

BETTER MONITORING . . . WORSE PRODUCTIVITY?

JOHN Y. ZHU¹

November 23, 2018

Abstract

A principal privately monitors an agent's hidden efforts. The agent's contract depends on the principal's unverifiable reports of her private information. How does the quality of the principal's private monitoring affect the optimal contract and agent productivity? When monitoring generates information that is at once imprecise (weak statistical power) but sensitive to effort (strong incentive power) the principal is unable to commit not to be too tough on the agent. Improving a well-functioning monitoring system by adding new information that is weak in statistical power but strong in incentive power can cause the optimal dynamic contract to collapse and induce zero effort.

JEL Codes: C73, D86, J41, M5

Keywords: monitoring, moral hazard, information design, private monitoring, optimal contracts, principal agent,

¹I thank Dilip Abreu and David Pearce.

1 Introduction

In principal-agent relationships with moral hazard, the provision of incentives relies on information generated by monitoring. Because of this dependence, there is a common assumption that better monitoring must mean better productivity. In this paper, I describe a channel through which better monitoring can sometimes lead to lower agent productivity and, ultimately, a worse outcome for the principal.

My better monitoring/worse outcome result is quite general. The main thing it relies on is monitoring being private, or at least having a private component: This means as the agent repeatedly exerts hidden effort, the principal privately observes information about those efforts. Over time, the principal makes unverifiable reports of her private information that then affect the agent's payoff through an optimal dynamic contract specifying pay and termination all as a function of those reports.

In this setting how can better monitoring lead to a worse outcome? Since monitoring is private, the link between information and incentives is filtered through the principal's strategic reports. And now what can happen is there are types of information that, if added to the monitoring system, would allow the principal to be "too tough" on the agent ex-interim – that is, induce a lot of effort but at great cost to efficiency. Because being too tough is inefficient, the principal ex-ante wants to commit not to be too tough in the future. But this is tantamount to committing to a particular way to make unverifiable reports which is not credible. The lack of credibility causes the contracting parties to preemptively agree to an optimal contract that curtails the principal's ability to punish the agent, hindering incentive provision. Consequently, effort drops and the principal ends up worse off than before.

What improvements to monitoring are counterproductive? *Adding new information with weak statistical power but strong incentive power is counterproductive.*

I will define statistical and incentive power later. For now, here is an intuitive way to think about those concepts. Imagine the following toy setting: An agent has two hidden effort choices $e \in \{0, 1\}$ labelled by their costs to the agent. Effort 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 inflict a punishment p on the agent (representing the amount of agent utility that is destroyed) whenever b is realized. In this setting, one question that could be asked is how much expected punishment does the principal need to inflict on the agent in order to induce effort? The answer measures the statistical power of information. The smaller the expected punishment, the more statistically powerful is the information. Another question one could ask is how large does the punishment need to be in order to induce effort? The smaller is the punishment the more incentive power is contained by the information. An easy computation shows that ordering information based on $\frac{q_0}{q_1}$ respects statistical power while ordering information based on $q_0 - q_1$ respects incentive power.

A perfectly informative signal has the strongest possible statistical power and incentive power. A completely uninformative signal has the weakest possible statistical

power and incentive power. Because statistical and incentive power coincide at the extremes, it is easy to confound the two concepts, mistakenly assuming that a weak (strong) signal must mean weak (strong) incentives. But it is apparent that many types of information differ greatly along these two measures of information power and my paper demonstrates that such information can play a pivotal role in determining welfare. In particular, I will show that under private monitoring adding new information that is sufficiently weak in statistical power but sufficiently strong in incentive power to a well-functioning monitoring system will cause the optimal dynamic contract to completely collapse into a trivial arrangement that induces zero effort from the agent at all times.

1.1 Related Literature

Conceptually, optimal contracting papers can be seen 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 optimal contract written based on that information is determined. Much of the literature focuses on the second step.

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 do contracts react to incompleteness, 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. Questions regarding the frequency of performance evaluations, organizational transparency, and the use of monitoring software are all issues of monitoring design and are relevant to how the principal optimally contracts with the agent. In this paper I take a step toward developing a theory of monitoring design by highlighting why distinguishing between the statistical and incentive power of information is important to understanding how the information content of monitoring is determined.

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.

Another way to position my better monitoring/worse outcome result, which is established under private monitoring, is to look at related results in the public mon-

itoring 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 highlight the role played by information that is weak in statistical power but strong in incentive 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. The refinement is similar in spirit to the one in Dewatripont (1987) that selects perfect equilibria in sequential models of spatial competition.

2 Model and Optimal Contract

This section introduces the model and characterizes the optimal contract holding the monitoring structure fixed. In subsequent sections I then explore how changes to the information content of monitoring affects productivity and firm value.

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 $t = 0, \Delta, 2\Delta, \dots$. 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 $\lim_{a_t \rightarrow 1} h(a_t) = \infty$. After A exerts effort, P monitors A : First, P observes a private signal X_t taking finitely many values. a_t determines the distribution of X_t . I assume effort has a *monotone* effect on X_t : $Im(X_t)$ can be divided into disjoint subsets *Good* and *Bad* such that $\mathbf{P}(X_t = x \mid a_t)$ is strictly increasing (decreasing) in a_t if and only

if $x \in Good$ ($x \in Bad$). Given X_t , P 's utility is $u(X_t)$. I assume $\mathbf{E}_{a_t}u(X_t)$ is a strictly increasing, weakly concave function of effort and $\mathbf{E}_0u(X_t) > 0$.² 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, τ) 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, τ) , 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_{t-1}^A 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_t^P of observations $\{X_s\}_{s \leq t}$. The system of beliefs consists of a belief about $H_{t-\Delta}^P$ at each decision node $(H_{t-1}^A, h_{t-\Delta})$ of A , and a belief about H_t^A at each decision node $(H_t^P, h_{t-\Delta})$ 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

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

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

2.1 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.³

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) and is somewhat tangential to understanding better monitoring/worse

²In the next section when I look at how changes to X_t affect outcomes it will be assumed that the function $\mathbf{E}_{a_t}u(X_t)$ remains unchanged.

