Better Monitoring ... Worse Productivity?

John Y. Zhu[1]

April 15, 2019

Abstract

The *incentive power* and *statistical power* of information are defined. In moral hazard models with private monitoring improving the monitoring technology by bringing in new information with strong incentive power but weak statistical power can significantly reduce productivity and surplus. When monitoring is sampling a stochastic process tracking cumulative productivity, reducing sampling frequency – while information destroying – can reduce an imbalance between incentive and statistical power leading to better outcomes. These results, that arise because of the monitor's ability to abuse private information with strong incentive power but weak statistical power, provide an argument for a certain level of worker privacy and an arm's-length management style.

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

[1]I thank Dilip Abreu, Davide Cianciaruso, Tommaso Denti, Bill Sandholm, David Pearce and seminar participants at UC Boulder, Cornell, University of Rochester, University of Minnesota, University of Washington at St Louis, INSEAD, HEC, University of Maryland, UC Riverside, University of Wisconsin, Financial Research Association Conference, Richmond Fed, University of Kentucky, Michigan State.

1 Introduction

Imagine you monitor a worker, generating information you then use to reward/punish the worker and induce costly effort. You are given the chance to upgrade your monitoring technology to one with greater incentive power – meaning given any target effort level to be induced, the information generated by the new monitoring technology requires the worker to be exposed to less variation in utility. Would you upgrade?

In general, the answer is not necessarily. And the reason, valid across a broad range of games, is because what matters for surplus is a combination of the incentive power of information and the statistical power of information – a measure of power based on viewing the information generated by monitoring as a hypothesis test. Increasing incentive power at the expense of statistical power can often be counterproductive.

To fix ideas, consider the following toy model: An agent has two hidden action choices $e \in \{0, 1\}$ labelled by their costs to the agent – think of them as shirking and effort. e affects the distribution of a binary signal – g or b – where q_e is the probability of b given e and $q_0 > q_1$. There are no monetary transfers but the principal can choose a punishment p (representing the amount of agent utility destroyed) that is inflicted on the agent whenever b is realized. Clearly, as $q_0 - q_1$ increases the size of p needed to induce effort decreases. Thus, $q_0 - q_1$ measures the *incentive power* of information. At the same time, the signal g or b can also be viewed as a hypothesis test where the null is shirking and the alternative is effort. The *statistical power* of a hypothesis test is one minus the probability of a type II error – in this case q_1. The smaller is q_1 the greater the statistical power of information.

Notice inducing effort requires an expected destruction in utility $q_1 \cdot p$. This implies, when effort is induced optimally, the surplus is a monotonic function of the likelihood ratio $\frac{q_0}{q_1}$ or equivalently $\frac{q_0}{q_1} - 1 = \frac{q_0 - q_1}{q_1}$ which is the "product" of incentive power and statistical power. If an increase in incentive power is accompanied by an increase in statistical power, then surplus unambiguously increases. But if an increase in incentive power is accompanied by a sufficiently strong decline in statistical power, then surplus will decrease. If the imbalance between incentive and statistical power is severe enough, it may not even be efficient to induce effort anymore, in which case, despite the greater incentive power, productivity declines as well.

This idea that increasing incentive power can be counterproductive if accompanied by a significant reduction in statistical power is relevant across a broad range of settings including various games featuring public monitoring. The present paper explores the implications for moral hazard under private monitoring. When monitoring is private, a commitment problem on the part of the principal/monitor to truthfully report her private information arises, generating very sharp examples of how increasing incentive power can nevertheless reduce surplus and productivity.

In particular, I show that surplus and productivity can decline significantly even if the total amount of information – and therefore the incentive power of information – generated by monitoring increases. Conventional wisdom suggests inundating a monitor with lots of noisy information about worker performance is counterproductive. This result provides an explanation for why that is the case when the monitor has discretion in the way information is used to provide incentives.

To demonstrate the intuition, I start with a bad news Poisson monitoring technology under which the optimal dynamic contract induces positive effort. Depending on how the rest of the model is parameterized the induced effort under the optimal contract can be made arbitrarily high. I then increase the total amount of information generated by monitoring by adding a conditionally independent Brownian signal of effort. I show that under the improved monitoring technology the optimal dynamic contract collapses, inducing zero effort at all times.

Relative to the bad news Poisson monitoring technology already in place, Brownian information has very strong incentive power but very weak statistical power. Adding such imbalanced information to a well functioning monitoring technology causes the information generated by the improved monitoring technology to also be imbalanced, leading to a negative effect similar to the one described in the toy model analysis above. Of course, if the principal could just commit to ignore the Brownian information, then surplus and productivity would not decline much less collapse. But such a commitment is tantamount to committing to a way to report unverifiable information which is not credible. Although ex-ante surplus would increase under such a commitment, ex-interim the principal is strictly better off exposing the agent's payoff to Brownian information's very strong incentive power because this extracts more effort from the agent. Anticipating this hold-up on the part of the principal, the agent ex-ante demands a contract that restricts (in this case completely) the principal's ability to punish, hindering efficient incentive provision. Consequently, despite the better monitoring, surplus and productivity decline.

Conversely, surplus and productivity can increase despite a reduction in the amount of information generated by monitoring if an imbalance between incentive power and statistical power is also reduced. As an example, I consider a canonical model in which the agent's hidden efforts control the drift of a Brownian motion and monitoring consists of the principal privately sampling the Brownian motion. Brownian information has an extreme imbalance between incentive and statistical power. Consequently, if the principal sampled continuously, the optimal contract would collapse. The imbalance is substantially reduced and surplus and productivity increased by reducing sampling frequency even though this leads to strictly less information generated by monitoring.

The optimality of infrequent sampling suggests that it can often be better to employ an arm's-length management style where a worker is evaluated every once in a while based on his overall performance rather than all the time based on his detailed day-to-day activities.

Worth noting here is that what the literature calls infrequent monitoring – continuously sampling but only occasionally observing the results of those samplings – does not improve surplus or productivity and is in fact counterproductive in my model.

Infrequent monitoring does not change the type of information observed by the principal, just the timing of observation. If the information was imbalanced before, it is still imbalanced after making monitoring infrequent. What makes infrequent sampling useful is that it changes the very nature of the private information observed by the monitor. This contrasts with well-known results about the benefits of infrequent monitoring when monitoring is purely public. See, for example, Abreu, Milgrom and Pearce (1991).

1.1 Related Literature

This paper studies how surplus and productivity change with respect to changes in the monitoring technology in a moral hazard setting with optimal contracting. To put this work in context, one can imagine a typical contracting paper as proceeding in two steps. In the first step, the monitoring technology that will generate contract-relevant information is determined. Then, in the second step, the contracting technology (i.e. what can be conditioned on the information generated by monitoring) is determined and the optimal contracting problem is solved. Much of the literature is focused on the second step, exploring how various contracting technologies can help solve an array of agency problems or explain certain real life contracts.

For example, papers have asked how should contracts split cash flow, leading to a theory of capital structure and security design (e.g. Townsend, 1979); or how can contracts overcome constraints on completeness, shedding light on the allocation of control rights (e.g. Aghion and Bolton, 1992); or how should contracts set pay-to-performance sensitivity, leading naturally to a theory of executive compensation. See, for example, Sannikov (2008), Edmans et al (2012), or Zhu (2013, 2018a).

In contrast, for the first step, the monitoring technology is usually exogenously fixed. This approach may be fine in situations where contracts are narrowly defined only over objectively measurable performance measures like stock price. However, in many real life contractual relationships the monitoring technology itself is, at least partially, a choice variable, and the optimal contracting problem should include a discussion of monitoring design (if not a full-blown optimal monitoring design problem). Questions regarding the frequency of performance evaluations, the degree of organizational transparency, and the use of monitoring software are all issues of monitoring design and are relevant to how a firm optimally contracts with its workers. My paper with its insights about better monitoring/worse outcome and the benefits of infrequent sampling sheds light on these issues. More broadly, I take a step toward developing a theory of monitoring design by highlighting how the interplay between the incentive and statistical power of information is important to understanding the equilibrium information content of monitoring in settings where private forms of monitoring such as subjective evaluation are important.

