, the noisy data may still trigger abusive reporting behavior and lead to a worse contracting outcome.

A more plausible way to overcome the better monitoring/worse outcome result 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 intelligently censor the noisy input and turn it into information that represent better quality signals of performance.[14]

In general, the best way to censor an influx of performance data will depend on the nature of the information being gathered, as well as the goals of the organization beyond just incentive provision. This represents an interesting avenue for future research as new monitoring technologies continue to come on the market. In the last part of this paper, I explore a particularly simple and widely applicable way to censor the information generated by monitoring by restricting *when* the principal is allowed to monitor.

5 Periodic Performance Appraisal

I now consider a canonical setting in which there is a stochastic process tracking *cumulative* productivity and P monitors by *sampling* that process. I will 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

[14]For example, the Oakland Police Department has recently adopted the *Vision* monitoring technology. *Vision* collects data from police stops, citizen complaints, body camera footage, and the use of force. It then flags potential outlier officers, at which point a supervisor can come in for a closer examination. (Cassidy, 2019)

frequency at which P observes new information. In the discussion above, I am suggesting that it is the first effect that is the source of potential benefits from monitoring infrequently: Reducing sampling 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 public 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 the principal and agent play the credible threats equilibrium.

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 $t_2 - \Delta$, which 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 1 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, batching usually strictly decreases surplus. 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 equilibrium 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 Frequency – Quantity vs Quality of Incentives

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: $t = 0, \Delta, \ldots T$. 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:

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

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 public history dependent $\mathcal{M}$, w, and τ also specifies a public history dependent 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 equilibrium admits a natural generalization by equating sampling periods with dates in the original dynamic model. 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. *In general, the optimal contract features a deterministic sequence of sampling times. As $\Delta \to 0$ and $T \to \infty$, there exists a D^* such that 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 a deterministic effort sequence $a^*(D)$ of length D, a real number $\rho^*(D)$ and a probability $p^*(D)$ such that the optimal contract in the continuous time infinite horizon limit given sampling frequency $\frac{1}{D}$ has the following structure:*

- *A exerts effort sequence $a^*(D)$ in every sampling period,*
- *For each $k \in \mathbb{Z}^+$, $\mathcal{M}_{kD} = \{pass, fail\}$,*
- *$m_{kD} = fail$ iff $Z_{kD} - Z_{(k-1)D} \leq \rho^*(D)$,*
- *If $m_{kD} = pass$ then A is retained for the sampling period $[kD, (k+1)D)$,*
- *If $m_{kD} = fail$ then A is terminated at date kD with probability $p^*(D)$,*
    - *If A is not terminated then it is as if P reported pass.*

The effort sequence $a^*(D)$ is completely pinned down by an initial effort level $a_0^*(D)$ and the first-order condition

$$e^{-rt}h'(a_t^*(D)) = h'(a_0^*(D)) \text{ for all } t \in (0, D).$$

The credible threats equilibrium 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 an endogenous to-be-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 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 product of incentive power and statistical power to measure the quality of *Bad* signals. As sampling frequency increases to infinity, the quality of *Bad* signals worsens to zero. Theorem 2 can be easily adapted to then show that in the continuous sampling limit providing incentives by punishing *Bad* signals is so inefficient that the optimal contract collapses. As sampling frequency decreases, the quality of *Bad* signals improves. If there were no discounting, P's flow payoff would increase to the first-best level as sampling frequency approaches zero. However, since there is discounting, as the length of a sampling period increases the quantity of incentives that can be provided begins to erode: In the beginning of a long sampling period, the threat of termination in the distant future when the period concludes has little effect on the continuation payoff of A today if A discounts. Thus, as sampling frequency goes to zero, the optimal contract again induces zero effort. Optimal sampling frequency, which is between zero and infinity, is determined by balancing the quantity versus the quality of incentive provision.

One concern about maintaining a policy of infrequently sampling 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 she would want to do so. An unravelling argument then implies that the contracting parties would 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 Some Final Thoughts About Monitoring Design

"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." (Sloan Management Review, 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." (Harvard Business Review, Cappelli and Tavis, 2016)

This paper has focused on how non-contractible information is used to *assess past performance* and *motivate hidden effort*. The key finding is that more informative monitoring can lead to *more distorted reporting* of past performance and result in an optimal contract that induces *less effort*. In the introduction, I explained my motivation for looking at this issue is driven by the rapidly expanding scope of monitoring in the workplace. A major reason for this expansion has been the need for greater worker development, as the above two quotes attest. In an increasingly dynamic world where, to quote Schrage et al (2019), "the half-life of skills is going to get shorter," monitoring will be increasingly focused on providing enough information about performance to enable managers to be effective coaches.