³It 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.

outcome. To avoid breaking the flow of the paper, I now go directly to describing the optimal contract. Since I have not yet defined what it means for a contract to be incentive compatible, I can only highlight ways in which the optimal contract's strategy profile "seems reasonable" given its contract game. The formal definition of incentive compatibility, taken from Zhu (2018b), is developed in Appendix A.

The optimal contract is a simple efficiency wage contract in the spirit of Shapiro and Stiglitz (1984) that induces effort through the threat of termination.

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 $w_{severance}$.

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 agent performance. 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. By the same argument, P must be indifferent between reporting *pass* and *fail* at all times,

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

Otherwise, P would always want to report the message that led to the higher $V_{t+\Delta}$. 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,

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

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

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

The solution can be recursively computed by setting $S_0^* = \mathbf{E}_{a_t=0} u(X_t)$ 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 $i = 0, 1, 2 \dots$ and $S^* = S_\infty^*$. Finally, w_{salary} is determined by A 's binding ex-ante participation constraint $W_0 = 0$,

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

The tractability of the optimal contract is due in part to P 's simple report strategy: P fails A at date t if and only if she sees something bad at date t . To see why P always reports this way, let us pick an arbitrary date t and examine P 's date t payoff,

$$V_t = \mathbf{E}_{a_t|m_t} u(X_t) + e^{-r\Delta} V_{t+\Delta}. \quad (1)$$

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. By design, the second component does not depend on m_t . 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,

Remark 1. *At date t , P will report in a way that maximizes date t effort incentives.*

This key property of P 's report strategy will be deduced in the appendix as a consequence of my formal definition of incentive compatibility. What kind of m_t maximizes date t effort incentives? Intuitively, m_t should depend only on X_t instead of the entire history of information $\{X_s\}_{s \leq t}$ generated by monitoring: Older information is not informative of effort today and conditioning today's report on older information will only dilute effort incentives today. Of course, there are still many ways for m_t to depend only on X_t . However, if the goal is to maximize effort incentives, it is

clear P should reward A whenever something good happens and punish A whenever something bad happens, which is precisely what P does in the optimal contract.

3 Better Monitoring/Worse Outcome

With the optimal contract characterized, I am now ready to perform the main comparative statics exercise of the paper: How do S^* (P 's payoff) and a^* (A 's productivity) change as I change the monitoring technology $\{X_t\}_{t \geq 0}$?

Theorem 1 hints at how better monitoring might lead to a worse outcome. Recall, P fails A today if and only if she observes a *Bad* signal today. This is a simple 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 system generates *Bad* signals that are mostly strong – that is, if the likelihood of occurring rapidly declines as A increases effort – then intuitively the answer is yes. Where this strategy becomes inefficient is when the monitoring system generates *Bad* signals that are mostly weak. In this case, one would like to see P be a little more discriminating and fail A only if she sees a strong *Bad* signal, or at least wait until she has seen multiple weak *Bad* signals across many dates before failing A . But 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. 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 P and A understand that in the future, if the monitoring system will generate lots of weak *Bad* signals, P will likely over-fail A . To counteract this, the contracting parties then preemptively write 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 really weak, 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.

To recap, I have argued that Theorem 1 suggests that when monitoring generates mostly weak *Bad* signals then one should expect effort to drop (possibly all the way to zero) and P to be worse off relative to when monitoring generates mostly strong *Bad* signals.

Is it possible to take a monitoring system that generates mostly strong *Bad* signals and *improve* it to the point where it generates mostly weak *Bad* signals? Because if it is possible, then better monitoring/worse outcome is implied.

In the next section I show that it is generically possible to take a well-functioning monitoring system and improve it in a way so that *Bad* signals become diluted, the optimal contract collapses, and P becomes worse off. Moreover, I show that

it is the same basic strategy every time: Introduce new information that is weak in statistical power but, relative to the old information, strong in incentive power. Before establishing this result in general, let us first work through an explicit example that demonstrates the basic idea.

3.1 An Example

In this example I will begin with a bad news Poisson monitoring system. “Bad news” means that the Poisson event is indicative of lower effort rather than higher effort. I show that the bad news Poisson monitoring system generates a *Bad* signal that is strong in some formal sense and consequently the optimal contract induces positive effort. I then improve the monitoring system by adding a conditionally independent Brownian signal where the drift of the Brownian motion is controlled by A ’s effort. I explain that Brownian motion has weak statistical power but very strong incentive power. I then show that in the improved monitoring system that generates both bad news Poisson information and Brownian information, the typical *Bad* signal suddenly becomes extremely weak. Consequently, the optimal contract collapses and P becomes worse off.