Two other recent papers – Georgiadis and Szentes (2018) and Li and Yang (2018) – also explore monitoring design in a principal-agent model, albeit from a different perspective emphasizing information costs. In those papers, better monitoring always leads to a weakly better outcome and the information content of monitoring is determined by balancing the benefits of information against the costs of acquiring it. Also, in those papers, complete optimal monitoring design problems are set up. In contrast, my paper performs more of a comparative statics exercise, except in the last section where I solve for the optimal sampling frequency in a setting where monitoring is sampling a stochastic process tracking cumulative productivity.

Another way to position my better monitoring/worse outcome result, which is established when monitoring is private, is to look at related results in the public monitoring sphere. In an optimal contracting model with public monitoring, Holmstrom (1979) shows that adding new information that makes monitoring more informative of effort generically improves the optimal contract. In a repeated games setting with public monitoring, Kandori (1992) shows that making monitoring more informative in the sense of Blackwell (1950) causes the pure-strategy sequential equilibrium payoff set to expand in the sense of set inclusion. Both of these results are better monitoring/better outcome type results.

My paper explores how giving the principal too much information can be counterproductive. Various related literatures, including those on intrinsic motivation, mediation, and career concerns have also explored from different angles how giving the principal and/or the agent(s) too much information can be counterproductive. See, for example, Cremer (1995), Aghion and Tirole (1997), Burkart, Gromb and Panunzi (1997), Holmstrom (1999), Benabou and Tirole (2003), and Prat (2005). See also Hirshleifer (1971). My contribution to this literature is to take an information design perspective and highlight the role played by information that is strong in incentive power but weak in statistical power.

My work is also part of the literature looking at optimal contracting under private monitoring. See, for example, Levin (2003), MacLeod (2003), and Fuchs (2007). In those papers better monitoring always leads to a weakly better outcome. The technical reason why better monitoring/worse outcome does not appear in those papers is because in those papers incentive compatibility means sequential equilibrium whereas in my paper I use a refinement of sequential equilibrium to define incentive compatibility. In a companion paper, Zhu (2018b), I argue that in the private monitoring setting many sequential equilibria allow for a type of commitment behavior on the part of the principal that is implausible. I then develop the refinement used in the current paper that removes those implausible sequential equilibria.

2 Incentive Power and Statistical Power

Summary: The incentive power and statistical power of information are defined for binary action models. The definitions are then generalized for models with a continuum of actions. To achieve the generalization, a restriction needs to be imposed on how actions affect information. The restriction is satisfied if, for example, a common

MLRP condition holds.

There are two actions: $a \in \{0, 1\}$ with costs $h(1) > h(0)$. A piece of *information* is defined to be a random variable X taking values in some measurable space $Im(X)$ controlled by a – that is, the distribution of X is determined by a. Let f denote the Radon-Nikodym derivative of the measure induced by $a = 1$ with respect to the measure induced by $a = 0$. I assume almost surely (a.s.) every signal is informative: $f \neq 1$ a.s.

A reward function $R(X)$ is a measurable mapping from X to $\mathbb{R}$. $R(X)$ induces 1 if

$$1 \in \arg\max_a \mathbf{E}_a R(X) - h(a)$$

Define $|R(X)|$ to be $\sup_{x \in Im(X)} R(x) - \inf_{x \in Im(X)} R(x)$. Consider the set of non-negative reals D such that there exists an $R(X)$ inducing 1 with $|R(X)| = D$. Define $D_X(1)$ to be the infimum of this set.

Definition. *Given two pieces of information X and Y, Y is said to have weakly more incentive power than X if $D_X(1) \geq D_Y(1)$.*

Lemma 1. *There exist reward functions $R(X)$ that induce 1 and $|R(X)| = D_X(1)$. For any such $R(X)$, there exists a real w such that $R(X) = R_w(X)$ a.s. where*

$$R_w(x) = \begin{cases} w & \text{if } f(x) > 1 \\ w - D_X(1) & \text{if } f(x) \leq 1 \end{cases}$$

The proof is obvious and now one can naturally view X as a hypothesis test and define statistical power $\Pi_X(1) := \mathbf{P}(f \leq 1 | a = 1)$.

Definition. *Given two pieces of information X and Y, Y is said to have weakly more statistical power than X if $\Pi_X(1) \geq \Pi_Y(1)$.*

Comment: When the action space is binary, it is very simple to rank pieces of information by incentive power and statistical power. Almost no restrictions are placed on how actions affect information and the rankings based on incentive and statistical power are complete. The rest of the paper looks at moral hazard models in which the agent has a continuum of effort choices. Thus, I need to generalize the ways pieces of information are ranked based on incentive and statistical power. The generalized rankings are no longer always complete (although within some special parametric famliies they are) and I will also need to impose further restrictions on how actions affect information so that a version of Lemma 1 holds which is required to have a meaningful definition of statistical power.

A piece of information X is now smoothly controlled by an effort parameter $a \in [0, 1)$

with smooth, strictly convex cost function $h(a)$ where $h'(0) = 0$, and $\lim_{a\to 1} h'(a) = \infty$. Smoothly controlled means that for any measurable set $A \subset Im(X)$, $\mathbf{P}(X \subset A)$ is a smooth function of a. For each $D \geq 0$, consider the set of a such that there exists an $R(X)$ inducing a with $|R(X)| \leq D$. Let $a_X(D)$ denote the supremum of this set.

Definition. *Given two pieces of information X and Y, Y is said to have weakly more incentive power than X if $a_Y(D) \geq a_X(D)$ for all $D \geq 0$.*

A piece of information X *can be viewed as a hypothesis test* if the following condition holds: For every $D > 0$ there exist $R(X)$ inducing $a_X(D)$ with $|R(X)| \leq D$, and any such $R(X)$ takes only two values a.s.

Some commonly used pieces of information can be viewed as hypothesis tests:

Example. *If $Im(X) = \mathbb{R}$, X satisfies strict MLRP with respect to a, and the distribution given each a has a density, then X can be viewed as a hypothesis test. If X is strictly monotone in the sense that $Im(X) = Good \cup Bad$ with $\mathbf{P}(X = x)$ strictly increasing (decreasing) and continuous in a if $x \in Good$ (Bad), then X can be viewed as a hypothesis test.*

From now on I restrict attention to information that can be viewed as hypothesis tests. For each $D > 0$, consider the set of $R(X)$ inducing $a_X(D)$ satisfying $|R(X)| \leq D$. By assumption, each such $R(X)$ is binary valued and can be associated with a probability $\mathbf{P}(R(X) = \min_{x\in Im(X)} R(x)|a_X(D))$. Define $\Pi_X(D)$ to be the infimum of this set of probabilities.

Example. *Suppose X is strictly monotone. Fix a D and suppose $R(X)$ induces $a_X(D)$ with $|R(X)| \leq D$. Then $R(X) = \inf_{x\in Im(X)} R(x)$ ($= \sup_{x\in Im(X)} R(x)$) if and only if $x \in Bad$ ($\in Good$). Thus, $\Pi_X(D) = \mathbf{P}(X \in Bad|a_X(D))$.*

Definition. *Given two pieces of information X and Y, Y is said to have more statistical power than X if for each D there exists a $\hat{D}$ such that $a_X(D) = a_Y(\hat{D})$ and $\Pi_X(D) \geq \Pi_Y(\hat{D})$.*

3 Model and Optimal Contract

I consider a dynamic contracting model between a principal P (she) and an agent A (he). The horizon is infinite and dates are of length $\Delta > 0$, denoted by $t = 0, \Delta, 2\Delta, \ldots$. The discount factor is $e^{-r\Delta}$ for some $r > 0$.

At the beginning of each date t, P pays A some amount $w_t \in \mathbb{R}$. 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 *monitors* A: First, P observes a private signal X_t smoothly controlled by a_t. I assume X_t is strictly monotone. Given a_t, P's unobserved utility is $u(a_t)\Delta$. I assume $u(a_t)$ is a strictly increasing, weakly concave function of effort and $u(0) > 0$. Next, P reports a public message m_t selected from a contractually pre-specified finite set of messages $\mathcal{M}$; then, a public randomizing device is realized; finally, A is randomly 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.

