Efficiency in General Agency Models with Imperfect Public Monitoring

Anqi Li*

December 2014

Abstract

In T-period agency relationships between a risk-neutral principal and a risk-averse agent where signals can depend on past actions and exhibit serial correlation, near-efficiency obtains when T is large if the monitoring technology satisfies two basic properties: *concentration of measure* and *informativeness.* The tension between these conditions is used to determine the boundary at which asymptotic efficiency does and does not obtain in agency models with frequent actions. Results deepen and extend our understanding of varying efficiency results in the agency literature, quantify the value of knowing details of the monitoring technology and help solve incentive issues when the monitoring technology is highly persistent.

*Department of Economics, Washington University in St. Louis. angellianqi@gmail.com. I thank Ilya Segal for inspiration, guidance and support. I also thank Marinho Bertanha, Aislinn Bohren, George-Levi Gayle, Johannes Hörner, SangMok Lee, Huiyu Li, Andy Skrzypacz, Juuso Toikka, Yiqing Xing, Takuro Yamashita and workshop participants at Stanford and SAET 2014 for helpful discussions. This work is very preliminary. Comments and suggestions are highly welcomed.

1 Introduction

What basic properties of the monitoring technology help sustain near-efficient outcomes in agency relationships? What is the value of identifying detailed features of the monitoring technology such as how past actions persistently affect future signals and how signals are themselves correlated? While these questions have significant implications for practical issues such as what information should be included in the performance evaluation of employees whose contribution takes time to realize and how social media can help monitor reform plans which take long to generate breakthroughs, they have not been examined systematically by existing studies in the agency literature, most of which focus on individual cases where near-efficiency does or does not obtain and do not handle the persistence of the monitoring technology.

The current paper addresses these questions in a T-period agency model between a risk-neutral principal and a risk-averse agent where signals can depend on past actions and exhibit serial correlation. The main result attributes the implementability of near-efficiency when T is large to two basic properties of the monitoring technology: *concentration of measure* and *informativeness*, but not to other technological details. The tension between these properties is used to determine the boundary at which asymptotic efficiency does and does not obtain in agency models with frequent actions. Results deepen and extend our understanding of varying efficiency results in the agency literature, quantify the value of knowing details of the monitoring technology and help solve incentive issues in the presence of technological persistence.

The key step in estalibshing the efficiency result is to examine a simple contract, *test contract*, which delivers a fixed consumption to the agent in the first $T-1$ period and penalizes the agent severely in the last period if the sample value of a performance test statistic is well below its mean at the efficient action profile — see Figure 1 for a graphically illustration. Theorem 1 shows that this contract is near-efficient when T is large so long as the monitoring technology satisfies the aforementioned properties.

Two new insights play central roles in characterizing the performance of test contract. Intuitively, the contract is near-efficient if it provides a near-optimal incentive and almost full insurance to the agent at the same time. These goals can be achieved simultaneously if we manage to bound the probability of mistakenly passing (failing) the agent when he shirks (works) a lot and to limit his gain from fine-tuning the action choice with past actions and signals despite that the equilibrium strategy is difficult to

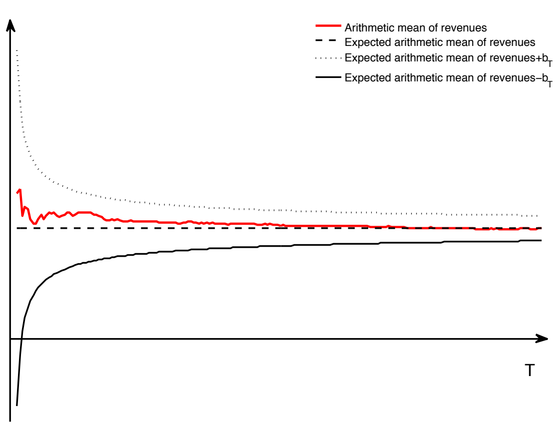

Figure 1 Mean of revenues: $\|\Omega_T\|/T \sim o\left(T^{-\frac{1}{2}}\right)$

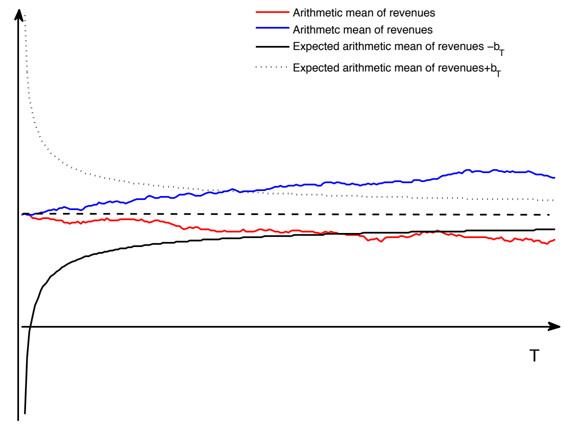

Figure 2 Mean of revenues: $\|\Omega_T\|/T \sim \Theta\left(T^{-\frac{1}{2}}\right)$

characterize directly. *Concentration of measure*, a robust property of a large class of test statistics under general signal processes, says that there is a uniform lower bound for the probability of the event where the sample test statistic is highly concentrated around its mean over all feasible data generating action profiles which converges to one as T grows to infinity. The uniformity of the concentration bound allows us to draw robust inferences about the equilibrium probability of committing Type I and II errors without meticulous calculations of the equilibrium strategy which is difficult to characterize directly because of the persistence of the monitoring technology. Meanwhile, *informativeness* says that if the expected value of the test statistic at the true action profile is close to its counterpart at the efficient action profile, then the true action profile improves the principal's payoff significantly and incurs a nontrivial action cost to the agent. This condition is merely a local restriction on the first moment of the test statistic and hence is much weaker than the usual likelihood ratio based assumptions in the agency literature. It is satisfied by most existing studies and holds automatically in some cases because of the interesting properties of other model primitives.

The reason why concentration of measure and informativeness are sufficient for asymptotic near-efficiency is straightforward. Intuitively, the equilibrium consists of two types of outcomes depending on whether the sample test statistic is concentrated around its mean or not. Concentration of measure says that the second type of outcome is rare however complex the equilibrium strategy is whereas informativeness implies that at the first type of outcome, since the agent gains little from fine-tuning the action choice with past actions and signals, a severe enough penalty suffices to deter persistent deviations from the efficient action profile.