In a bad news Poisson monitoring system, 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. How strong is the bad news Poisson *Bad* signal? The formal measure is the negative effort elasticity of $\mathbf{P}(Bad)$:

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

This measure corresponds to the intuitive measure of statistical power discussed in the introduction and so from now on the statistical power of information means the strength of the *Bad* signal.⁴ A simple computation shows that the negative effort-

⁴In my private monitoring setting, the question analogous to the one posed in the introduction concerning statistical power is: How little additional expected surplus destruction does there need to be in order to induce an additional increment of effort? The answer is captured by (i.e. a monotonic function of) (2). This measure of statistical power is effort-dependent. However, when compared to the Brownian information that will be introduced shortly, the statistical power of bad news Poisson information is greater no matter the effort level. Thus, I can talk about bad news Poisson information being more statistically powerful than Brownian information without reference to effort level. Similarly, Brownian information has stronger incentive power than bad news Poisson information irrespective of effort level. In the general analysis, I will continue to restrict attention to a family of signals that can be ranked by statistical power and incentive power without reference

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 that the *Bad* signal of bad news Poisson monitoring is a strong one and it is not too inefficient to let P punish A nontrivially whenever she sees it. Consequently, the optimal contract under bad news Poisson monitoring can induce positive effort. Later I generalize this result by showing that for a broad class of monitoring systems if the negative effort elasticity of $\mathbf{P}(Bad)$ is not “too small,” then there exist parameterizations of the rest of the model such that the optimal contract induces positive effort.

Let us now see what happens when the bad news Poisson monitoring system 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 a weak *Bad* signal, and because it is a weak *Bad* signal, it is important that P be discriminating when using Brownian information. This basically means that P needs to commit to be patient and wait until she has seen many 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 maximizing 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 system (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 to effort level.

in place – then it doesn't matter how noisy it is, P will not ignore it.

Is it possible for a piece of information to simultaneously have weak statistical power and strong incentive power? 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 weak statistical power. The formal measure of incentive power, which matches the one in the introduction, is the negative derivative of $\mathbf{P}(Bad)$ with respect to effort:

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

While the measure for incentive power looks like the measure for statistical power, they are not the same, and it is very easy to come up with information such that the first derivative is very small but the second one is very large.

Now to show that P will not ignore Y_t let us compute the incentive power of both X_t and Y_t . The incentive power of X_t is

$$\lambda\Delta.$$

The incentive power of Y_t is

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

Thus, when a bad news Poisson monitoring system is improved by adding a Brownian component, there is no way P will ignore the Brownian component.⁵

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* 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* 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.

⁵In general, the new information does not need to have 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.

But now we have a problem: This combined *Bad* signal is quite common – occurring about half the time no matter what effort A puts in. Consequently, punishing A whenever this very common, very weak *Bad* signal occurs is extremely inefficient and should be avoided at all costs. But as I’ve said before – there is nothing one 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 bad news Poisson event occur, in which case A “deserves” to get punished, or if the only thing P saw was the 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.

3.2 A More General Analysis

What aspects of Brownian information made introducing it into a Poisson monitoring system so counterproductive?

One important property of Brownian information that emerged in the analysis is that the Brownian *Bad* signal is weak:

1. New information has weak statistical power.

Intuitively, the weak *Bad* signal of the new information can help cause the improved monitoring system to generate mostly weak *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 strong incentive power.

Here, the idea is even if the new information has extremely weak *Bad* signals, if P ignores the new information, then introducing it makes no difference. To ensure P doesn’t ignore the new information, it must have sufficiently strong incentive power.

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 very weak and even if P fails A based off of it, if the signal is extremely rare 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.

I now show in a formal sense that it is these three and only these three conditions that matter for better monitoring/worse outcome.

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

$$\begin{aligned}\lim_{\Delta \rightarrow 0} \mathbf{E}_{a_t} u(X_t) &= \Theta(\Delta) \\ \lim_{\Delta \rightarrow 0} -\frac{d}{da_t} \mathbf{P}(X_t = \text{Bad} \mid a_t) &= \Theta(\Delta^\alpha) \text{ for some } \alpha \geq 0 \\ \lim_{\Delta \rightarrow 0} \mathbf{P}(X_t = \text{Bad} \mid a_t) &= \Theta(\Delta^{\gamma^b}) \text{ for some } \gamma^b \geq 0 \\ \lim_{\Delta \rightarrow 0} \mathbf{P}(X_t = \text{Good} \mid a_t > 0) &= \Theta(\Delta^{\gamma^g}) \text{ for some } \gamma^g \geq 0.\end{aligned}$$

Obviously, $\alpha \geq \max\{\gamma^b, \gamma^g\}$, and γ^g or γ^b is equal to zero. α measures the incentive power of information – the lower is α the more effort that can be induced. $\alpha - \gamma^b$ is the exponent associated with the negative effort elasticity of $\mathbf{P}(\text{Bad})$. It measures the strength of a typical *Bad* signal. The smaller is $\alpha - \gamma^b$, the more statistically powerful is the information. Information with “large” $\alpha - \gamma^b$ but “small” α (i.e. weak statistical power but strong incentive power) will play a central role in the better monitoring/worse outcome result.

This class of monitoring technologies includes the familiar cases where effort affects the drift of a Brownian motion:

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

the intensity of a good news Poisson process:

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

and the intensity of a bad news Poisson process:

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

Definition. X_t is *Brownian* if $(\alpha, \gamma^b, \gamma^g) = (.5, 0, 0)$, *good news Poisson* if $(\alpha, \gamma^b, \gamma^g) = (1, 0, 1)$, *bad news Poisson* if $(\alpha, \gamma^b, \gamma^g) = (1, 1, 0)$.

Theorem 2. If $\alpha - \gamma^b = 0$ and $\gamma^b \leq 1$ 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. □

Given that $\alpha - \gamma^b$ measures the strength of a typical *Bad* signal, Theorem 2 formalizes an intuition sketched out at the beginning of this section: When a typical *Bad* signal is strong (and not too rare) the optimal contract can induce positive effort; when a typical *Bad* signal is weak (or too rare), the optimal contract collapses and 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 statistical power and incentive power, is different from most work in the repeated games literature which emphasizes only statistical power. In particular, whereas better monitoring can lead to a worse outcome in my setting, in most repeated games models improvements to the information content of monitoring always weakly improve the scope for cooperation.