A *contract game* $(\mathcal{M}, w, \tau)$ specifies a finite message space $\mathcal{M}$, a payment plan w, and a termination clause τ. Let h_t denote the public history of messages and public randomizing devices up to the end of date t. w consists of an $h_{t-\Delta}$-measurable payment w_t to the agent for each t. τ is a stopping time where $\tau = t+\Delta$ is measurable with respect to h_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-1} of prior effort choices. m consists of a message choice m_t for each t depending on $h_{t-\Delta}$ and P's private history H^P_t of observations $\{X_s\}_{s\leq t}$. The system of beliefs consists of a belief about $H^P_{t-\Delta}$ at each decision node $(H^A_{t-1}, h_{t-\Delta})$ of A, and a belief about H^A_t at each decision node $(H^P_t, h_{t-\Delta})$ of P.

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

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

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

3.1 The Optimal Contract

The optimal contracting problem is to find an incentive compatible contract that maximizes V_0 subject to the agent's ex-ante participation constraint $W_0 \geq 0$ and an interim participation constraint $W_t + V_t \geq 0$ for all t. Intuitively, if the interim participation constraint were violated then both parties could be made strictly better off by separating under some severance pay.[2]

Incentive compatibility typically means that the principal's report strategy and the agent's effort strategy comprise some sort of equilibrium behavior. A detailed discussion of what is the right equilibrium concept is the subject of a companion paper Zhu (2018b) which argues that a particular refinement of sequential equilibrium is the right equilibrium concept.

The idea is this: Imagine a game in which after player 1 moves, player 2 is indifferent between all of her actions. Dewatripont (1987) and Tranaes (1998) argue that in such games, player 2 will credibly commit to a way to choose among her many

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

best-response actions *before* player 1 moves in an effort to induce player 1 to choose an action that is more preferred by player 2 from an ex-ante perspective. Zhu (2018b) applies this idea to the private monitoring model described above. Since monitoring is private, to induce P to report truthfully at the end of date t she must be made indifferent between all of her date t reports. This implies the higher is A's date t effort, the higher is P's payoff standing at the beginning of date t. Thus, applying the logic of Dewatripont (1987) and Tranaes (1998), at the beginning of date t P commits to the date t report strategy that maximizes the date t effort induced from A. Given what we know about the incentive power of strictly monotonic information from the previous section, it is now clear that in any incentive compatible contract and at each date t the principal will report the message that leads to the highest (lowest) possible agent continuation payoff if and only if $X_t \in Good$ ($X_t \in Bad$). No other messages will be reported. This characterization of P's report strategy in any incentive compatible contract implies:

Theorem 1. *The optimal contract has the following structure:*

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

Proof. See appendix A. □

A similar result holds more generally if information can be viewed as a hypothesis test and any two $R(X)$ that induce $a_X(D)$ with $|R(X)| = D$ differ by a constant a.s. This is satisfied if for example a linearly controls the mean of a normal random variable: $X \sim N(b + ma, \sigma)$ for some constants b, m, and σ.

We already understand that P reports in a simple pass/fail way in order to take full advantage of the incentive power of information. Let us now get a feel for the other features of the optimal contract.

Notice, the optimal contract is a wage contract. At each date t, conditional on still being employed, the agent is paid the same amount regardless of performance history. There is a good reason for this. Suppose instead there was an additional message that leads to A receiving a big bonus which P is supposed to report if she observes some really positive information about A's performance (i.e. a *Good* signal whose probability increases sharply as a increases). The problem with this altered contract is that its strategy profile would not satisfy any reasonable notion of incentive compatibility: Because monitoring is private, P can always claim she didn't see the really positive information even if she did and thereby avoid having to pay A the big bonus. In general, P must be indifferent between reporting different messages that occur on the equilibrium path, which means in the optimal contract

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

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

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

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

Next, consider A's effort incentives. Since the optimal contract is a wage contract, one might wonder where are the effort incentives coming from? The answer is through the threat of termination. In my model, termination destroys surplus – by assumption, even zero effort generates positive surplus. Since P is completely insured against any surplus destruction, this means it is A who bears the cost of inefficient termination,

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

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

$$h'(a^*)\Delta = -\frac{d\mathbf{P}(Bad)}{da}|_{a=a^*} p^* S^*.$$

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

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

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

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

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

4 Better Monitoring Worse Outcome

As I discussed in the introduction, it is generally possible for an increase in the incentive power of information to lead to lower surplus and productivity if it is accompanied by a sufficient decline in statistical power. However, this result has limited scope when players can ignore information. Such is the case in standard repeated games and moral hazard models with public monitoring. In those settings if an increase in incentive power is part of an overall increase in the quantity of information generated by monitoring, then it is impossible for surplus and productivity to decline no matter how much statistical power declines.

In contrast, when monitoring is private so that the principal always reports in a way that maximizes effort incentives then better monitoring can lead to a worse outcome.

Proposition 1. *Fix a monitoring technology X satisfying $\mathbf{P}(X \subset Bad|a = 0) < 1$. Then there exist monitoring technologies Y that are strictly more informative and have strictly more incentive power than X but generate strictly lower surpluses.*

Proof. I will prove it in the case that X takes finitely many values. The proof can be easily generalized to the case when $Im(X)$ is infinite. By assumption there exists a *Good* signal x with $\mathbf{P}(X_t = x|a = 0) = \beta > 0$. Split x into two signals x^g and x^b. Have it so that $\mathbf{P}(X_t = x^g|a) = \mathbf{P}(X_t = x|a) - \beta + f(a)$ where $f(a)$ is a strictly increasing nonnegative function with $\lim_{a\to 1} f(a)$ "very close to 0" both in absolute terms and compared to β. Call this new monitoring technology Y. Y is obviously strictly more informative than X, and is strictly monotonic with x^g a *Good* signal and x^b a *Bad* signal. Y has strictly more incentive power than X due to f strictly increasing in a. Y has strictly less statistical power than X due to $\beta > 0$. Consider the effort a^*_Y induced by the optimal contract $\mathscr{C}^Y$ under Y. Since X's incentive power is close to Y, it is possible to induce an effort level a_X very close to a^*_Y under X using a contract with a pass/fail contract game similar to the contract game of $\mathscr{C}^Y$. However, in this contract the fail message is employed much less often in the sense that the surplus saved from reporting fail less outweighs the slight decline in effort induced from a^*_Y to a_X. This comes from the fact that $|f|$ is small compared to β. Thus, the Pareto-optimal surplus under X must be strictly larger than that under Y. □

The proposition while quite general does not give the reader a good sense of just how much worse things can get when monitoring is improved. Moreover, in the proof of the corollary it is hard to map the splitting of x into x^g and x^b to a real life way in which monitoring might be improved. In the remainder of this section I address both of these issues.

I now show how starting with a monitoring technology under which the optimal contract induces arbitrarily high effort it is possible to add a conditionally independent signal of effort and cause the optimal contract to collapse into a trivial arrangement that induces zero effort from the agent at all times. Note such a collapse necessarily means surplus declines because inducing zero effort at all times with a trivial contract is always an option for the contracting parties.

To get a sense of how this is possible, let us revisit Theorem 1. Recall, P fails A at date t if and only if she observes a *Bad* signal at date t. This simple report strategy is a consequence of P always wanting to maximize effort incentives. The question is: Is this the efficient thing to do? Put another way, if P were a benevolent social planner instead of a utility maximizer would she still report in this way or something close given the contract game? The answer is it depends. If the monitoring technology generates *Bad* signals that are *really bad* (of course, we will have to be formal about what "really bad" means) then intuitively the answer is yes. Where this strategy becomes inefficient is when the monitoring technology generates *Bad* signals that are for the most part only *marginally bad*. In this case, one would like to see P be a little more discriminating and fail A only if she sees a really *Bad* signal, or at least wait until she has seen marginally *Bad* signals across many dates before failing A. But we already know P is unable to be discriminating: Sure, at the time of contracting, P would like to commit to be discriminating in the future. The problem is, once the contract is written, P can't help but change her report strategy to an indiscriminate one that maximizes effort incentives by punishing A maximally any time a *Bad* signal no matter how marginal occurs. Since changing a report strategy amounts to changing a function over private, unverifiable information, it is not something that can be contracted away.