Our sufficient conditions seem to be minimal as they help identify the boundary at which asymptotic near-efficiency does and does not obtain in agency problems with frequent actions. Specifically, consider the typical setting of discrete-time approximations of continuous-time agency models with Brownian motion signals (e.g., Hellwig and Schmidt (2002), Biais et al. (2007) and Sadzik and Stacchetti (2013)) where the horizon is a unit time interval that is divided into T subintervals and the noisiness of the signal per subinterval measured by $\|\Omega_T\|/T = (\sup \Omega_T - \inf \Omega_T)/T$ matches its counterpart under Brownian motion $\Theta\left(T^{-\frac{1}{2}}\right)$.[1] Exactly at this noise level, many interesting test statistics cannot have sufficiently concentrated measures and be informative enough simultaneously, and it is shown by previous authors that even the optimal contract is typically inefficient when T is large. This is not a coincidence. Since Brownian motion process has poor concentration properties, it is typical for the test statistic to drift away from its mean as time proceeds. As a result, we need b_T to be large in order to make the concentration bound tight. But then we lose informativeness because the agent can deviate to inferior action profiles but still pass the test easily — see Figure 2 for a graphical illustration. In contrast, if signals follow time-inhomogeneous Poisson processes and delivers either good news or bad news about the true effort, i.e., $\Omega_T = \{0, 1\}$ (e.g., Radner (1985), Biais et al. (2010), Myerson (2010) and the example of Section 2), or more generally if $\|\Omega_T\|/T \sim o\left(T^{-\frac{1}{2}}\right)$ is slightly smaller than its counterpart under Brownian motion processes, then it is straightforward to construct test statistics that satisfy our sufficient conditions even in the presence of technological persistence. In these situations, the principal's gain from specifying details of the monitoring technology decays very fast as the number of interactions increases.

The result is useful for practitioners in several ways. First, it distinguishes contractual frictions that arise from the monitoring technology itself from those that are due to other modeling or practical considerations. Thus, it helps practitioners understand the very nature of the frictions that they face and offer guidance to the contracting techniques that they should use. Second, it quantifies the principal's gain from knowing details of the monitoring technology and helps decide if such invest-

[1] Throughout, use $\mathcal{O}(\cdot)$, $o(\cdot)$ and $\Theta(\cdot)$ to denote "at most the order of", "smaller than the order of" and "exactly the order of", respectively.

ment is worthwhile. Finally, it provides a toolbox for solving incentive issues in a large class of environments with technological persistence. In this way, it sheds light on practical issues such as what information should be included in employee evaluations (e.g., Poisson signals such as breakthroughs or Brownian motion signals such as noisy cash flows) and how to understand the role of social media in monitoring reform plans especially when the agent's contribution takes long to realize and when success (failure) begets success (failure).

1.1 Related Literature

Efficiency in agency models and mechanisms The current analysis deepens and extends our understanding of varying efficiency results in the agency literature, including Mirrlees (1974), Müller (1997), Rubinstein and Yaari (1983), Radner (1981) and Radner (1981), among many others. In static contracting problems where the agent's utility function is unbounded from below, the first order approach is valid and the likelihood ratio becomes infinitely negative at the lower bound of the signal space, Mirrlees (1974) and subsequently Müller (1997) show that a two-step scheme that pays a fixed consumption most of the time and penalizes only terrible outcomes is near-efficient. Concentration of measure, which requires persistent deviations from the target action profile to cause drastic changes in the tail distribution of the performance test statistic, can be thought of as a dynamic analog of their likelihood ratio condition, though it is a more common and robust feature of the monitoring technology that holds uniformly over a large class of test statistics.

The result suggests that the insight of Radner (1981) and Radner (1985) carries over even if the monitoring technology exhibits persistence and signals are much noisier than the author assumes. Radner (1981) considers a finite-horizon agency game where signals depend only on the concurrent action and take value in a uniformly bounded space that is independent of the number of interactions. In that setting, he examines a different question of whether exerting high effort all the time until the mean of signals falls short of a threshold for the first time is ex-ante ε-optimal when the horizon is long. And his construction makes use of point-wise convergence theorems (e.g., the Law of the iterated logarithm) which are silent about the relative convergence rate of the test statistic at varying action profiles and hence cannot be used immediately for equilibrium characterizations especially when action choices can depend heavily on past actions and signals. Radner (1985) uses the analysis of Radner (1981) as a building block to construct an equilibrium in an infinite-horizon agency game with a high discount factor under the assumption that signals depend only on the concurrent action. The current work shows that transitory technology and uniformly bounded signal space are superfluous for attaining asymptotic near-efficiency in his setup.

Several authors have demonstrated how to attain asymptotic near-efficiency by matching the empirical and theoretical distributions of preference types in multi-agent screening problems without monetary transfers, including Jackson and Sonnenschein (2007) and Escobar and Toikka (2013), among others.[2] In these studies, the bounds on the empirical distribution of reported types are meant to overcome the challenge due to the lack of monetary transfers. Furthermore, preference types are assumed to follow exogenous i.i.d. or Markov processes and details of the type processes are being exploited to obtain equilibrium characterizations. In contrast, I consider a single-agent moral hazard problem with monetary transfers where signals depend endogenously on the agent's actions, identify basic properties of the monitoring technology that help sustain asymptotic near-efficiency and use concentration bounds to draw robust inferences about the equilibrium outcome.

Contracting with frequent actions Several authors, including Hellwig and Schmidt (2002), Biais et al. (2007) and Sadzik and Stacchetti (2013), have considered the discrete-time analog of continuous-time agency models with Brownian motion signals. All these studies match the noisiness of the signal per time interval of length $1/T$ the to its counterpart under Brownian motion. The most relevant to the current work is Sadzik and Stacchetti (2013) which shows that at this noise level, the optimal contract is typically inefficient even when T is large in general infinite-horizon agency models with transitory technologies. The current analysis shows that exactly at this a noise level, many heuristic test statistics cannot have sufficiently concentrated measures and be enough informative at the same time. This suggests using the tension

[2] Jackson and Sonnenschein (2007) shows that if players simultaneously observe their i.i.d. preference types over many replicas of the same social choice problem and must announce types whose empirical distribution fits exactly the theoretical distribution of true types, then Pareto-efficiency is virtually implementable when the number of problems is sufficiently large. Escobar and Toikka (2013) extends Jackson and Sonnenschein (2007) to a dynamic setting with AR(1) types and shows that when the horizon is long but finite, reporting as truthfully as possible is an ε-equilibrium of a game where the conditional distribution of announced types given other players' announcements must be bounded tightly around the theoretical distribution of true types given others' true types.

between basic properties of the monitoring technology to identify the boundary at which asymptotic efficiency does and does not obtain.

Contracting with persistent monitoring technologies The uniformity of the concentration bound dispenses us with the burden of knowing signal processes at deviation action profiles. This feature distinguishes the test contract from standard incentive contracts which vary the payment with the outcome of a likelihood ratio test that compares the signal distribution at the target action profile versus deviation action profiles. While the latter method yields useful characterizations of the optimal policy in certain environments (e.g., static settings (Laffont and Tirole (1993)) or dynamic settings where actions affect only the concurrent outcome (Spear and Srivastava (1987), Sannikov (2008)), it becomes restrictive when signals can depend on past actions because then the principal has to keep track of the action history in order to construct the relevant likelihood ratio test, a problem which quickly becomes intractable when the monitoring technology is highly persistent. Test contract circumvents this dimensionality issue because the concentration bound holds uniformly over all feasible actions profiles. Sannikov (2013) examines contracting with persistent technologies by allowing actions to affect the mean of future revenues. He assumes that signals follow Brownian motion processes and makes detailed parametric assumptions whereas the current analysis does exactly the opposite.