This is not to say that moral hazard will become unimportant. But it does suggest that the design of an organization's monitoring system will increasingly emphasize factors beyond just motivating effort. At the margin, optimal monitoring design may now involve significant trade-offs between moral hazard concerns and other concerns such as worker development.

For example, consider the issue of how frequently to appraise worker performance. The need for greater worker development has driven numerous organizations to implement continuous performance management (CPM) practices. As I mentioned in the introduction, while 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, there are concerns that all this extra information, that must be filtered through a manager's discretionary reports, might be abused and end up hurting incentives. The results derived in this section speak to those concerns and clarify why optimally solving a moral hazard problem may conflict with worker development by calling for less frequent performance appraisals that can mitigate the principal's abusive reporting tendencies.

An interesting and important monitoring design problem now arises, where the competing demands of worker development and incentive provision tug in opposite directions. How firms resolve conflicts like this one can have far reaching ramifications for organizational structure and human capital. Important considerations include what kinds of information should be collected, to what extent can feedback for worker development and incentive provision be disentangled, and how might job design interplay with monitoring design?[15]

Much work remains to be done to answer these questions. The present paper contributes by highlighting a cost that non-contractible information can impose on incentive provision that will be useful to know when tackling the more ambitious optimal monitoring design problem sketched above. I end by discussing one way in which my work can shed light on that problem.

Given that CPM is useful for worker development but potentially problematic for incentive provision, one potential strategy for resolving the conflict is to complement a CPM component meant for worker development 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).

Of course, an obvious concern about such a hybrid arrangement is whether or not the two modes of performance appraisals will interfere with each other. My analysis of batching in Section 5.1 highlights a potential pitfall of trying to provide incentives periodically while monitoring continuously. If the periodic appraisal of past performance that determines worker incentives is simply a review of past continuous appraisals used for worker development then the periodic appraisal can be counterproductive for incentive provision – the principal's abusive tendency will lead her to overemphasize performance near the end of a performance appraisal period at the expense of recognizing earlier performance.

[15] One emerging way job design can relieve some pressure on the monitoring system needing to address both worker development and incentive provision demands is through internal platform-based talent markets. With jobs becoming increasingly dynamic and multifaceted, rather than have one agent continue to perform an ever-changing job and therefore requiring constant development and monitoring, organizations are exploring how to split jobs into discrete tasks and then matching the tasks with different employees who have the skills to perform them through an internal market clearing mechanism. (Smet et al, 2016).

6 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

$$\begin{aligned}
a_t(H^A_{t-\Delta}, h_{t-\Delta}) &= a^*_t(h_{t-\Delta}) \\
m_t(H^P_t, h_{t-\Delta}) &\in \mathcal{M}^*_t(h_{t-\Delta}) \\
\mathbf{E}_{\xi_t} W^*_{t+\Delta}(h_{t-\Delta} m_t(H^P_t, h_{t-\Delta})\xi_t) &= W^{pass}_{t+\Delta}(h_{t-\Delta}) \text{ if } X_t \in Good \\
\mathbf{E}_{\xi_t} W^*_{t+\Delta}(h_{t-\Delta} m_t(H^P_t, h_{t-\Delta})\xi_t) &= W^{fail}_{t+\Delta}(h_{t-\Delta}) \text{ if } X_t \in Bad
\end{aligned}$$

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 $h_{t-\Delta}$, $H^P_{t-\Delta}$, and $x_1, x_2 \in Im(X_t)$, $m_t(H^P_{t-\Delta}x_1, h_{t-\Delta}) = m_t(H^P_{t-\Delta}x_2, h_{t-\Delta})$ 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 a Pareto-optimal contract game $\mathscr{C}$ in the T-model with the following properties: For all $t < T - \Delta$, there exist numbers $p^*_{t+\Delta}$ such that

- If $m_t = pass$ then A is retained for date $t + \Delta$,
- 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*,

Let $(V^{\mathscr{C}}_0, W^{\mathscr{C}}_0)$ denote the ex-ante payoff of $\mathscr{C}$.