Now at the time of contracting A understands that in the future, if the monitoring technology will generate lots of marginally *Bad* signals, P will likely over-fail A. To counteract this, the contracting parties then preemptively agree to an optimal contract that reduces the pain of failure. That means setting p^* to be a low value. And in some cases when the typical *Bad* signal is extremely marginal, it might even be optimal to lower p^* all the way to zero. Of course, once p^* hits zero failing becomes equivalent to passing and A will exert zero effort.

Is it possible to take a monitoring technology that generates mostly really *Bad* signals and *improve* it to the point where it generates mostly extremely marginal *Bad* signals? Because if it is, then better monitoring can indeed lead to a complete collapse in the optimal contract.

I now show that such an improvement can be achieved in a very simple way: Introduce new information that is, relative to the old information, very strong in incentive power but very weak in statistical power. Before establishing this result at a reasonably general level, let us first work through an explicit example that demonstrates the basic ideas.

4.1 An Example

In this example I will begin with a bad news Poisson monitoring technology. "Bad news" means that the Poisson event is indicative of lower effort rather than higher effort. I show that the bad news Poisson monitoring technology generates a really *Bad* signal in some formal sense and consequently the optimal contract induces positive effort – depending on the intensity of the Poisson process this effort can be made to be arbitrarily high. I then improve the monitoring technology by adding a conditionally independent Brownian signal where the drift of the Brownian motion is controlled by A's effort. I explain that Brownian information has, relative to bad news Poisson information, very strong incentive power but very weak statistical power. I then show that in the improved monitoring technology that generates both bad news Poisson information and Brownian information, the typical *Bad* signal suddenly becomes extremely marginal. Consequently, the optimal contract collapses and P becomes worse off.

In a bad news Poisson monitoring technology, each date the incremental information X_t is

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

Here, Δ is understood to be small and it is evident that the Poisson event itself is the *Bad* signal whereas no event is the *Good* signal. Just how bad (e.g. really bad or marginally bad) is the *Bad* signal of bad news Poisson information? The formal measure is the negative effort-elasticity of $\mathbf{P}(Bad)$:

$$-\frac{d\log\mathbf{P}(Bad)}{da}. \tag{1}$$

In the subsequent general analysis I will justify why this is the natural measure for how bad is the typical *Bad* signal. For now let us take it as given. This elasticity corresponds to the likelihood ratio discussed in the introduction. Just like the likelihood ratio, it is a barometer for how imbalanced are the incentive and statistical power of information. Very strong incentive power plus very weak statistical power will mean a very small negative effort-elasticity relative to medium incentive power and statistical power. A simple computation shows that the negative effort-elasticity of $\mathbf{P}(Bad)$ is

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

What matters about this quantity is that it remains bounded away from zero as Δ becomes small no matter the effort level. This means bad news Poisson information is not too imbalanced and features a really *Bad* signal. Thus, it is not too inefficient to have a contract that lets P punish A nontrivially whenever she sees it, and the optimal contract under bad news Poisson monitoring can induce positive effort. By making λ sufficiently large, the induced optimal effort level can be made arbitrarily high. Later I generalize this result by showing that for a broad class of monitoring technologies in the continuous time limit if the negative effort-elasticity of $\mathbf{P}(Bad)$ is not vanishingly small then there exist parameterizations of the rest of the model such that the optimal contract induces arbitrarily high effort.

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

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

Each date the Brownian signal is a single step of an extremely fine random walk. Whenever the random walk goes up it is a *Good* signal, whenever it goes down it is a *Bad* signal. For Brownian information, the negative effort-elasticity of $\mathbf{P}(Bad)$ is

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

Unlike before, it is clear that this elasticity goes to zero as Δ goes to zero no matter the effort level. This means the Brownian *Bad* signal is an extremely marginal *Bad* signal, and because it is extremely marginal, it is important that P be discriminating when using Brownian information to justify punishing A. This basically means that P needs to commit to be patient and wait until she has seen many extremely marginal Brownian *Bad* signals before deciding to fail A based on Brownian information. But as I explained earlier being patient is not something that is compatible with P's desire to maximize effort incentives at all times.

But if P is unwilling to be patient when using Brownian information, then she should not be using it at all. In other words, under the improved monitoring technology (X_t, Y_t) that generates both bad news Poisson information and Brownian information, it will be best if P simply ignores Y_t.

Will P ignore Y_t? Not if Y_t has sufficiently strong incentive power. Recall, ultimately what P cares about is maximizing effort incentives and so if a new piece of information has strong incentive power – at least relative to the information already in place – then it does not matter how noisy it is, P will not ignore it.

Is it possible for a piece of information to simultaneously have very strong incentive power but feature extremely marginal *Bad* signals? As we shall see shortly, the answer is yes, and a canonical example of such information is Brownian information. We already know that Brownian information has a vanishingly small negative effort-elasticity meaning its *Bad* signal is extremely marginal. A sufficient statistic for incentive power is the sensitivity of $\mathbf{P}(Bad)$ with respect to effort:

$$-\frac{d\mathbf{P}(Bad)}{da}. \tag{2}$$

While the measure for incentive power looks like the measure of how bad is the *Bad* signal, they are not the same, and it is very easy to come up with information such that (1) is very small but (2) is very large.

Now to show that P will not ignore Y_t let us compare the incentive power of X_t and Y_t. The sensitivity of $\mathbf{P}(Bad)$ with respect to effort under X_t is

$$\lambda\Delta.$$

The sensitivity of $\mathbf{P}(Bad)$ with respect to effort under Y_t is

$$\frac{\sqrt{\Delta}}{2} \gg \lambda\Delta.$$

Thus, Y_t has much greater incentive power than X_t which means when a bad news Poisson monitoring technology is improved by adding a Brownian component, there is no way P will ignore the Brownian component.[3]

In fact, an easy application of the product rule shows that, whereas before the improvement, P would fail A whenever a bad news Poisson event occurred, after the improvement, P now fails A whenever the bad news Poisson event occurs or whenever the Brownian random walk goes down. In particular, even if the bad news Poisson event doesn't occur but the Brownian random walk goes down A is still failed. This is noteworthy, because the combination of the bad news Poisson event not occurring and the Brownian random walk going down is a combination of a *Good* bad news Poisson signal and a *Bad* Brownian signal. A priori, it may not be clear how to interpret such a combination – one could make the argument that seeing one good signal and one bad signal should constitute a neutral signal overall. But that is not the case here: The combination of the *Good* bad news Poisson signal and the *Bad* Brownian signal is unambiguously a *Bad* signal overall because its likelihood of occurring unambiguously decreases as effort increases – *not by much* – but it does decrease. And the main reason for this decrease is due to the very strong incentive power of Brownian information.

But now we have a problem: This composite *Bad* signal is quite common – occurring about half the time no matter what effort A puts in. Consequently, punishing A whenever this common, extremely marginal *Bad* signal occurs is extremely inefficient and should be avoided at all costs. But as I've said before – there is nothing one

[3] In general, the new information does not need to have weakly stronger incentive power than the old information in order for P not to ignore it. As I will show in the general analysis, the new information's incentive power just needs to be above a certain threshold that is increasing in the incentive power of the old information.

can do to avoid this inefficiency. Monitoring is private. When P fails A the whole point is one cannot tell if it is because P saw the really bad bad news Poisson event occur, in which case A "deserves" to get punished, or if the only thing P saw was the marginally bad Brownian random walk go down. The only way to imperfectly counteract P's inevitable over-failing of A is to make failing painless. That means setting $p^* = 0$.

Thus, when bad news Poisson monitoring is improved by adding a Brownian component, the optimal contract collapses into a trivial arrangement that always pays A a flat wage w_{salary} and never terminates A. A best responds by putting in zero effort, and P despite her better monitoring becomes worse off.

4.2 Continuous Monitoring

What aspects of Brownian information made adding it to a bad news Poisson monitoring technology so counterproductive?

One important property of Brownian information that emerged in the analysis is that it features a marginally *Bad* signal:

1. New information has a marginally *Bad* signal.

Intuitively, the marginally *Bad* signal of the new information can help the improved monitoring technology generate information that also features mostly marginally *Bad* signals which, recall in the intuition sketched out in the beginning of this section, is a precondition for the optimal contract to collapse.