The rest of the paper proceeds as follows: Section 2 presents a simple example; Section 3 lays down the statistical background; Section 4.1 describes the baseline model (Section 4.1); Section 4.2 states the main result; Section 4.3 illustrates the theoretical and practical implications of the main result in special cases; Section 5 considers an extension of the baseline model; Section 6 concludes. See Appendix A for simulation results, Appendix B for omitted proofs and the online appendix for robustness properties of the test contract.

2 An Example

This section uses a simple example to illustrate the key idea behind our main result. A risk-neutral principal and a risk-averse agent interact over finite T instances without discounting. They face zero outside options at the outset when they sign a binding contract. At each instant $t = 1, \cdots, T$, the agent exerts high or low effort $a_t \in \mathcal{A} = \{0, 1\}$ which costs him $c(1) = c > c(0) = 0$ and generates a random revenue X_t that is defined on $\Omega_T = \{H, L\}$ and can depend on the past only through the effort history in an *arbitrary* manner. The agent's payoff equals the utility of consumption minus the effort cost $\sum_{t=1}^{T} u(\psi_t) - c(a_t)$ where the flow utility function $u(\cdot)$ is increasing, concave and unbounded from below. The principal's profit equals the revenue minus the consumption payment $\sum_{t=1}^{T} X_t - \psi_t$. The high effort profile is efficient and yields a per-period expected revenue w. In the complete information benchmark, it is optimal for the principal to elicit this effort profile, pay a fixed consumption $u^{-1}(c)$ to the agent and earn a per-period expected profit $w - u^{-1}(c)$. The current analysis concerns how to approximate this profit level as $T \to \infty$ when the principal observes only the revenue for arbitrary technologies that satisfy the aforementioned assumptions.

Consider the *average-revenue test contract* which delivers a fixed consumption ψ to the agent at each date $t = 1, 2, \cdots, T$ and tests whether the arithmetic mean of realized revenues $\varphi(\mathbf{X}) = \frac{1}{T}\sum_{t=1}^{T} X_t$ is above or below a threshold $w - b_T$ at date T where $b_T = T^{-\frac{1}{2}+\varepsilon}$ for some arbitrary $\varepsilon \in \left(0, \frac{1}{2}\right)$. If the result is affirmative, then the agent passes the test and earns the same consumption ψ as before. Otherwise he fails the test and earns a lower consumption $\underline{\psi}$ which causes him a utility loss $u(\psi) - u(\underline{\psi}) = \alpha c T$ for some arbitrary $\alpha > 1$.

Albeit simple, this contract allows the principal to extract almost full surplus when T is large. Intuitively, near-efficiency obtains if the contract provides a near-optimal incentive and almost full insurance at the same time. In the current context, these two goals can be achieved simultaneously if the contract effectively deters two types of deviations: passing the test by sheer luck or by manipulating the test outcome via fine-tuning the effort choice with past efforts and revenues. Fortunately, *concentration of measure* says that the first type of deviation is rare whereas *informativeness* suggests that the gain from the second type of deviation is limited.

First, consider the event where the mean of revenues differs from its expected value by more than b_T. Since revenues are bounded independent random variables for each given effort profile, it follows from Hoeffding (1963)'s concentration inequality for sums of bounded independent random variables that the probability of this event is bounded uniformly from above by $\mu_T = 2\exp\left(-\frac{2Tb_T^2}{(H-L)^2}\right)$ over all feasible effort profiles. The uniformity of the concentration bound allows us to bound the equilibrium probability of this event by μ_T, too, without going through meticulous calculations of the equilibrium strategy which is difficult to characterize directly because the effort choice can vary subtly with the history of efforts and revenues.

Second, suppose that the mean of revenues is b_T-concentrated around its expected value and the agent passes the test, i.e., $\mathbb{E}[\frac{1}{T}\sum_{t=1}^{T} X_t|a^T] + b_T \geq \frac{1}{T}\sum_{t=1}^{T} X_t \geq w - b_T$. By assumption, we have $\mathbb{E}[\frac{1}{T}\sum_{t=1}^{T} X_t|a^T] \geq w - 2b_T$. Then the *informativeness* of the revenue process which, in the current example, follows from the very definition of efficiency, implies that $\frac{1}{T}\sum_{t=1}^{T} c(a_t) \geq u \circ (u^{-1}(c) - 2b_T)$. That is, the agent cannot save much effort cost however carefully he adjusts the effort choice with past efforts and revenues. The reason is simple: if the contrary is true, then there is a much cheaper way to generate more or less the same revenue as by exerting the efficient effort profile, a contradiction.

The uniformity of the probability and payoff bounds leads to an analytic lower bound for the equilibrium profitability of the average-revenue test contract for each $T \in \mathbb{N}$ which converges to $w - u^{-1}(c)$ as $T \to \infty$. Intuitively, the equilibrium consists of two types of outcomes depending on whether the mean of revenues is b_T-concentrated around its expected value or not. The probability of the second type of outcome is at most μ_T and hence is negligible when T is large. At the first type of outcome, the agent either passes the test and spends approximately $u \circ (u^{-1}(c) - 2b_T)$ per-period on the effort cost or fails the test and pays a per-period penalty αc, and it is clearly optimal for her to pass the test most of the time. Indeed, we can quantify the equilibrium probability of failure based on the observation that a lower bound for the agent's equilibrium loss per-period $\pi_T \cdot \alpha c + (1 - \mu_T - \pi_T) \cdot u \circ (u^{-1}(c) - 2b_T)$ where $\pi_T = \mathbb{P}\{\text{Failure at the concentration event}\}$ is weakly smaller than an upper bound for his per-period loss at the efficient action profile $c + \frac{1}{2}\mu_T \cdot \alpha c$. Straightforward algebra yields the result.

3 Statistical Background

Concentration Inequality Concentration of measure is fairly general phenomenon which, roughly speaking, says that a well-behaved function defined on a high dimensional probability space almost always takes values that are close to its mean. Concentration inequalities prescribe for each sample size $T \in \mathbb{N}$ and $\varepsilon > 0$ a uniform upper bound for the probability that a large class of well-behaved functions differ from their expected values by more than ε. Until recently, researchers have developed varying methods of obtaining concentration bounds. Interested readers should consult Boucheron et al. (2003) and Kontorovich and Ramanan (2008) for thorough reviews on existing results and methodologies. Most examples considered below make use of McDiarmid (1989)'s concentration inequality for independent (but not necessarily identical) random variables and Kontorovich and Ramanan (2008)'s concentration inequality for Markov chains with bounded contraction coefficients.