The existence of such a Pareto-optimal contract game is trivially true when $T = \Delta$. For general T, let S_T^* denote the surplus generated by any Pareto-optimal contract in the T-model.

Now fix a Pareto-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 $\mathscr{C}$. Next, change the portion of $(\mathcal{M}, w, \tau)$ following the date 0 $fail$ report to a randomization between $\mathscr{C}$ and termination with final payment $w = -V_0^{\mathscr{C}}$ where the randomization is based on ξ_0 and is structured so that A's expected date Δ payoff equals $W_0^{\mathscr{C}} - (W_\Delta^{pass} - W_\Delta^{fail})$. Finally shift $w_0\Delta$ so that A's ex-ante payoff remains unchanged. 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 a Pareto-optimal contract game in the $T - \Delta$. Moreover, it has the same structure that I am trying to prove by induction. This completes the induction.

Fix a Pareto-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 a Pareto-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 surplus. This implies the stationary structure described in the Theorem.

□

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 $\alpha \leq 1$. Finally, if these conditions hold, then by making u a sufficiently steep function of a_t, one can ensure that $\lim_{\Delta \to 0} a_t^*(\Delta)$ is arbitrarily close to 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: $\gamma_1^g - \gamma_2^b \le 0 < \alpha_1 - \alpha_2 < \gamma_1^b - \gamma_2^g$.

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

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

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

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

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

Proof of Theorem 3. First consider the case $\alpha_2 \le \alpha_1$. Then we are in the situation characterized by Lemma 2. Since by assumption $\alpha_1 = \gamma_1^b$, we can only be in Case 1. In this case, $\alpha > \gamma^b$ becomes equivalent to $\alpha_2 > \min\{\gamma_1^b, \gamma_2^b\}$. But since $\gamma_1^b = \alpha_1 \ge \alpha_2$. This condition becomes equivalent to $\gamma_2^b < \min\{\gamma_1^b, \alpha_2\}$. Notice, if we reverse the inequality $\gamma_2^b < \min\{\alpha_2, \gamma_1^b\}$, then it must be that $\gamma_2^b = \alpha_2 > \gamma_1^b = \alpha_1$ which violates the assumption $\alpha_2 \le \alpha_1$.

To apply Lemma 2 to the other case, it suffices to reverse the roles of X_1 and X_2. So assume $\gamma_2^b = \alpha_2 \le \alpha_1$. This implies $\gamma_2^g = 0$. Now we can only be in Cases 1 or 3. But Case 3 implies $\alpha = \gamma^b$. We are in Case 1 if and only if $\alpha_1 < \gamma_1^b + \alpha_2$. But recall, the roles of X_1 and X_2 have been reversed. So, really, the condition is $\alpha_2 < \alpha_1 + \gamma_2^b$ which is the first inequality of the theorem. If this inequality is reversed, then with the roles reversed, it becomes $\alpha_1 > \gamma_1^b + \alpha_2$ which takes us to Case 3. This establishes one part of the tightness claim. Finally, in Case 1, $\alpha > \gamma^b$ becomes equivalent to $\alpha_2 > \min\{\gamma_1^b, \gamma_2^b\}$. But since $\alpha_2 = \gamma_2^b$, this is equivalent to $\gamma_1^b < \gamma_2^b$. Now remembering the roles are reversed, the actual condition is $\gamma_2^b < \gamma_1^b$. Finally, remember we are in the case where $\alpha_2 \ge \alpha_1 = \gamma_1^b$. So $\gamma_2^b < \gamma_1^b$ can be written as $\gamma_2^b < \min\{\gamma_1^b, \alpha_2\}$. If this inequality is reversed, then $\gamma_2^b > \gamma_1^b$. Under the reversed roles, this becomes $\gamma_1^b > \gamma_2^b$. Now Lemma 2 implies $\alpha = \gamma$. This establishes the second part of the tightness claim.