Another important property that emerged from the analysis is that Brownian information has strong incentive power:

2. New information has sufficiently strong incentive power.

Here, the idea is even if the new information has extremely marginal *Bad* signals, if P ignores the new information, then introducing it makes no difference. To ensure P does not ignore the new information, it must have sufficiently strong incentive power since what P cares about is maximizing effort incentives.

Finally, recall the Brownian *Bad* signal is quite common, occurring about half the time no matter A's effort:

3. New information has sufficiently common *Bad* signals.

Intuitively, even if a new *Bad* signal is extremely marginal and even if P fails A based off of it, if the signal almost never occurs then the inefficiency remains small and P can afford to continue to punish A non-trivially, in which case the optimal contract does not collapse.

Properties 1. and 3. together say that the statistical power of information is sufficiently weak. Thus, a compact way to state these three properties is to say that relative to the information already in place *new information has strong incentive power but weak statistical power.*

I will now establish a general result about how improving a well functioning monitoring technology by adding new information strong in incentive power but weak in statistical power can cause the optimal contract to completely collapse.

In this analysis I consider the set of all monitoring technologies in the continuous-time limit satisfying the following regularity conditions: For all a_t,

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

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

$$X_t = \begin{cases} g & \text{with probability} \quad a_t\lambda\Delta \\ b & \text{with probability} \quad 1 - a_t\lambda\Delta \end{cases}$$

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.

Theorem 2. *Assume $\alpha \leq 1$. If $\alpha = \gamma^b$ then the model can be parameterized so that the optimal contract induces non-zero effort. Otherwise the optimal contract induces zero effort.*

Proof. See appendix. □

Recall, in the introduction I argued that an increase in incentive power can reduce surplus and productivity if accompanied by a sufficiently strong decrease in statistical power. Theorem 2 is a particularly strong manifestation of this result: Imagine one starts with a monitoring technology where $\alpha = \gamma^b$ and positive effort is induced. Now lower α (increase incentive power). If γ^b drops (statistical power decreases) sufficiently so that $\gamma^b < \alpha$ then the optimal contract induces zero effort which means both surplus and productivity strictly decline despite the greater incentive power.

Theorem 2 also justifies my use of the negative effort-elasticity of $\mathbf{P}(Bad)$ as a measure of how "bad" is the typical Bad signal. Given my regularity assumptions, the negative effort-elasticity of $\mathbf{P}(Bad) = \Theta(\Delta^{\alpha-\gamma^b})$. Thus, Theorem 2 can be reworded as saying if the negative effort-elasticity is bounded away from zero then the optimal contract induces positive effort. If it converges to zero as Δ tends to zero then the optimal contract induces zero effort.

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.*

This corollary matches classic results from the literature on repeated games with 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 work, with its emphasis on the distinction between incentive and statistical power, differs from related work 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 monitoring in the models described above always weakly improve the scope for cooperation.

Armed with Theorem 2 I can now investigate how improvements to the monitoring technology affect optimality. I begin with a binary valued monitoring technology $X_{1t} \in \{b_1, g_1\}$ with associated exponents $(\alpha_1, \gamma_1^b, \gamma_1^g)$. I then improve it by adding a conditionally independent binary valued monitoring technology $X_{2t} \in \{b_2, g_2\}$ with associated 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_{1t}, X_{2t}) generated by the improved monitoring technology. Thus, (X_{1t}, X_{2t}) 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.

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

Lemma 2. *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 3. *Given X_{1t} and X_{2t} with associated 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 associated exponents of the vector-valued* (X_{1t}, X_{2t}) *are*

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

Proposition 3 only considers the case where $\alpha_1 \geq \alpha_2$. The other case, $\alpha_2 \geq \alpha_1$, is implied by symmetry.

Proof. See appendix. □

Lemma 3 yields an explicit characterization of counterproductive improvements to the monitoring system.

Proposition 2. *Suppose* $\alpha_1 = \gamma_1^b$. *If* $\alpha_2 < \alpha_1 + \gamma_2^b$ *and* $\gamma_2^b < \min\{\gamma_1^b, \alpha_2\}$, *then* $\alpha > \gamma^b$. *The result is tight in the sense that if either of the inequalities is reversed then* $\alpha = \gamma^b$.

The two inequalities, $\alpha_2 < \alpha_1 + \gamma_2^b$ and $\gamma_2^b < \min\{\gamma_1^b, \alpha_2\}$, of Proposition 2 formalize how introducing new information that is, relative to the information already in place, sufficiently strong in incentive power but sufficiently weak in statistical power is counterproductive. The first inequality corresponds to property 2. from the beginning of this subsection. If one breaks up $\gamma_2^b < \min\{\gamma_1^b, \alpha_2\}$ into the two component inequalities $\gamma_b^2 < \alpha_2$ and $\gamma_b^2 < \gamma_b^1$ one recovers properties 1. and 3. from the beginning of this subsection.

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

5 Infrequent Sampling

The better monitoring/worse outcome result suggests that there is value to limiting the information observed by P if such a limitation can reduce an imbalance between incentive and statistical power. Such an imbalance can seriously impact surplus and productivity, for example causing the optimal contract to induce zero effort under a wide range of monitoring technologies in the continuous-time limit. How might surplus and productivity be restored by limiting information? One obvious way, if applicable, is to simply reverse the process that caused the optimal contract to collapse in the previous section. However, there are other practical ways to beneficially limit information by controlling when the principal monitors.

The repeated games literature has highlighted the benefits of *infrequent monitoring* so that, say, every 10 units of time, P observes the information generated from the past 10 units of time all at once. Does this type of infrequent monitoring help in my setting? Even though my model, as it is currently defined, does not allow the contract to control when the principal sees X_t, my analysis has already indirectly provided an answer to this question.

Releasing information in batches as suggested by infrequent monitoring is equivalent to releasing information as it is generated but restricting the players to respond to new 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 new information every once in a while: For example, the contract game could be structured so that pay and termination do not depend on any report made between t_1 and $t_2 - \Delta$. In this case, the contract game does not allow P to react to new information between t_1 and $t_2 - \Delta$ and it is equivalent to batching the information generated between t_1 and t_2 and releasing it all at once at date t_2. Since Theorem 2 is a result about optimal contracting, contract games that allow P to react to new information only every once in a while are already folded into the analysis. Thus, my optimality result indirectly implies that infrequent monitoring cannot increase surplus.

In fact, choosing a contract game that allows P to react to new information only every once in a while is not only not helpful, it is usually hurtful. Suppose the contract game does not allow P to react to new information between t_1 and $t_2 - \Delta$. On date t_2 when P finally has the opportunity to affect A's continuation payoff through her reports, all of A's efforts before date t_2 have been sunk. 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. As I explained earlier when discussing incentive compatibility, 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 $t_2 - \Delta$.

This discussion of infrequent monitoring in conjunction with the better monitoring/worse outcome result shows how the relationship between the information content of monitoring and surplus/productivity is fundamentally different in a moral hazard model with private monitoring compared to most repeated games models with purely public monitoring. In most of the repeated games literature, players have some ability to ignore or not ignore information ex-interim based on ex-ante considerations. Consequently, the following two important comparative statics results concerning the information content of monitoring emerge:

- Holding the frequency of monitoring fixed, increasing information content never hurts.
- Holding the information content of monitoring fixed, decreasing frequency often helps.

In contrast, in my setting, the inability of P to commit to any behavior that deviates from maximizing today's effort incentives every day means that the above two comparative statics get almost completely reversed:

- Holding the frequency of monitoring fixed, increasing information content often hurts.
- Holding the information content of monitoring fixed, decreasing frequency never helps.