McDiarmid (1989) Fix $T \in \mathbb{N}$ and let $\mathbf{X} = (X_1, X_2, \cdots, X_t)$ be T independent random variables defined on a sample space Ω. If a test statistic $\varphi : \Omega^T \to \mathbb{R}$ has *bounded differences*, i.e.,

$$\left|\varphi(x_1, \cdots, x_t, \cdots x_T) - \varphi(x_1, \cdots, x_t', \cdots, x_T)\right| \leq \gamma_t, \forall x_t, x_t', x_s, s \neq t \tag{3.1}$$

for some $\gamma_1, \cdots, \gamma_T > 0$, then for any $\varepsilon > 0$, we have

$$\mathbb{P}\left\{\varphi(\mathbf{X}) - \mathbb{E}\varphi(\mathbf{X}) \geq \varepsilon\right\} \leq \exp\left(-\frac{2\varepsilon^2}{\sum_{t=1}^T \gamma_t^2}\right) \tag{3.2}$$

and

$$\mathbb{P}\left\{\varphi(\mathbf{X}) - \mathbb{E}\varphi(\mathbf{X}) \leq -\varepsilon\right\} \leq \exp\left(-\frac{2\varepsilon^2}{\sum_{t=1}^T \gamma_t^2}\right) \tag{3.3}$$

Remark 1. If $\Omega \subset \mathbb{R}$ is bounded and $\varphi(\mathbf{X}) = \frac{1}{T}\sum_{t=1}^T X_t$, then the right hand side of (3.2) and (3.3) equals $\exp\left(-\frac{2T\varepsilon^2}{\|\Omega\|^2}\right)$ where $\|\Omega\| = \sup \Omega - \inf \Omega$.

Remark 2. If $\Omega \subset \mathbb{R}$ is bounded and $\varphi(\mathbf{X}) = \frac{\sum_{t=1}^T \exp\left(-\frac{rt}{T}\right) X_t}{\sum_{t=1}^T \exp\left(-\frac{rt}{T}\right)}$ for some $r > 0$, then the right hand side of (3.2) and (3.3) equals approximately $\exp\left(\frac{-2T\varepsilon^2 g(r)}{\|\Omega\|^2}\right)$ for some $g(r)$ that depends only on r when T is large.

Remark 3. If $\Omega \subset \mathbb{R}$ is bounded and $\varphi(\mathbf{X}) = \frac{1-\delta}{1-\delta^T}\sum_{t=1}^T \delta^{t-1} X_t$ for some $\delta \in (0, 1)$, then the right hand side of (3.2) and (3.3) equals $\exp\left(-\frac{2\varepsilon^2(1-\delta^T)(1+\delta)}{\|\Omega\|^2(1+\delta^T)(1-\delta)}\right)$.

Kontorovich and Ramanan (2008) Fix an arbitrary $T \in \mathbb{N}$ and let $\mathbf{X} = (X_1, X_2, \cdots, X_t)$ be a (possibly inhomogeneous) Markov chain that is defined on a countable sample space Ω and a Markov measure $\mathbb{P}$. Denote the initial distribution, the transition kernels and the contraction coefficients by $p_0(\cdot)$, $\{p_t(\cdot|\cdot)\}_{t=1}^{T-1}$ and

$\{\theta_t\}_{t=1}^{T-1}$, respectively.[3] If $\varphi : \Omega^T \to \mathbb{R}$ is *c-Lipschitz with respect to the normalized Hamming metric* for some $c > 0$,[4] i.e.,

$$\frac{|\varphi(\mathbf{x}) - \varphi(\mathbf{x}')|}{\frac{1}{T}\sum_{t=1}^{T} \mathbb{1}_{x_t \neq x_t'}} \leq c \text{ for all } \mathbf{x}, \mathbf{x}' \in \Omega^T \tag{3.4}$$

then for any $\varepsilon > 0$,

$$\mathbb{P}\left\{\varphi(\mathbf{X}) - \mathbb{E}\varphi(\mathbf{X}) \geq \varepsilon\right\} \leq \exp\left(-\frac{T\varepsilon^2}{2c^2 M_T^2}\right) \tag{3.5}$$

and

$$\mathbb{P}\left\{\varphi(\mathbf{X}) - \mathbb{E}\varphi(\mathbf{X}) \leq -\varepsilon\right\} \leq \exp\left(-\frac{T\varepsilon^2}{2c^2 M_T^2}\right) \tag{3.6}$$

where

$$M_T = \max_{1 \leq t \leq T-1} (1 + \theta_t + \theta_t\theta_{t+1} + \cdots + \theta_t \cdots \theta_{T-1}) \tag{3.7}$$

In particular, if there exists $\theta \in (0,1)$ such that $\theta_t \leq \theta$ for all $t = 1, 2, \cdots, T$, then $M_T \leq \frac{1}{1-\theta}$ and hence (3.5) and (3.6) are bounded from above by $\exp\left(-\frac{(1-\theta)^2 T\varepsilon^2}{2c^2}\right)$.

4 The Model

4.1 Setup

Environment A principal (she) and an agent (he) interact over finite T instances and face zero outside options at the outset when they sign a binding contract. At each instant $t = 1, \cdots, T$, the agent takes an action $a_t \in \mathcal{A} = \mathbb{R}_+$ and earns a consumption $\psi_t \in \mathbb{R}$. His total payoff over T instances is $U(\psi_1, \cdots, \psi_T) - C(a_1, \cdots, a_T)$ where $U : \mathbb{R}^T \to \mathbb{R}$ is the utility of consumption and $C : \mathcal{A}^T \to \mathbb{R}_+$ is the cost of taking actions that satisfies $C(0, \cdots, 0) = 0$. The principal earns an expected payoff $W(a_1, \cdots, a_T)$ from a T-period action profile $a^T = (a_1, \cdots, a_T)$. The cheapest way of awarding the agent a gross utility level V is to pay him a stream of fixed consumptions which costs the principal $\Psi(V)$ for some increasing function $\Psi : \mathbb{R} \to \mathbb{R}$. The current framework

[3] $\mathbb{P}\{(X_1, X_2, \cdots, X_t) = (x_1, x_2, \cdots, x_t)\} = p_0(x_1)\prod_{s=1}^{t-1} p_s(x_{s+1}|x_s)$ for each $t = 1, \cdots, T$ and $(X_1, \cdots, X_t) \in \Omega^t$. $\theta_t = \sup_{x,x' \in \Omega} \|p_t(\cdot|x) - p_t(\cdot|x')\|_{\text{TV}}$ where "TV" stands for the total variation metric.

[4] For example, the function in Remark 1 is $\|\Omega\|$-Lipschitz with respect to the normalized Hamming metric.

nests the example of Section 2 where $W\left(a^T\right) = \mathbb{E}[\sum_{t=1}^T X_t|a^T]$, $C\left(a^T\right) = \sum_{t=1}^T c(a_t)$, $U\left(\psi_1, \cdots, \psi_T\right) = \sum_{t=1}^T u(\psi_t)$ and $\Psi(V) = Tu^{-1}\left(\frac{V}{T}\right)$ but allows for more sophisticated inter-temporal preferences such as discounting and non-separability.