Armed with Theorem 2 I can now investigate how improvements to the monitoring system 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 monotone effect on the vector valued information (X_{1t}, X_{2t}) generated by the improved monitoring system. 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 $\Delta \rightarrow 0$, 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 $A(\Delta) = \Theta(\Delta^{\alpha_1 + \gamma_1^b})$ 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 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 monotonic function of a_t in the continuous-time limit is $\alpha_1 + \gamma_2^g \neq \gamma_1^b + \alpha_2$. Thus,

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

Proposition 1. *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*

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

Proposition 1 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. □

Proposition 1 yields an explicit characterization of counterproductive improvements to the monitoring system.

Corollary 2. *Suppose $\alpha_1 - \gamma_1^b = 0$. If $\alpha_2 - \gamma_2^b > 0$, $\alpha_2 < \alpha_1 + \gamma_2^b$, and $\gamma_2^b < \gamma_1^b$, then $\alpha - \gamma^b > 0$. If any of the three inequalities is reversed then $\alpha - \gamma^b = 0$.*

Corollary 2 formalizes how introducing new information that is weak in statistical power but strong in incentive power is counterproductive. It implies the earlier example showing that improving bad news Poisson monitoring by adding a Brownian component causes the optimal contract to induce zero effort.

Given a monitoring system X_{1t} under which the optimal contract induces positive effort, Corollary 2 lists three inequalities that, if satisfied by the new information X_{2t} being added, leads to the worst outcome in which the optimal contract induces zero effort. These three inequalities correspond to the three intuitive conditions for better monitoring/worse outcome listed in the beginning of this subsection. Given Theorem 2 the first inequality says that the new information has a weak *Bad* signal, which is condition 1 from before. The second inequality, which is an upper bound on α_2 , says that the new information has sufficiently strong incentive power relative to the original information which is condition 2 from before. Finally, the third inequality, which is an upper bound on γ^b , says that the *Bad* signal of the new information is sufficiently common which is condition 3 from before.

It is worth emphasizing that for the second inequality the upper bound on α_2 is $\alpha_1 + \gamma_2^b$, not α_1 : For better monitoring/worse outcome to appear, the new information has to have sufficiently strong incentive power but it does not necessarily

have to have stronger incentive power than the information generated by the original monitoring system. The reason I mention this is because Proposition 1 implies that $\alpha = \min\{\alpha_1, \alpha_2\}$ which gives the incorrect impression that the principal just picks the X_{it} with the strongest incentive power and reports solely based off of that. If that were the case then improving monitoring by adding an X_{2t} with strictly weaker incentive power ($\alpha_2 > \alpha_1$) could never cause the optimal contract to collapse, contradicting Corollary 2.

4 Dynamic Monitoring Design

I now use my better monitoring/worse outcome result to shed some light on optimal monitoring design.

The better monitoring/worse outcome result suggests that there is value to limiting the information observed by P . Theorem 2 shows that in the continuous-time limit the optimal contract induces zero effort under a wide range of monitoring technologies such as Brownian monitoring. Given this negative result, how might positive effort be restored by limiting information? The repeated games literature has emphasized the idea 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 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 for granted, 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 make P better off.

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. Thus, P 's desire to maximize date t_2 effort incentives will lead her to report in a way so that A 's date $t_2 + \Delta$

continuation payoff depends only on X_{t_2} . Anticipating this, A exerts zero effort from t_1 to $t_2 - \Delta$. More generally, on any date t where P 's date t report has no payoff impact, A 's date t effort is zero.

This discussion of infrequent monitoring as well as the better monitoring/worse outcome result shows how the relationship between the information content of monitoring and firm value is fundamentally different in my model compared to most repeated games models. In most of the repeated games literature, the only thing that matters is the statistical power of information. 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, it is the interplay between statistical power and incentive power that matters, and 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/samples, 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/sampling from the type that is typically referred to in the repeated games literature, I will refer to this type of infrequent monitoring/sampling simply 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 , the private informative signal she observes is no longer X_t where

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

but, rather, $Y_t = \sum_{s \leq t} X_s$.

In this new model, a contract game, in addition to specifying \mathcal{M} , w , and τ , also specifies a *monitoring design* e , consisting of a predictable sequence of random monitoring times $e_1 < e_2 < \dots$. Predictable means that e_{i+1} is measurable with respect to h_{e_i} . An assessment is defined similarly to before except P 's decision nodes only occur on monitoring dates. Nevertheless, for each date t , W_t and V_t can be defined similar to before. Throughout the analysis I will assume an infinitesimal Δ .

Notice, in this new model, if one restricts attention to contracts that have P monitor every date, then the optimal contracting problem becomes identical to the original one and Theorem 2 implies that A exerts zero effort. However, it turns out, monitoring every date is in general sub-optimal:

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

- P monitors A every Δ^* units of time: $e = \{\Delta^*, 2\Delta^*, 3\Delta^* \dots\}$
- $\mathcal{M} = \{\text{pass}, \text{fail}\}$.
- For each $k \in \mathbb{Z}^+$, $m_{k\Delta^*} = \text{fail}$ iff $Y_{k\Delta^*} - Y_{(k-1)\Delta^*} \leq \rho^*$.
- w consists of a pair of constants $w_{\text{salary}}, w_{\text{severance}}$.
- For each $k \in \mathbb{Z}^+$, if $m_{k\Delta^*} = \text{pass}$ then A is retained for the monitoring period $(k\Delta^*, (k+1)\Delta^*]$ and is paid a stream $w_{\text{salary}}dt$.
- For each $k \in \mathbb{Z}^+$, if $m_{k\Delta^*} = \text{fail}$ 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 $w_{\text{severance}}$.