Putting the two cases together, the (tight) conditions we have established for $\alpha > \gamma$ are the ones claimed in the Theorem. This completes the proof. $\square$

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* $f \in \mathcal{F}_W$ *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

$$\mathbf{E} W 1_{\rho'}(N(g(a'), \sigma^2) = \mathbf{E} f(N(g(a'), \sigma^2)).$$

By construction, we have $a < a' \Rightarrow \mathbf{E} W 1_{\rho'}(N(g(a), \sigma^2) \leq \mathbf{E} f(N(g(a), \sigma^2))$ and $a > a' \Rightarrow \mathbf{E} W 1_{\rho'}(N(g(a), \sigma^2) \geq \mathbf{E} f(N(g(a), \sigma^2))$. 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. $\square$

Definition. *Fix a $D > 0$ divisible by Δ and an $a \in [0, 1)$. For $t = 0, \Delta, 2\Delta \dots$ 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 depend on the public history and are therefore 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

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

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

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

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

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

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

$$W^*_{e_k+\Delta}(h_{e_k}) = \sum_{t=e_k+\Delta}^{e_{k+1}} (w_t(h_{e_k}) - h(a^*(h_{e_k})(t-e_k-\Delta)))\Delta + e^{-r(e_{k+1}-e_k)}\Big[W^{pass}_{e_{k+1}+\Delta}(h_{e_k}) - \mathbf{P}\Big(N(a^{e_{k+1}-e_k}, e_{k+1}-e_k) \leq \rho(h_{e_k})\Big) \cdot (W^{pass}_{e_{k+1}+\Delta}(h_{e_k}) - W^{fail}_{e_{k+1}+\Delta}(h_{e_k}))\Big]$$

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

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

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

for all t, $h_{t-\Delta}$ and sampling dates e_k with corresponding histories h_{e_k}. 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 $Z_{e_{k+1}} - Z_{e_k} > \rho(h_{e_k})$.

Suppose by induction Lemma 3 is true for all models of length $\leq T$. Consider the $T+\Delta$-model. Fix a contract game $(\mathcal{M}, w, \tau, e)$. Change the contract game as follows: Replace the subgames following histories of the form $pass\ \xi_{e_1}$ with a common Pareto-optimal contract of the $T + \Delta - (e_1 + \Delta)$-model with some ex-ante payoff (W', V'). Replace the subgame following $fail$ with the randomization between the one following $pass\ \xi_{e_1}$ and termination with $w_{e_1+\Delta} = -V'$ calibrated so that A's expected payoff is $W' - (W^{pass}_{e_1+\Delta} - W^{fail}_{e_1+\Delta})$. 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 the unchanged. 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 $\Delta \to 0$.

$\square$

Proof of Theorem 5. By Lemma 4 and the characterization of credible threats equilibria, given fixed sampling frequency $\frac{1}{D}$, 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, $\mathbf{P}(Z_{kD} - Z_{(k-1)D} \leq \rho^*(D) \mid a^*(D)) = \Theta(1)$ as $D \to 0$. This implies that the quality of Bad signals goes to zero as $D \to 0$. Now by the argument used in the proof of Theorem 2 the effort induced by the optimal contract goes to zero as $D \to 0$. The case when $D \to \infty$ is obvious. $\square$

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] Cassidy, M. (2019, November 23) "Oakland to Use Technology that can Flag Potentially Bad Police Officers," Retrieved from `https://www.sfchronicle.com/crime/article/Oakland-to-use-technology-that-can-flag-14856949.php`

[10] Chassang, S. and G. Padró i Miquel (2019) "Crime, Intimidation, and Whistleblowing: A Theory of Inference from Unverifiable Reports," *Review of Economic Studies* Vol. 86, pp. 2530-2553

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

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

[13] 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

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

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

[16] Gigler, F. and T. Hemmer (1998) "On the Frequency, Quality, and Informational Role of Mandatory Financial Reports," *Journal of Accounting Research* Vol. 36, pp. 117-147

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

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

[19] 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`

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

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

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

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

[24] Lizzeri, A., M. A. Meyer, and N. Persico (2002) "The Incentive Effects of Interim Performance Evaluations," *working paper*

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

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

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

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

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

[30] 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/`

[31] Smet, A. D., S. Lund, and W. Schaninger (2016 January) "Organizing for the Future," Retrieved from https://www.mckinsey.com/business-functions/organization/our-insights/organizing-for-the-future

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

[33] 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

[34] Williamson, O. E. (1979) "Transaction-Cost Economics: The Governance of Contractual Relations," *Journal of Law and Economics* Vol. 22, pp. 233-261

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

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