The principal observes neither the agent's action nor her payoff directly. At the end of each instant $t = 1, \cdots, T$, a random signal X_t defined on a measurable space $(\Omega_T, \mathcal{F}_T)$ is publicly realized. Denote by $\left\{X_1, \cdots, X_T : a^T\right\}$ the signal process generated by an arbitrary action profile a^T. Define the *monitoring technology* by the collection of signal processes at varying action profiles:

$$I(T) = \left\{\left\{X_1, \cdots, X_T : a^T\right\} : a^T \in \mathcal{A}^T\right\} \tag{4.1}$$

Assume throughout that $I(T)$ is common knowledge.

A *test statistic* $\varphi_T : \Omega_T^T \to \mathbb{R}$ is a mapping from the sample space to the reals. For an arbitrarily action profile a^T, I say that the sample test statistic $\varphi_T(\mathbf{X})$ generated by a^T is *b_T-concentrated around its mean* for some $b_T > 0$ (denote this event by $\mathcal{E}\left(a^T; b_T, \varphi_T\right)$) if

$$\left|\varphi_T(\mathbf{X}) - \mathbb{E}\left[\varphi_T(\mathbf{X})|a^T\right]\right| < b_T \tag{4.2}$$

Likewise, I say that the sample test statistic generated by a^T is *b_T-semi-concentrated around its mean* for some $b_T > 0$ (denote this event by $\mathcal{E}^-\left(a^T; b_T, \varphi_T\right)$) if

$$\varphi_T(\mathbf{X}) - \mathbb{E}\left[\varphi_T(\mathbf{X})|a^T\right] > -b_T \tag{4.3}$$

Key Assumptions For illustrative purpose, I assume that monitoring has no direct welfare impacts, i.e., $W(\cdot)$ and $C(\cdot)$ depend only on the agent's actions, and restrict attention to the commonly considered case in the agency literature with pure moral hazard where there is a pure action profile maximizes the social surplus (subsequently referred to as the *target action profile*):

Assumption 1. *There exists $a^{*T} = (a^*_{1,T}, \cdots, a^*_{T,T})$ such that*

$$a^{*T} \in \underset{\sigma_T \in \Delta(\mathcal{A}^T)}{\arg\max}\, \mathbb{E}\left[W\left(a^T\right)|\sigma_T\right] - \Psi\left(\mathbb{E}\left[C\left(a^T\right)|\sigma_T\right]\right)$$

These assumptions will be relaxed in Section 5 where I allow monitoring to have direct welfare impacts and consider the virtual implementation of arbitrary target strategies.

Define $W_T^* = W\left(a^{*T}\right), C_T^* = C\left(a^{*T}\right)$ and $\Pi_T^* = W_T^* - \Psi\left(C_T^*\right)$ as the payoff to the principal, the cost to the agent and the social surplus that is generated by a^{*T}, respectively, all of which are assumed to be strictly positive. Let $\underline{W}_T = \inf_{\sigma_T \in \Delta(\mathcal{A}^T)} \left\{\mathbb{E}\left[W\left(a^T\right)|\sigma_T\right]\right\} > -\infty$ and use $\eta_T = \frac{\underline{W}_T}{W_T^*}$ to denote the ratio between the infimum of the principal's payoff over all feasible action profiles and its counterpart at the target action profile.

The monitoring technology satisfies the following properties:

Assumption 2. *There exist* $\{b_T, \varphi_T\}_{T=1}^{\infty}$ *such that*

(i) (Concentration of measure) For each $T \in \mathbb{N}$, $\mathbb{P}\left(\mathcal{E}\left(a^T; b_T, \varphi_T\right)|a^T\right) > 1 - \mu_T$ *and* $\mathbb{P}\left(\mathcal{E}^-\left(a^T; b_T, \varphi_T\right)|a^T\right) > 1 - \mu_T^-$ *for all* $a^T \in \mathcal{A}^T$ *for some* $\mu_T \geq \mu_T^- > 0$ *where* $\lim_{T\to\infty} \mu_T = \lim_{T\to\infty} \mu_T^- = 0$.

(ii) (Informativeness) For each $T \in \mathbb{N}$, *if* $\mathbb{E}\left[\varphi_T(\boldsymbol{X})|a^T\right] \geq \mathbb{E}[\varphi_T(\boldsymbol{X})|a^{*T}] - 2b_T$, *then* $W\left(a^T\right) \geq W_T^*(1 - w_T)$ *and* $C\left(a^T\right) \geq C_T^*(1 - c_T)$ *where* $\lim_{T\to\infty} w_T = \lim_{T\to\infty} c_T = 0$.

Part (i) says that there is a test statistic that is b_T-concentrated around its mean with a probability that is bounded uniformly from below over all feasible action profiles and converges to one as the as the sample size increases. This assumption holds for a wide range of test statistics under general signal processes especially when the correlation between signals at each given action profile is not too strong and individual signals have small and even contributions to the variation of the aggregate test statistic. It substantially generalizes many commonly made assumptions about the monitoring technology in the agency literature by allowing signals to depend arbitrarily on past actions and to exhibit moderate serial correlation.

Example 1 (Conditionally independent signals)**.** If $X_1, \cdots, X_T$ are independent random variables for each $a^T \in \mathcal{A}^T$ and if φ_T has bounded differences defined by (3.1), then one can apply McDiarmid (1989)'s concentration bound prescribed by (3.2) - (3.3). In the example of Section 2, since the revenue space $\Omega_T = \{H, L\}$ is uniformly bounded and hence the contribution of individual revenues to the variation of the test statistic $\varphi_T(\mathbf{X}) = \frac{1}{T}\sum_{t=1}^{T} X_t$ is at most $\frac{H-L}{T}$, it follows that $\mu_T = 2\exp\left(-\frac{2Tb_T^2}{(H-L)^2}\right) = 2\mu_T^-$.

Example 2 (Moderate serial correlation)**.** Suppose that Ω_T is countable and X_t depends only on $a^t, X_{t-1}, \cdots, X_{t-d}$ where $d \in \mathbb{N}$ is independent of T for each $t =$

$1, \cdots, T$. Define $Z_t = (X_t, \cdots X_{t-d+1})$, taking obvious care when dealing with $Z_1, \cdots, Z_d$. Clearly, $\left\{Z_1, \cdots, Z_T : a^T\right\}$ is a Markov chain for each $a^T \in \mathcal{A}^T$. Denote the t^{th} contraction coefficient of this Markov chain by $\theta_t\left(a^T\right)$. Define $M_T\left(a^T\right) = \max_{1 \leq t \leq T-1} 1 + \theta_t\left(a^T\right) + \theta_t\left(a^T\right)\theta_{t+1}\left(a^T\right) + \cdots + \theta_t\left(a^T\right)\cdots\theta_{T-1}\left(a^T\right)$. If $M_T\left(a^T\right) \leq M_T$ for all $a^T \in \mathcal{A}^T$ for some M_T and $\varphi_T : \Omega_T^{dT} \to \mathbb{R}$ is Lipschitz with respect to the normalized Hamming metric defined by (3.4), then one can apply Kontorovich and Ramanan (2008)'s concentration bound prescribed by (3.5) - (3.6).