Notice this optimal contract is virtually identical to the original optimal contract – the main difference being that in the original model the length of a monitoring period was exogenously fixed to be Δ whereas in the new model it is endogenously determined.

The magnitude Δ^* of this endogenously determined monitoring 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 weak and consequently,

insisting on failing A whenever these “*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^* becomes large, statistical power goes up but now discounting begins eroding the incentive power of information: In the beginning of a monitoring period, the threat of termination in the distant future when the monitoring period concludes has little effect on the continuation payoff of A today. The optimal Δ^* balances these two opposing forces: The desire for greater statistical power on the one hand versus the desire for greater incentive power on the other.

5 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 statistical power versus more incentive power.

6 Appendix

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

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

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

Proof of Proposition 1. 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 $\Delta \rightarrow 0$, 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 $(g_1, b_2) \in \text{Bad}$. By the product rule, as $\Delta \rightarrow 0$, 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 $(b_1, g_2) \in \text{Bad}$.

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

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

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

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

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

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

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

I am now ready to discuss incentive-compatibility in settings where the monitoring system 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 $T \rightarrow \infty$ 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 : Looking at (1), the sequential equilibrium concept allows P to commit to a report strategy m_t that does not maximize $a_t | m_t$. 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. $W_t(H_{t-\Delta}^{A+}, h_{t-\Delta})$ is belief-free if it does not depend on A 's beliefs at all succeeding $(H_t^A, h_{t-\Delta})$. W_t is public given $h_{t-\Delta}$ if $W_t(H_{t-\Delta}^{A+}, h_{t-\Delta})$ is constant across all $H_{t-\Delta}^{A+}$, in which case I simplify $W_t(H_{t-\Delta}^{A+}, h_{t-\Delta})$ 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 incentive power 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

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

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

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

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 incentive power 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 2. *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 $h_{t-\Delta}$, $(a, m) \in \mathcal{E}_1 \cup \mathcal{E}_2$, $H_{t-\Delta}^P$, and a commitment $\hat{m}_t(h_{t-\Delta})$ such that P has an opportunity conditional on $h_{t-\Delta}$ and $V_t(h_{t-\Delta})|\hat{m}_t(h_{t-\Delta}) > V_t(H_{t-\Delta}^P, h_{t-\Delta})$.

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 $V_t(H_{t-\Delta}^P, h_{t-\Delta})$ 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. \square

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

Proposition 2. *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 $m_t(H_t^P, h_{t-\Delta}) = f^{R(h_{t-\Delta})}(X_t)$, $a_t(h_{t-\Delta}) = a^{R(h_{t-\Delta})}$, and

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

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 $m_{T-\Delta}(H_{T-\Delta}^P, h_{T-2\Delta}) = f^{R(h_{T-2\Delta})}(X_{T-\Delta})$ for all $H_{T-\Delta}^P$ and $a_{T-\Delta}(H_{T-\Delta}^A, h_{T-2\Delta}) = a^{R(h_{T-2\Delta})}$ for all $H_{T-\Delta}^A$. 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 $\mathcal{E}^* \subset \mathcal{E}_0 \subset \dots \subset \mathcal{E}_{T-\Delta} \subset \mathcal{E}_T$. 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 $\mathcal{E}^* = \mathcal{E}_0$. \square

Despite my conservative approach to removing sequential equilibria, Proposition 2 implies that all the “complex” sequential equilibria involving P trying to keep A in the dark about his own continuation payoff are removed.

Proposition 2 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 2 basically formalizes Remark 1.

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:*

- $\mathcal{M} = \text{Im}(X_t)$ and $m_t(H_t^P, h_{t-\Delta}) = X_t$.
- For each $t < T$ there is a pair of constants $w_t^{\text{salary}}, w_{t+\Delta}^{\text{severance}}$ such that A is paid w_t^{salary} at date t for working and is paid a severance $w_{t+\Delta}^{\text{severance}}$ 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 2 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 \mathcal{C}^Δ . Now for any realization of X_0 change the contract so that after P reports $m_0(X_0)$ the contract randomizes between \mathcal{C}^Δ 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 monotone effect on X_t for all t , then the optimal contract converges to the one characterized in Theorem 1 as $T \rightarrow \infty$.*

B Incentive Compatibility under Infrequent Sampling

Because the infrequent sampling model is a generalization of the original model, the original refined equilibrium concept also needs to be generalized. The same conservative approach as before can be used to develop the notion of maximizing effort incentives in the infrequent sampling model. In the special case when the contract game has P monitor each date, the current model's refined equilibria will reduce to the original model's refined equilibria.

I begin by proving that the current model's monitoring system satisfies something akin to Assumption 1.