Despite the ineffectiveness of the type of infrequent monitoring typically considered in the repeated games literature, there is another, arguably more natural, way to infrequently monitor that can help improve outcomes in my model: In many situations, the relevant stochastic information process tracks some notion of cumulative productivity and monitoring is *sampling* that process. Under this definition of monitoring, when P infrequently monitors, not only is the release of information being delayed as in the repeated games literature, but also the quantity of information generated declines unlike in the repeated games literature. To distinguish this type of infrequent monitoring from the type that is typically referred to in the repeated games literature, I will refer to this type of infrequent monitoring as *infrequent sampling.*

To explore the costs and benefits of infrequent sampling, I now consider a canonical setting where the stochastic information process is Brownian motion with the drift being controlled by effort. In this new model, the timing of events at each date t is the same as in my original model except P may or may not monitor A. If P does monitor A, she no longer observes X_t where

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

but, rather, samples the process $Y_t = \sum_{s \leq t} X_s$ tracking cumulative productivity. From now on I refer to monitoring as sampling.

In this new model, a contract game, in addition to specifying $\mathcal{M}$, w, and τ, also specifies a sequence of *sampling times* $e_1 < e_2 < \ldots$. An assessment is defined similarly to before except P's decision nodes only occur on sampling dates. Throughout the analysis I will assume an infinitesimal Δ.

In this more flexible model, if one restricts attention to contracts that have P sample every date, then the optimal contracting problem becomes identical to the original one and Theorem 2 implies that A exerts zero effort due to Brownian information's extreme imbalance between incentive power and statistical power. However, this imbalance can be mollified by sampling only every once. As a result, infrequent sampling is optimal:

Theorem 3. *There exist $\Delta^* > 0$, ρ^*, and p^* such that the optimal contract has the following structure:*

- *P samples every Δ^* units of time: $e = \{\Delta^*, 2\Delta^*, 3\Delta^* \dots\}$*
- *$\mathcal{M} = \{pass, fail\}$.*
- *For each $k \in \mathbb{Z}^+$, $m_{k\Delta^*} = fail$ iff $Y_{k\Delta^*} - Y_{(k-1)\Delta^*} \leq \rho^*$.*
- *w consists of a pair of constants $w_{salary}, w_{severance}$.*
- *For each $k \in \mathbb{Z}^+$, if $m_{k\Delta^*} = pass$ then A is retained for the sampling period $(k\Delta^*, (k+1)\Delta^*]$ and is paid a stream $w_{salary}dt$.*
- *For each $k \in \mathbb{Z}^+$, if $m_{k\Delta^*} = fail$ then A is terminated with probability p^*.*
  - *If A is not terminated then it is as if P reported pass.*
  - *If A is terminated then he is paid a lump sum $w_{severance}$.*

Most of the work in proving Theorem 3 has already been done. P report pass or fail based on whether or not the threshold ρ^* is reached is a consequence of maximizing effort incentives when information satisfies MLRP with respect to effort. The evenly spaced aspect of optimal sampling frequency is a consequence of the infinite time horizon and the fact that the continuation contract after pass is itself Pareto-optimal.

The main difference between this optimal contract and the original optimal contract is that in the original model the length of a sampling period was exogenously fixed to be Δ whereas in the new model it is endogenously determined.

The magnitude Δ^* of this endogenously determined sampling period length is pinned down by an intuitive tradeoff: The set of signal realizations $Y_{k\Delta^*} - Y_{(k-1)\Delta^*} \leq \rho^*$ is similar in spirit to the *Bad* set of signal realizations for X_t in the original model. As Δ^* shrinks, this set of "*Bad*" signals becomes increasingly marginal meaning the negative effort-elasticity of $\mathbf{P}(Bad)$ becomes vanishingly small – and consequently, insisting on failing A whenever these marginal "*Bad*" signals occur becomes increasingly inefficient. As we now understand, this type of over-failing of A is counterproductive and will cause the optimal contract to collapse. On the other hand, as Δ^* increases away from zero discounting begins eroding the incentive power of information: In the beginning of a sampling period, the threat of termination in the distant future when the sampling period concludes has little effect on the continuation payoff of A today. The optimal Δ^* balances these two opposing forces: The desire for greater balance between incentive and statistical power on the one hand versus the desire for greater incentive power in absolute terms on the other.

6 Conclusion

This paper studied how changes to the information content of private monitoring in a moral hazard setting impacts productivity and surplus. I showed that improving the information content of monitoring by introducing new information that is weak in statistical power but strong in incentive power can backfire, leading to a decline in productivity and surplus. In some cases, improvements to monitoring can cause the optimal contract to collapse into a trivial contract that induces zero effort. Delaying the monitor's ability to react to the information generated by monitoring only makes things worse. On the other hand, in settings where monitoring is sampling the current value of a fixed stochastic process tracking cumulative productivity, infrequent sampling can be beneficial. Optimal sampling is periodic with the period length determined by an intuitive tradeoff between more incentive power versus more balance between incentive and statistical power.

7 Appendix

Proof of Theorem 2. Let $a_t^*(\Delta)$ denote the effort induced by the optimal contract at date t. 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,

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

must hold. Since marginal cost $= \Theta(\Delta)$, therefore marginal benefit $= \Theta(\Delta)$. Since $e^{-r\Delta}S^*(\Delta) = \Theta(\Delta^0)$ and, by assumption, $-\frac{d}{da}\mathbf{P}(X_t \in Bad \mid a_t^*(\Delta), \Delta) = \Theta(\Delta^\alpha)$, 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 $\gamma^b + 1 - \alpha \geq 1 \Rightarrow \alpha - \gamma^b = 0$. Feasibility of $p^*(\Delta) = \Theta(\Delta^{1-\alpha})$ implies $\alpha \leq 1$. □

Proof of Proposition 3. Case 1a: $\gamma_1^g - \gamma_2^b < \alpha_1 - \alpha_2 = 0 < \gamma_1^b - \gamma_2^g$.

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

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

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

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

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

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

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

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

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

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

A Incentive Compatibility

The section generalizes the first half of Zhu (2018b). I begin with a generalization of the class of monitoring systems considered:

Assumption 1. *Let R be any non-empty, finite set of real numbers. Let ξ be any full-support finite-valued random variable whose distribution does not depend on a_t. There exists a unique function $f^R(X_t)$ taking values in R with the following two properties:*

- *The set* $\arg\max_{a_t} \mathbf{E}_{a_t} f^R(X_t) - h(a_t)\Delta$ *contains a maximal element a^R.*
- *For any function $g(X_t, \xi)$ taking values in R, if it is not true that $g(X_t, \xi) = f^R(X_t)$ for all X_t and ξ, then a^R is strictly larger than any element of* $\arg\max_{a_t} \mathbf{E}_{a_t,\xi} g(X_t, \xi) - h(a_t)\Delta$.

One can think of R as a set of possible rewards for A, g as a performance-sensitive reward function designed to induce effort from A, and ξ as noise. When Assumption 1 is used in the analysis below, R will correspond to the set of possible discounted date $t + \Delta$ continuation payoffs for A, g will be P's date t report strategy, and ξ will be P's private history leading up to date t. Assumption 1 says to maximize effort A's performance-sensitive reward cannot depend on noise.

Assumption 1 holds under many natural models of how effort affects the distribution of X_t, including the special case where effort has a strictly monotone effect on X_t. In this case, f^R takes at most two values, the maximal and minimal values of R, with f^R taking the minimal value of R if and only if $X_t \in \mathit{Bad}$.

I am now ready to discuss incentive compatibility in settings where the monitoring technology satisfies Assumption 1. I assume that the model has a terminal date $T < \infty$ unlike in the body of the paper. Once I define incentive compatibility and characterize the optimal contract, I will then show that as $T \to \infty$ the optimal contract converges to the one in Theorem 1.

In most contracting models, incentive compatibility means the assessment is a sequential equilibrium. However, in my setting, many sequential equilibria feature implausible behavior by P: Pick an arbitrary date $t < T$ and examine P's date t payoff,

$$V_t = (-w_t + u(a_t|m_t)\Delta) + e^{-r\Delta}V_{t+\Delta}. \tag{3}$$

V_t is the sum of two components – her expected date t utility as a function of A's date t effort and her discounted date $t + \Delta$ continuation payoff. It is without loss of generality to assume the second component does not depend on m_t – otherwise P would only report the messages that maximized the second component. Thus, the only thing m_t affects is A's date t effort. Looking at the first component, it is clear the higher is a_t the higher is V_t. Thus, at the beginning of date t, P should want to commit to m_t that maximizes date t effort incentives.