A sufficient condition for M_T to exist is that $\theta_t(a^T) \leq \theta$ for some $\theta \in (0,1)$ for all t and a^T. That is, the serial correlation is not too strong and is relatively homogeneous across different action profiles. Meanwhile, the Lipschitz condition is satisfied if individual signals have small and even contribution to the variation of φ_T. For example, if Ω_T is bounded and $\varphi_T(\mathbf{Z}) = \frac{1}{T}\sum_{t=1}^{T} X_t$, then φ_T is $\|\Omega_T\|$-Lipschitz with respect to the normalized Hamming metric.

Part (ii) says that if mean of the test statistic at the true action profile is close to its counterpart at the target action profile, then the true action profile significantly improves the principal's payoff and incurs non-trivial cost to the agent. This part is merely a restriction on the first moment of the test statistic and hence is much weaker than the usual density-based on assumptions in the agency literature. In certain cases, it holds automatically because of the interesting properties of other model primitives.

Remark 4. In the example of Section 2, since $\varphi_T(\mathbf{X})$ equals the mean of revenues, $\mathbb{E}[\varphi_T(\mathbf{X})|a^T] \geq w - 2b_T$ implies $W(a^T) = \mathbb{E}[\frac{1}{T}\sum_{t=1}^{T} X_t|a^T] \geq w \cdot \left(1 - \frac{2b_T}{w}\right)$. Furthermore, it follows from the efficiency of the target action profile that $\frac{1}{T}\sum_{t=1}^{T} c(a_t) \geq u \circ (u^{-1}(c) - 2b_T) \approx c \cdot \left(1 - \frac{u'\left(u^{-1}(c)\right)}{c} \cdot 2b_T\right)$. Thus, we have $w_T = \frac{2b_T}{w}$ and $c_T = \frac{u'\left(u^{-1}(c)\right)}{c} \cdot 2b_T$.

Finally, assume that the agent can be effectively penalized by decreases in the last period consumption:

Assumption 3. *There exists $\alpha > 1$ independent of T such that for each $T \in \mathbb{N}$, $U(\psi, \cdots, \psi, \psi) = C_T^*(1 + \alpha\mu_T^-)$ and $U(\psi, \cdots, \psi, \psi) - U(\psi, \cdots, \psi, \psi') = \alpha C_T^*$ for some ψ and ψ'.*

This assumption is not as restrictive as it seems. In particular, if players have future interactions, then the agent can be penalized by changes in the continuation value instead. See Section 4.3 for further discussions on this subject.

4.2 Main Result

A *test contract* $\left(\varphi_T(\cdot), \mathcal{R}_T, \psi_T, \underline{\psi}_T\right)$ consists of (1) a test statistic φ_T, (2) a rejection region

$$\mathcal{R}_T = \left(-\infty, \mathbb{E}[\varphi_T(\mathbf{X})|a^{*T}] - b_T\right] \tag{4.4}$$

and (3) a pair of consumptions $(\psi_T, \underline{\psi}_T)$ where

$$U(\psi_T, \cdots, \psi_T) = C_T^*(1 + \alpha\mu_T^-) \tag{4.5}$$

and

$$U\left(\psi_T, \cdots, \psi_T, \psi_T\right) - U\left(\psi_T, \cdots, \psi_T, \underline{\psi}_T\right) = \alpha C_T^* \tag{4.6}$$

The contract pays a performance-independent consumption ψ_T to the agent at each $t = 1, \cdots, T-1$ and computes the sample test statistic $\varphi_T(\mathbf{X})$ at the end of instant T. If $\varphi_T(\mathbf{X}) \in \mathcal{R}_T^c$, then the agent passes the test and consumes ψ_T as before. Otherwise he fails the test and consumes $\underline{\psi}_T$ instead. The contract induces a dynamic game where the agent's strategy $\sigma_T = \{\sigma_{t,T} : \mathcal{A}^{t-1} \times \Omega^{t-1} \to \Delta(\mathcal{A})\}$ is a collection of history-dependent action choice rules. The solution concept is *Bayesian Nash equilibrium*. The equilibrium strategy σ_T^* is defined by

$$\sigma_T^* \in \arg\min_{\sigma_T} \mathbb{E}\left[C\left(a^T\right) + \mathbb{1}_{\varphi_T(\mathbf{X})\in\mathcal{R}_T} \cdot \alpha C_T^* \middle| \sigma_T\right] \tag{IC}$$

And the contract satisfies the agent's ex-ante participation constraint if

$$U(\psi_T, \cdots, \psi_T) - \mathbb{E}\left[C\left(a^T\right) + \mathbb{1}_{\varphi_T(\mathbf{X})\in\mathcal{R}_T} \cdot \alpha C_T^* \middle| \sigma_T^*\right] \geq 0 \tag{IR}$$

The main result of this paper establishes an analytic lower bound for the profitability of the test contract and shows that this bound converges to the full surplus under mild assumptions about players' preferences and the production technology:

Theorem 1. *Under Assumptions 1 - 3, there exists a Bayesian Nash equilibrium σ_T^* of $\left(\varphi_T, \mathcal{R}_T, \psi_T, \underline{\psi}_T\right)$ where for each $T \in \mathbb{N}$,* *(i) The probability of failure is bounded from above by*

$$\mathbb{P}\left(\varphi_T(\boldsymbol{X}) \in \mathcal{R}_T | \sigma_T^*\right) \leq \mu_T + \frac{\alpha\mu_T^- + \mu_T(1-c_T) + c_T}{\alpha - 1 + c_T}$$

(ii) The ratio between the expected profit and the full surplus is bounded from below by

$$1 - \frac{1}{\Pi_T^*}\left\{ \begin{matrix} W_T^* \cdot \left[w_T + \left(\mu_T + \frac{\alpha\mu_T^- + \mu_T(1-c_T)+c_T}{\alpha-1+c_T} \right)(1 - w_T - \eta_T) \right] \\ +\Psi\left(C_T^* \cdot (1+\alpha\mu_T^-)\right) - \Psi\left(C_T^*\right) \end{matrix} \right\}$$

As $T \to \infty$,

(a) The probability of failure is vanishing: $\mathbb{P}\left(\varphi_T(\boldsymbol{X}) \in \mathcal{R}_T | \sigma_T^*\right) \sim \mathcal{O}(\max\{\mu_T, c_T\})$;

(b) The profit-surplus ratio converges to one if the second term of Part (ii) is vanishing.