Lemma 3. *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 the sequence of efforts $a_{(s,t]}$ from date $s + \Delta$ through date t . There exists a unique function $f^{R, t-s}(Y_t - Y_s)$ taking values in R with the following two properties:*

- *The set $\arg \max_{a_{(s,t]}} e^{-r(t-s)} \mathbf{E}_{a_{(s,t]}} f^{R, t-s}(Y_t - Y_s) - \sum_{t'=s+\Delta}^t e^{-r(t'-s)} h(a_{(s,t]}(t')) \Delta$ contains a maximal element/effort sequence $a^{R, t-s}$.*
- *For any function $g(Y_t - Y_s, \xi)$ taking values in R , if it is not true that $g(Y_t - Y_s, \xi) = f^{R, t-s}(Y_t - Y_s)$ for all $Y_t - Y_s$ and ξ , then $a^{R, t-s}$ is strictly larger than any element of $\arg \max_{a_{(s,t]}} e^{-r(t-s)} \mathbf{E}_{a_{(s,t]}, \xi} g(Y_t - Y_s, \xi) - \sum_{t'=s+\Delta}^t e^{-r(t'-s)} h(a_{(s,t]}(t')) \Delta$.*

$f^{R, t-s}(Y_t - Y_s)$ is characterized by a threshold $\rho^{R, t-s}$ such that

$$f^{R, t-s}(X_t - X_s) = \begin{cases} \max_{w \in R} w & \text{if } Y_t - Y_s > \rho^{R, t-s} \\ \min_{w \in R} w & \text{if } Y_t - Y_s \leq \rho^{R, t-s} \end{cases}$$

Moreover, $\arg \max_{a_{(s,t]}} e^{-r(t-s)} \mathbf{E}_{a_{(s,t]}} f^{R, t-s}(Y_t - Y_s) - \sum_{t'=s+\Delta}^t e^{-r(t'-s)} h(a_{(s,t]}(t')) \Delta$ generically has a single element. When the set does not have a single element, decreasing $\min_{w \in R} w$ infinitesimally will cause the set to have a single element that is infinitesimally close to $a^{R, t-s}$.

Proof. Without loss of generality, assume the minimum and maximum values of R are 0 and 1, and assume $s = -\Delta$. Let \mathcal{F} denote the set of all function f of the form

$$f(X_t) = \begin{cases} 1 & \text{if } Y_t - Y_s > \rho \\ 0 & \text{if } Y_t - Y_s \leq \rho \end{cases}$$

for some ρ .

Let $g(Y_t, \xi)$ be any function taking values in R that is not completely independent of ξ . Let $f(Y_t)$ be the unique function $\in \mathcal{F}$ such that

$$1 - F_t \left(\rho - \sum_{t'=0}^t a_{[0,t]}^g(t')\Delta \right) = \mathbf{E}_{a_{[0,t]}^g, \xi} g(Y_t, \xi).$$

Here F_t is the cdf of a normal random variable with mean zero and variance t . Let $a_{[0,t]}^f$ be the largest effort sequence induced by f . Then $a_{[0,t]}^f > a_{[0,t]}^g$.

I now prove there exists a unique $f^* \in \mathcal{F}$ such that the largest effort sequence induced by f^* is strictly larger than any effort sequence induced by any other $f \in \mathcal{F}$. The first part of the proof then implies that $f^* = f^{R, t-s}$.

The proof is as follows: Suppose there are two functions f^* and $f \in \mathcal{F}$ such the largest effort sequence induced by f^* equals the largest effort sequence induced by f . I now show that there must be another function $\hat{f} \in \mathcal{F}$ such that largest effort sequence induced by \hat{f} is strictly larger than the largest effort sequence induced by f^* and f :

Let ρ^* and ρ be the thresholds of f^* and f , and let $a_{[0,t]}^f$ be their common largest induced effort sequence. Define $a := \sum_{t'=0}^t a_{[0,t]}^f(t')\Delta$. If $a \geq \rho^*$ and $\geq \rho$, then $\rho^* = \rho$, which is a contradiction. So, without loss of generality, assume $a < \rho$. Then the function \hat{f} with threshold $\hat{\rho} = a$ induces a unique effort sequence that is strictly larger than $a_{[0,t]}^f$. \square

Lemma 3 speaks of higher and lower effort sequences. This is well-defined since first-order conditions imply that given two best response effort sequences $a_{(s,t]}$, $a'_{(s,t]}$, if $a_{(s,t]}(t') > a'_{(s,t]}(t')$ for some $t' \in (s, t]$ then $a_{(s,t]}(t') > a'_{(s,t]}(t')$ for all $t' \in (s, t]$.

I will use Lemma 3 to show that at the end of a monitoring period running from $s + \Delta$ to t the principal's report strategy is characterized by a function $f^{R, t-s}$. This mirrors what I did in the original model except I did not have to worry about the agent having multiple best response efforts. I will use the second half of Lemma 3 to justify assuming that if the agent has multiple best response effort sequences, he will choose the maximal effort sequence $a^{R, t-s}$ which is most preferred by the principal. This is because the second half of Lemma 3 implies that the principal can ensure something close to $a^{R, t-s}$ is the unique best response of the agent by reducing the agent's outside option by ε .

The definitions of belief-free continuation payoff, public continuation payoff, and

belief-free (a, m) are all unchanged.

Definition. Fix a set of sequential equilibria \mathcal{E} and a public history h_s such that P monitored at date s given h_s and t is the next monitoring date. P is said to have an opportunity to maximize effort incentives at the beginning of date $s + \Delta$ given \mathcal{E} and conditional on h_s if, for every h_t compatible with h_s , all $(a, m) \in \mathcal{E}$ are belief-free given h_t and share the same belief-free, public continuation payoff $(W_{t+\Delta}(h_t), V_{t+\Delta}(h_t))$ on date $t + \Delta$ and all subsequent dates that start a monitoring period.

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

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

This definition of an opportunity to maximize effort incentives generalizes the original definition to the current model. As part of my conservative approach, I only give P an opportunity to maximize effort incentives at the beginning of a monitoring period, which is defined to be the time period between consecutive monitoring dates, and only if at the end of the period all elements of \mathcal{E} are basically the same and don't depend on what happens before conditional on the public history. In the special case when the monitoring period is a single date, the current definition of an opportunity reduces to the previous definition. The definition of \mathcal{M}^* is unchanged except it is written as $\mathcal{M}^*(h_s)$ instead of $\mathcal{M}^*(h_{t-\Delta})$ since nothing public occurs during the time interval (s, t) .