However, the sequential equilibrium concept allows P to use m_t that do not maximize $a_t \mid m_t$. My strategy for refining sequential equilibrium is to be conservative about using the idea of P wanting to maximize effort incentives to remove equilibria. This way when I do remove an equilibrium, it is hard to object. Then I show that given a contract game, the set of equilibria that survive my conservative refinement process all generate the same continuation payoff process. This means no matter how my "minimal" refinement is strengthened, as long as the strengthening does not remove all equilibria from a contract game then Pareto-optimal contracts are unchanged. This "squeeze" argument implies that my refinement and the resulting optimal contract are robust.

To operationalize my conservative approach to removing equilibria, I begin by defining some restrictive conditions on assessments that will need to be satisfied for there to be an opportunity for P to maximize effort incentives.

Definition. *$W_t(H^{A+}_{t-\Delta}, h_{t-\Delta})$ is belief-free if it does not depend on A's beliefs at all succeeding $(H^A_t, h_{t-\Delta})$. W_t is public given $h_{t-\Delta}$ if $W_t(H^{A+}_{t-\Delta}, h_{t-\Delta})$ is constant across all $H^{A+}_{t-\Delta}$, in which case I simplify $W_t(H^{A+}_{t-\Delta}, h_{t-\Delta})$ to $W_t(h_{t-\Delta})$. Define belief-free and public for V_t similarly.*

(a, m) is belief-free given h_t if at every succeeding decision node the corresponding player's set of best-response continuation strategies does not depend on that player's belief.

See Ely, Hörner and Olszewski (2005) for a discussion of belief-free equilibria in repeated games of private monitoring. To define when a sequential equilibrium is removed, I suppose the set of all sequential equilibria has already been whittled down to some subset $\mathcal{E}$. I then provide restrictive conditions as a function of $\mathcal{E}$ under which certain additional sequential equilibria can be removed.

Definition. *Fix a set sequential equilibria $\mathcal{E}$. P is said to have an opportunity to maximize effort incentives at the beginning of date t given $\mathcal{E}$ and conditional on $h_{t-\Delta}$*

if for every succeeding h_t, all $(a, m) \in \mathcal{E}$ are belief-free given h_t and share the same belief-free, public continuation payoff process $(W_{s+\Delta}(h_s), V_{s+\Delta}(h_s))_{s \geq t}$.

When P has an opportunity, her set of best response messages is

$$\mathcal{M}^*(h_{t-\Delta}) := \arg\max_{m' \in \mathcal{M}} \mathbf{E}[e^{-r\Delta} V_{t+\Delta}(h_t) \mid h_{t-\Delta} m'].$$

Notice P has an opportunity to maximize effort incentives at date t only when the equilibrium property of all $(a, m) \in \mathcal{E}$ starting from date $t + \Delta$ do not depend on what happens before date $t + \Delta$ and the continuation payoff processes of A and P starting from date $t+\Delta$ are uniquely determined and do not depend on what happens before date $t + \Delta$. Thus, when P has an opportunity at date t one can think of date t as the terminal date with the players receiving lump sum payments

$$(\mathbf{E}[e^{-r\Delta} W_{t+\Delta}(h_t) \mid h_{t-\Delta} m_t], \mathbf{E}[e^{-r\Delta} V_{t+\Delta}(h_t) \mid h_{t-\Delta} m_t])$$

at the end of date t after P makes her final report m_t.

Definition. *Suppose P has an opportunity to maximize effort incentives given $\mathcal{E}$ and conditional on $h_{t-\Delta}$. A commitment $\hat{m}_t(h_{t-\Delta})$ is a choice of a message $\in \mathcal{M}^*(h_{t-\Delta})$ for each $(H_t^P, h_{t-\Delta})$ that depends on H_t^P only up to X_t.*

Given a commitment, A's best response effort does not depend on A's belief about P's private history and is, therefore, public. Consequently, P's date t continuation payoff from making a commitment does not depend on P's belief about A's private history and is, therefore, belief-free and public:

Define $a_t | \hat{m}_t(h_{t-\Delta})$ to be the largest element of

$$\arg\max_{a'} \mathbf{E}_{a', \hat{m}_t(h_{t-\Delta})} \left[-h(a')\Delta + e^{-r\Delta} W_{t+\Delta}(h_t) \right]$$

where the expectation is computed using the distribution over the set of h_t compatible with $h_{t-\Delta}$ generated by a date t effort a' and $\hat{m}_t(h_{t-\Delta})$. Define

$$V_t(h_{t-\Delta}) | \hat{m}_t(h_{t-\Delta}) := \mathbf{E}_{a_t | \hat{m}_t(h_{t-\Delta}), \hat{m}_t(h_{t-\Delta})} \left[u(a_t | \hat{m}_t(h_{t-\Delta}))\Delta + e^{-r\Delta} V_{t+\Delta}(h_t) \right].$$

I now implicitly define when an equilibrium can be removed by defining when a set of equilibria can no longer be further refined:

Definition. *A set $\mathcal{E}$ of sequential equilibria maximizes effort incentives if whenever P has an opportunity to maximize effort incentives conditional on $h_{t-\Delta}$, there does not exist an $(a, m) \in \mathcal{E}$, $H_{t-\Delta}^P$, and a commitment $\hat{m}_t(h_{t-\Delta})$ such that $V_t(h_{t-\Delta}) | \hat{m}_t(h_{t-\Delta}) > V_t(H_{t-\Delta}^P, h_{t-\Delta})$.*

The order in which equilibria are removed under my conservative approach does not matter:

Lemma 4. *If $\mathcal{E}_1$ and $\mathcal{E}_2$ are sets of sequential equilibria that maximizes effort incentives, then so is $\mathcal{E}_1 \cup \mathcal{E}_2$. Thus, there is a unique maximal set $\mathcal{E}^*$ of sequential equilibria that maximizes effort incentives.*

Proof. Suppose $\mathcal{E}_1 \cup \mathcal{E}_2$ does not maximize effort incentives. Then there exists an $h_{t-\Delta}$, $(a, m) \in \mathcal{E}_1 \cup \mathcal{E}_2$, $H^P_{t-\Delta}$, and a commitment $\hat{m}_t(h_{t-\Delta})$ such that P has an opportunity conditional on $h_{t-\Delta}$ and $V_t(h_{t-\Delta})|\hat{m}_t(h_{t-\Delta}) > V_t(H^P_{t-\Delta}, h_{t-\Delta})$.

Without loss of generality, assume $(a, m) \in \mathcal{E}_1$. Then P has an opportunity given $\mathcal{E}_1$ and conditional on $h_{t-\Delta}$. Given $\mathcal{E}_1$, $\hat{m}_t(h_{t-\Delta})$ continues to be a commitment. Moreover, the payoffs $V_t(h_{t-\Delta})|\hat{m}_t(h_{t-\Delta})$ and $V_t(H^P_{t-\Delta}, h_{t-\Delta})$ are the same given $\mathcal{E}_1$ and $\mathcal{E}_1 \cup \mathcal{E}_2$. This contradicts the assumption that $\mathcal{E}_1$ maximizes incentive power. □

Definition. *A sequential equilibrium maximizes effort incentives if it is an element of $\mathcal{E}^*$.*

Proposition 3. *Fix a contract game. All $(a, m) \in \mathcal{E}^*$ are belief-free and generate the same belief-free, public continuation payoff process that can be computed recursively:*

When $\tau = t$, all $(a, m) \in \mathcal{E}^*$ generate the same belief-free, public continuation payoff $(W_t(h_{t-\Delta}), V_t(h_{t-\Delta})) = (w_t(h_{t-\Delta}), -w_t(h_{t-\Delta}))$. If $\tau > t$, then by induction suppose all $(a, m) \in \mathcal{E}^*$ generate the same belief-free, public continuation payoff $(W_{t+\Delta}(h_t), V_{t+\Delta}(h_t))$ for all h_t. Define

$$R(h_{t-\Delta}) := \{\mathbf{E}[e^{-r\Delta} W_{t+\Delta}(h_t) \mid h_{t-\Delta} m'] \mid m' \in \mathcal{M}^*(h_{t-\Delta})\}.$$