The second term of Part (ii) is vanishing if the social surplus is a non-trivial fraction of the principal's payoff, the infimum of the principal's payoff over all feasible action profiles is not too small and the compensation function is not increasing too fast in the agent's utility level such that $\frac{W_T^*}{\Pi_T^*} \sim \Theta(1)$, $\left(\mu_T + \frac{\alpha\mu_T^- + \mu_T(1-c_T)+c_T}{\alpha-1+c_T}\right)\eta_T \to 0$ and $\frac{\Psi\left(C_T^* \cdot (1+\alpha\mu_T^-)\right) - \Psi\left(C_T^*\right)}{\Pi_T^*} \to 0$ as $T \to \infty$. As far as the author is aware, these assumptions are satisfied by most existing studies in the agency literature.

In principle, one can certainly vary the incentive payment with the agent's performance score in a more continuous manner. While such subtlety may not help improve the asymptotic efficiency, it is certainly useful if parties have only a few number of interactions or if the probability and payoff bounds are not as tight as assumed.

4.3 Special Cases

This section illustrates how the tension between our sufficient conditions helps determine the boundary at which asymptotic efficiency does and does not obtain in a large class of agency problems.

Example 3 (Contracting with frequent actions)**.** Time evolves continuously over $[0, 1]$ and players share a common interest rate $r > 0$. The economy indexed by $T \in \mathbb{N}$ divides the unit time interval into T sub-intervals of equal length $\Delta = 1/T$. Over each sub-interval $[(t-1)\Delta, t\Delta]$, $t = 1, \cdots, T$, the agent spends an effort cost $c(a_t)\Delta$, generates a random public revenue $X_t\Delta$ and earns a net payoff $[u(\psi_t) - c(a_t)]\Delta$. For each T-period action profile a^T, define $W(a^T) = \mathbb{E}\left[\frac{\sum_{t=1}^T \exp(-rt\Delta)X_t}{\sum_{t=1}^T \exp(-rt\Delta)}\Big| a^T\right]$ as the per-period expected revenue and $C(a^T) = \frac{\sum_{t=1}^T \exp(-rt\Delta)c(a_t)}{\sum_{t=1}^T \exp(-rt\Delta)}$ as the per-period effort cost. Assume that X_t is defined on a bounded sample space $\Omega_T \subset \mathbb{R}$ and $W_T^*, C_T^* \sim \Theta(1)$ as $T \to \infty$.

Consider the *discounted average-revenue test contract* which tests if $\varphi_T(\mathbf{X}) = \frac{\sum_{t=1}^T \exp(-rt\Delta)X_t}{\sum_{t=1}^T \exp(-rt\Delta)}$ belongs to the corresponding rejection region or not. Since $w_T = \frac{2b_T}{W_T^*}$, $c_T = \frac{u'\left(u^{-1}(C_T^*)\right)}{C_T^*} \cdot 2b_T$ are both $\Theta(b_T)$ whereas $\mu_T \approx 2\exp\left(-\frac{2Tb_T^2 g(r)}{\|\Omega_T\|^2}\right)$ for some $g(r)$ that depends only on r when T is large (see Remarks 2 and 4), Assumption 2 is satisfied if and only if $\{b_T\}_{T=1}^\infty$ and $\{\mu_T\}_{T=1}^\infty$ are both vanishing in T, or

$$\frac{\|\Omega_T\|}{T} \sim o\left(T^{-\frac{1}{2}}\right) \text{ as } T \to \infty \tag{4.7}$$

Condition (4.7) says that the noisiness of the signal per-period does not increase too fast as the period length shrinks. It is satisfied by many theoretical and applied works, including the notable work of Radner (1981) which examines a finite-horizon agency game where the monitoring technology is transitory and the signal space does not expand as the horizon grows. Theorem 1 suggests that these assumptions are superfluous for attaining asymptotic efficiency in his setup. It also suggests that in a large class of environments where asymptotic efficiency does obtain, the value of knowing exact details of the monitoring technology decays very fast as the number of interactions increases.

Example 4 (Contracting with Poisson monitoring technology)**.** In the setup of Example 3, suppose that a binary signal $X_t \in \{0, 1\}$ is publicly realized at each instant $t\Delta, t = 1, \cdots, T$ where $X_t = 1$ conveys good news about the agent's action and occurs with probability $\lambda(a^t)\Delta$ over $[(t-1)\Delta, t\Delta]$. When Δ is small, Bernoulli process generated by a^T can be thought of as a discrete-time analog of a time-inhomogeneous Poisson process with arrival rates $\lambda(a^t), t = 1, \cdots, T$.

Consider a contract that tests the discounted mean of good signals $\varphi_T(\mathbf{X}) = \frac{\sum_{t=1}^T \exp(-rt\Delta)X_t}{\sum_{t=1}^T \exp(-rt\Delta)}$. Since $\|\Omega_T\| = 1$ satisfies Condition (4.7), there exist $\{b_T\}_{T=1}^\infty$ and $\{\mu_T\}_{T=1}^\infty$ that are both vanishing in T and hence Assumption 2 (i) is satisfied. If, in addition, that there exist vanishing sequences $\{w_T, c_T\}_{T=1}^\infty$ such that $\frac{\sum_{t=1}^T \exp(-rt\Delta)\lambda(a^t)}{\sum_{t=1}^T \exp(-rt\Delta)} \geq$

$\frac{\sum_{t=1}^{T}\exp(-rt\Delta)\lambda(a^{*t})}{\sum_{t=1}^{T}\exp(-rt\Delta)} - b_T$ implies $W(a^T) \geq W_T^*(1-w_T)$ and $C(a^T) \geq C_T^*(1-c_T)$, then Assumption 2 (ii) is also satisfied. As far as the author is aware, this is true for most existing studies on contracting with Poisson monitoring technologies, including Biais et al. (2010) and Myerson (2010).

An immediate corollary of this observation is that near-efficiency obtains in infinite-horizon contracting problems with Poisson monitoring technologies where parties share a common interest rate $r > 0$ and the agent has bounded flow utilities provided that we can divide the horizon into many small blocks and confine the technological persistence within individual blocks. Indeed, this observation is made by Biais et al. (2010) and Myerson (2010) for transitory monitoring technologies. To see why, let us fix an arbitrarily small $\varepsilon > 0$ and divide the infinite horizon into many blocks of length L where $\int_0^L e^{-rs}c(a_s^*)ds \ll \varepsilon \int_L^\infty e^{-rs}ds$. That is, from the agent's point of view, the cost of taking the target actions in current block is much smaller than the present value of earning a net payoff ε from the next block onward. Now consider a contract that implements the contract described in the previous paragraph in each block and rewards the agent a net flow payoff ε so long as the relationship continues but switches to no production and the lowest consumption payment when the agent fails the test for the first time. Standard arguments suggest that there is an equilibrium of this contract that yields almost full surplus when interactions are sufficiently frequent.