The definition of a commitment $\hat{m}_t(h_s)$ is the natural generalization of the original definition – the choice of message at date t depends on H_t^P only up to $Y_t - Y_s$ instead of X_t like before. Again, in the special case when the monitoring period is a single date this definition of a commitment reduces to the original definition. Define $a_{(s,t]}|\hat{m}_t(h_s)$ as the largest best-response effort sequence from date $s + \Delta$ through date t given $\hat{m}_t(h_s)$. It is the analogue of $a_t|\hat{m}_t(h_{t-\Delta})$ from before. Define

$$V_{s+\Delta}(h_s)|\hat{m}_t(h_s) = \mathbf{E}_{a_{(s,t]}|\hat{m}_t(h_s), \hat{m}_t(h_s)} \left[\sum_{s < t' \leq t} e^{-r(t'-s)} (Y_{t'} - w_{t'}(h_s)) + e^{-r(t-s)} V_{t+\Delta}(h_t) \right].$$

$V_{s+\Delta}(h_s)|\hat{m}_t(h_s)$ is the analogue of $V_t(h_{t-\Delta})|\hat{m}_t(h_{t-\Delta})$ from before.

The condition for when a set \mathcal{E} of sequential equilibria maximizes effort incentives is basically the same as before. Lemma 2 continues to hold. The definition of an equilibrium that maximizes effort incentives is unchanged.

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

$(W_t(h_{t-\Delta}), V_t(h_{t-\Delta})) = (w_t(h_{t-\Delta}), -w_t(h_{t-\Delta}))$ if the contract game terminates given $h_{t-\Delta}$. Fix a date $s + \Delta$ that begins a monitoring period ending on date t . Given

$(W_{t+\Delta}(h_t), V_{t+\Delta}(h_t))$ for all h_t succeeding h_s , define

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

Then $m_t(H_t^P, h_s) = f^{R(h_s), t-s}(Y_t - Y_s)$, $a_{(s,t]}(h_s) = a^{R(h_s), t-s}$, and $(W_{s+\Delta}(h_s), V_{s+\Delta}(h_s)) =$

$$\left(\sum_{s < t' \leq t} e^{-r(t'-s)} [w_{t'}(h_s) - h(a^{R(h_s), t-s}(t' - s)) \Delta] + e^{-r(t-s)} \mathbf{E}_{f^{R(h_s), t-s}(Y_t - Y_s)} W_{t+\Delta}(h_t), \right. \\ \left. \sum_{s < t' \leq t} e^{-r(t'-s)} [-w_{t'}(h_s) + \mathbf{E}_{a^{R(h_s), t-s} Y_t}] + e^{-r(t-s)} \mathbf{E}_{f^{R(h_s), t-s}(Y_t - Y_s)} V_{t+\Delta}(h_t) \right).$$

The proof mirrors that of Proposition 2. Proposition 3 generalizes Proposition 2 and, in particular, implies that any stronger refinement that doesn't remove all equilibria from a contract game cannot lower the Pareto-frontier.

Most of the work in proving Theorem 3 has already been done. Given that Proposition 3 implies the refined equilibrium concept of the original model admits a natural generalization to the current model, all but the first part of the theorem follow naturally from Theorem 1. The only thing worth remarking on is the first part which says that in the optimal contract (and more generally all Pareto-optimal contracts), monitoring occurs on non-random dates that are evenly spaced. Suppose after some monitoring date e_i , the next monitoring date e_{i+1} is random. For each realization of e_{i+1} , there is associated to it a continuation surplus starting from date $e_i + \Delta$. Pick the realization of e_{i+1} that generates the largest continuation surplus and change the Pareto-optimal contract so that after e_i , the next monitoring date is always that realization of e_{i+1} . This increases the expected continuation surplus at date $e_i + \Delta$ which can only help increase ex-ante surplus. The evenly spaced aspect of optimal monitoring is a consequence of the infinite time horizon.

B.1 Time Consistency

In generalizing the refined equilibrium concept of the original model to the infrequent sampling model I have so far glossed over an important issue. When a monitoring period lasts more than one date, a commitment should be time consistent. Suppose P makes a commitment at the beginning of the monitoring period and A takes it seriously and begins enacting his best response effort sequence. The fear is that at some point in the middle of the period, P is better off changing her strategy for how to select among her best response reports at the end of the period. If this were the case, P would not be able to commit not to change her commitment which means the commitment isn't really a commitment. Allowing P to make non time consistent commitments and using such commitments to remove sequential equilibria would not be consistent with my conservative approach to removing equilibria. This of course was not an issue previously when a monitoring period was a single date.

Below I develop a stringent condition for a commitment to be considered time-consistent (being as stringent as possible is logically consistent with my conservative approach to removing equilibria) and show that even if P is only allowed to make time-consistent commitments, the same set of equilibria survives. The key is to show that the report strategy characterized by $f^{R, t-s}$ is time-consistent.