Then $m_t(H^P_t, h_{t-\Delta}) = f^{R(h_{t-\Delta})}(X_t)$, $a_t(h_{t-\Delta}) = a^{R(h_{t-\Delta})}$, and

$$\begin{aligned} W_t(h_{t-\Delta}) &= w_t(h_{t-\Delta}) - h\left(a^{R(h_{t-\Delta})}\right)\Delta + e^{-r\Delta}\mathbf{E}_{a^{R(h_{t-\Delta})},\ f^{R(h_{t-\Delta})}(X_t)} W_{t+\Delta}(h_t), \\ V_t(h_{t-\Delta}) &= -w_t(h_{t-\Delta}) + \mathbf{E}_{a^{R(h_{t-\Delta})},\ f^{R(h_{t-\Delta})}(X_t)}\left[u(a^{R(h_{t-\Delta})})\Delta + e^{-r\Delta} V_{t+\Delta}(h_t)\right]. \end{aligned}$$

Proof. Begin with the set $\mathcal{E}_T$ of all sequential equilibria. It is easy to verify that P has an opportunity given $\mathcal{E}_T$ and conditional on any public history of the form $h_{T-2\Delta}$ satisfying $\tau(h_{T-2\Delta}) > T-\Delta$. Now, define a new set $\mathcal{E}_{T-\Delta} \subset \mathcal{E}_T$ of sequential equilibria as follows: $(a, m) \in \mathcal{E}_{T-\Delta}$ if and only if for each $h_{T-2\Delta}$ satisfying $\tau(h_{T-2\Delta}) > T - \Delta$ there is a $f^{R(h_{T-2\Delta})}(X_{T-\Delta})$ such that $m_{T-\Delta}(H^P_{T-\Delta}, h_{T-2\Delta}) = f^{R(h_{T-2\Delta})}(X_{T-\Delta})$ for all $H^P_{T-\Delta}$ and $a_{T-\Delta}(H^A_{T-\Delta}, h_{T-2\Delta}) = a^{R(h_{T-2\Delta}))}$ for all $H^A_{T-\Delta}$. By construction, $\mathcal{E}^* \subset \mathcal{E}_{T-\Delta}$.

Now it is easy to verify that P has an opportunity given $\mathcal{E}_{T-\Delta}$ and conditional on any public history of the form $h_{T-3\Delta}$ satisfying $\tau(h_{T-3\Delta}) > T - 2\Delta$. Similar to before, I can now define an $\mathcal{E}_{T-2\Delta}$ that contains $\mathcal{E}^*$. Proceeding inductively, I can define a nested sequence of sets of sequential equilibria $\mathcal{E}^* \subset \mathcal{E}_0 \subset \ldots \subset \mathcal{E}_{T-\Delta} \subset \mathcal{E}_T$. All equilibria in the set $\mathcal{E}_0$ are belief-free and generate the same belief-free, public continuation payoff process that is described in the proposition. It is easy to show $\mathcal{E}_0$ maximizes incentive power, which implies $\mathcal{E}^* = \mathcal{E}_0$. □

Despite my conservative approach to removing sequential equilibria, Proposition 3 implies that all the "complex" sequential equilibria involving P trying to keep A in the dark about his own continuation payoff are removed.

Proposition 3 says that P's report strategy at date t is characterized by the function $f^{R(h_{t-\Delta})}$. Given the definition of f^R in Assumption 1 and given that $R(h_{t-\Delta})$ is defined to be all the possible expected discounted date $t+\Delta$ continuation payoffs for A as a function of P's date t report, Proposition 3 basically formalizes the claim that P reports in a way that maximizes effort incentives at all times.

Definition. *A contract is incentive compatible if the assessment is a sequential equilibrium that maximizes effort incentives and* $W_t(h_{t-\Delta}) + V_t(h_{t-\Delta}) \geq 0$ *for all* $h_{t-\Delta}$.

The second part of the definition is an interim participation constraint. If it is violated both players are strictly better off terminating at the beginning of date t under some severance pay $\hat{w}_t$.

The Optimal Contracting Problem: For each point on the Pareto-frontier, find an incentive-compatible contract that achieves it.

Theorem 4. *Every payoff on the Pareto-frontier can be achieved by a contract with the following structure:*

- $\mathcal{M} = Im(X_t)$ *and* $m_t(H_t^P, h_{t-\Delta}) = X_t$.
- *For each* $t < T$ *there is a pair of constants* $w_t^{salary}, w_{t+\Delta}^{severance}$ *such that* A *is paid* w_t^{salary} *at date* t *for working and is paid a severance* $w_{t+\Delta}^{severance}$ *at date* $t+\Delta$ *if he is terminated at the beginning of date* $t+\Delta$. *Termination at date* $t+\Delta$ *occurs with some probability* $p_t^*(X_t)$.

Proof of Theorem 4. Proposition 3 implies there is an obvious correspondence between the portion of a contract after a history $h_{t-\Delta}$ – call it the date t continuation contract given $h_{t-\Delta}$ – and a contract in the version of the model with timeframe $[0, T-t]$.

The proof is by induction on the length of the model timeframe. Fix a Pareto-optimal contract. There is at least some realization of X_0 such that for all h_0 succeeding $m_0(X_0)$, the date Δ continuation contract given h_0 is a Pareto-optimal contract in the model with timeframe $[0, T-\Delta]$. Without loss of generality, it is the same Pareto-optimal contract $\mathscr{C}^\Delta$. Now for any realization of X_0 change the contract so that after P reports $m_0(X_0)$ the contract randomizes between $\mathscr{C}^\Delta$ and termination using the date 0 public randomizing device. This can be done in a way so that $\mathbf{E}[W_\Delta(h_0) \mid m_0(X_0)]$ and $\mathbf{E}[V_\Delta(h_0) \mid m_0(X_0)]$ remain the same. By construction, the altered contract remains incentive-compatible. Relabelling $m_0(X_0)$ as X_0 (if two realizations of X_0 lead to the same $m_0(X_0)$ then just create two separate messages – it won't affect anything), the contract now has the structure described in Theorem 1 at date 0. By induction, it has the structure described in Theorem 1 at all other dates. $\square$

Corollary 3. *When a_t has a strictly monotone effect on X_t for all t, then the optimal contract converges to the one characterized in Theorem 1 as $T \to \infty$.*

References

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

[2] Aghion, P. and P. Bolton (1992) "An Incomplete Contracts Approach to Financial Contracting," *Review of Economic Studies* Vol. 59. pp. 473-494

[3] Aghion, P. and J. Tirole (1997) "Formal and Real Authority in Organizations," *Journal of Political Economy* Vol. 105, pp. 1-29

[4] Bénabou, R. and J. Tirole (2003) "Intrinsic and Extrinsic Motivation," *Review of Economic Studies* Vol. 70, pp. 489-520

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

[6] Burkart, M., Gromb, D. and F. Panunzi (1997) "Large Shareholders, Monitoring, and the Value of the Firm," *Quarterly Journal of Economics* Vol. 97, pp. 693-728

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

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

[9] Edmans, A., X. Gabaix, T. Sadzik and Y. Sannikov (2012) "Dynamic CEO Compensation," *Journal of Finance* Vol. 67, pp. 1603-1647

[10] Ely, J., J. Hörner and W. Olszewski (2005) "Belief-Free Equilibria in Repeated Games," *Econometrica* Vol. 73, pp. 377-415

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

[12] Georgiadis, G. and B. Szentes (2018) "Optimal Monitoring Design," working paper

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

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

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

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

[17] Li, A. and M. Yang (2018) "Optimal Incentive Contract with Endogenous Monitoring Technology," working paper

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

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

[20] Sannikov, Y. (2008) "A Continuous-Time Version of the Principal-Agent Problem," *Review of Economic Studies* Vol. 75, pp. 957 - 984

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

[22] Townsend, R. (1979) "Optimal Contracts and Competitive Markets with Costly State Verification," *Journal of Economic Theory* Vol. 21, pp. 265-293

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

[24] Zhu, J. Y. (2013) "Optimal Contracts with Shirking," *Review of Economic Studies* Vol. 80, pp. 812-839

[25] Zhu, J. Y. (2018a) "Myopic Agency," *Review of Economic Studies* Vol. 85, pp. 1352-1388

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