To define the notion of a time consistent commitment $\hat{m}_t(h_s)$, begin by defining, for any $t' \in (s, t)$, $V_{t'+\Delta}(h_s)|\hat{m}_t(h_s)$. It is the natural extension of $V_{s+\Delta}(h_s)|\hat{m}_t(h_s)$ to date $t' + \Delta$. For any $t' \in (s, t)$, define $a_{(s,t')|\hat{m}_t(h_s)}$ to be the portion of $a_{(s,t]|\hat{m}_t(h_s)}$ ranging from date $s + \Delta$ through date t' . Let $\hat{m}'_t(h_s)$ be another commitment. Define $a_{(t',t]|\{\hat{m}'_t(h_s), a_{(s,t')|\hat{m}_t(h_s)}\}}$ to be the largest best-response effort sequence from date $t' + \Delta$ through date t given $\hat{m}'_t(h_s)$ and given that A previously chose $a_{(s,t')|\hat{m}_t(h_s)}$ in the current monitoring period. Define $V_{t'+\Delta}(h_s)|\{\hat{m}'_t(h_s), a_{(s,t')|\hat{m}_t(h_s)}\}$ to be P 's date $t' + \Delta$ continuation payoff under commitment $\hat{m}'_t(h_s)$ assuming A previously chose $a_{(s,t')|\hat{m}_t(h_s)}$ in the current monitoring period.

Definition. A commitment $\hat{m}_t(h_s)$ is time-consistent if there does not exist a date $t' \in (s, t)$ and commitment $\hat{m}'_t(h_s)$ such that $V_{t'+\Delta}(h_s)|\{\hat{m}'_t(h_s), a_{(s,t')|\hat{m}_t(h_s)}\} > V_{t'+\Delta}(h_s)|\hat{m}_t(h_s)$

One can now define a time-consistent version of the refined equilibrium concept by using the original definition except replace commitments with time-consistent commitments.

Lemma 4. Given a contract game the time-consistent equilibria that maximize effort incentives coincide with the equilibria that maximize effort incentives.

Proof. Fix a set \mathcal{E} of sequential equilibria and suppose P has an opportunity conditional on some h_s . It suffices to show that the commitment characterized by $f^{R(h_s), t-s}$ is time-consistent.

Without loss of generality, assume $s = -\Delta$ and the minimum and maximum values of $R(h_s)$ are 0 and 1. Since h_s is empty, I drop all mention of h_s below.

Suppose not. Then there exists a date $t^* < t$ and commitment m_t^* such that $V_{t^*+\Delta}|\{m_t^*, a_{[s,t^*]}|m_t\} > V_{t^*+\Delta}|m_t$. The proof of Lemma 3 implies without loss of generality, m_t^* viewed as a function of Y_t is an element of \mathcal{F} with some threshold ρ^* .

I now show that there is a commitment that induces a higher effort sequence than m_t which is a contradiction. Let $a_{[0,t]}^*$ be the largest effort sequence induced by m_t^* . Define F_t to be the cdf of a normal random variable with mean zero and variance t . Create the equation

$$\eta(x) := 1 - F_t \left(\rho^* - \sum_{t'=0}^t a_{[0,t]}^*(t')\Delta - x \right)$$

η is an increasing logistic-shaped function with a convex lower half and a concave

upper half. Since $a_{[0,t]}^*$ is a best-response to ρ^* , it must be that

$$\eta'(0) = e^{r(t-t')}h'(a_{[0,t]}^*(t')) \quad \forall t' \in [0, t] \quad (3)$$

The proof of Lemma 3 implies it is without loss of generality to assume 0 is in the domain of the concave upper half of $\eta(x)$.

The first-order conditions for effort imply that either $a_{[0,t]} \leq a_{[0,t]}^*$ or $a_{[0,t]} > a_{[0,t]}^*$. First suppose $a_{[0,t]} > a_{[0,t]}^*$. Now consider the scenario where the worker chooses to follow a up to date t^* but then switches to $a_{[0,t]}^*$ starting from date $t^* + \Delta$. At date $t^* + \Delta$, the marginal benefit of the date t' effort at level $a_{[0,t]}^*$ is $e^{-r(t-t^*)}\eta'(\sum_{k=0}^{t^*}(a_{[0,t]}(k) - a_{[0,t]}^*(k))\Delta)\Delta$ for any $t' \in [t^* + \Delta, t]$. The marginal cost is $e^{-r(t'-t^*)}h'(a_{[0,t]}^*(t'))\Delta$.

Since $\sum_{k=0}^{t^*}(a_{[0,t]}(k) - a_{[0,t]}^*(k))\Delta > 0$ and $\eta(0)$ is in the concave upper half, it must be that $\eta'(\sum_{k=0}^{t^*}(a_{[0,t]}(k) - a_{[0,t]}^*(k))\Delta) < \eta'(0)$. Combining this with (3) implies that under the proposed scenario, starting at date t^* , the marginal benefit of effort going forward is less than the marginal cost. This means that $a_{[t^*+\Delta,t]}|\{m_t^*, a_{[0,t^*]}|m_t\}$ must be smaller than $a_{[0,t]}^*$ on the interval $[t^* + \Delta, t]$. By assumption, this implies that $a_{[t^*+\Delta,t]}|\{m_t^*, a_{[0,t^*]}|m_t\} < a_{[0,t]}^* < a_{[0,t]}$ on the interval $[t^* + \Delta, t]$. Contradiction. So $a_{[0,t]}^* \geq a_{[0,t]}$. But $a_{[0,t]}^* \neq a_{[0,t]}$, so $a_{[0,t]}^* > a_{[0,t]}$. \square

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* Vol. 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] Shapiro, C. and J.E. Stiglitz (1984) “Equilibrium Unemployment as a Worker Discipline Device,” *American Economic Review* Vol. 74, pp. 433-444
- [23] Townsend, R. (1979) “Optimal Contracts and Competitive Markets with Costly State Verification,” *Journal of Economic Theory* Vol. 21, pp. 265-293

- [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