Existing models with Poisson monitoring technologies all assume additional contractual frictions to make the contract design problem non-trivial. For example, Biais et al. (2010) makes the agent more impatient than the principal whereas Myerson (2010) imposes an exogenous bound on the agent's continuation payoff. The current analysis explains why these modeling assumptions are necessary. It also helps practitioners differentiate the contractual frictions used by these authors from those that arise from the monitoring technology itself.

Example 5 (Discrete-time approximation of models with Brownian motion signals)**.** This example is taken from Sadzik and Stacchetti (2013)'s discrete-time approximation of continuous-time agency models with Brownian motion signals. In the setup of Example 3, suppose that $X_t = a_t + Z_t\sqrt{T}$ where $Z_1, \cdots, Z_T$ are i.i.d. random variables with mean 0 and variance σ^2. By independence, we have $\text{Var}\left(\varphi_T(\mathbf{X})|a^T\right) = \frac{\sum_{t=1}^{T}\exp(-2rt\Delta)\text{Var}(X_t|a_t)}{\left(\sum_{t=1}^{T}\exp(-rt\Delta)\right)^2} \approx \frac{\sigma^2}{g(r)}$ for all $a^T \in \mathcal{A}^T$ when T is large. Now fix any vanishing sequence $\{b_T\}_{T=1}^{\infty}$ that makes both $\{w_T\}_{T=1}^{\infty}$ and $\{c_T\}_{T=1}^{\infty}$ vanishing. Notice that there exists no vanishing sequence $\{\mu_T\}_{T=1}^{\infty}$ that satisfies Assumption 2 (i) because if the contrary is true, then $\frac{\sigma^2}{g(r)} = \text{Var}\left(\varphi_T(\mathbf{X})|a^T\right) \leq (1-\mu_T)b_T^2 + \mu_T \frac{\sigma^2}{g(r)} \to 0$ as $T \to \infty$, a contradiction to the assumption that $\{b_T\}_{T=1}^{\infty}$ and $\{\mu_T\}_{T=1}^{\infty}$ are both vanishing. Intuitively, since Brownian motion tends to drift apart from its mean, we need to keep b_T large in order to keep the concentration bound μ_T small. But then we lose informativeness. See Figure 2 for a graphical illustration.

Indeed, the same is true for Hellwig and Schmidt (2002)'s discrete-time approximation of Holmstrom and Milgrom (1987) where $X_t \in \{x_1\sqrt{T}, \cdots, x_N\sqrt{T}\}$, and for Biais et al. (2007)'s discrete-time analog of DeMarzo and Sannikov (2006) where $X_t \in \{x_1\sqrt{T}, x_2\sqrt{T}\}$. All these authors essentially assume that $\frac{\|\Omega_T\|}{T} \sim \Theta(T^{-\frac{1}{2}})$. Condition (4.7) suggests that *exactly* at this noise level, it is hard to conceive a test statistic that has sufficiently concentrated measures and conveys enough information about the true action at the same time. As a result, asymptotic efficiency may not obtain even if T is large.

This intuition is confirmed by Sadzik and Stacchetti (2013) which examines an infinite horizon agency model where the monitoring technology is transitory and the agent's flow utility is bounded. More precisely, suppose to the contrary that there is a near-efficient test contract when T is large in the setup of Sadzik and Stacchetti (2013). Then the principal can use the trick discussed in the previous example to extract almost full surplus when interactions are sufficiently frequent. But this contradicts with the result of Sadzik and Stacchetti (2013) which shows that the optimal contract is strictly inefficient when T is large so long as the Fisher information metric of Z_t is finite, an assumption that is satisfied by many common distributions.

Example 6 (Fixed discount factor)**.** In the setup of Section 2, suppose that parties discount the future by a common factor $\delta \in (0,1)$. Notice that the discounted average-revenue test contract with $\varphi_T(\mathbf{X}) = \frac{1-\delta}{1-\delta^T}\sum_{t=1}^{T}\delta^{t-1}X_t$ does not satisfy Assumption 2 (i) because the test statistic assigns a disproportionately high weight to early signals such that for each vanishing sequence $\{b_T\}_{T=1}^{\infty}$, we have $\lim_{T\to\infty}\mu_T = \lim_{T\to\infty}\exp\left(-\frac{b_T^2(1-\delta^T)(1+\delta)}{(H-L)^2(1+\delta^T)(1-\delta)}\right) = 1$ (see Remark 3). Nevertheless, the simulation result of Section A suggests that the contract still yields a high profit guarantee for the principal when T is neither too big nor too small under reasonable assumptions about the production technology, discount factor and payoff functions.

5 Extension

This section allows monitoring to have direct welfare impacts and considers the virtual implementation of arbitrary target strategies. In the setup of Section 4, suppose that X_t have full support on Ω_T for all $t = 1, \cdots, T$ and $a^T \in \mathcal{A}^T$. Each profile of actions and realized signals $(a^T, \mathbf{x})$ improves the principal's payoff by $W\left(a^T, \mathbf{x}\right)$ and incurs a total cost $C\left(a^T, \mathbf{x}\right)$ to the agent, where $\inf_{a^T \in \mathcal{A}^T, \mathbf{x} \in \Omega_T^T} W\left(a^T, \mathbf{x}\right) > -\infty$ and $\inf_{a^T \in \mathcal{A}^T, \mathbf{x} \in \Omega_T^T} C\left(a^T, \mathbf{x}\right) = 0 = C(0, \cdots, 0, \mathbf{x})$ for all $\mathbf{x} \in \Omega_T^T$.

Take an arbitrary target strategy $\hat{\sigma}_T$. Define $\hat{W}_T = \mathbb{E}\left[W\left(a^T, \mathbf{X}\right) | \hat{\sigma}_T\right]$, $\hat{C}_T = \mathbb{E}\left[C\left(a^T, \mathbf{X}\right) | \hat{\sigma}_T\right]$ and $\hat{\Pi}_T = \hat{W}_T - \Psi(\hat{C}_T)$ as the expected payoff to the principal, the expected cost to the agent and the expected social surplus that are generated by $\hat{\sigma}_T$, respectively. Define $\eta_T = \frac{\inf_{a^T, \mathbf{x}} W\left(a^T, \mathbf{x}\right)}{\hat{W}_T}$ as the infimum of the principal's payoff over all feasible profiles of actions and signals. Under $\hat{\sigma}_T$, an profile of actions and signals $(\hat{a}^T, \mathbf{x})$ are *mutually consistent* if they occur with a positive probability.

In the remainder of this section, assume that

Assumption 4. *There exist* $\{b_T, \varphi_T\}_{T=1}^{\infty}$ *such that*

(i) (Concentration of measure) For each $T \in \mathbb{N}$, $\mathbb{P}\left(\mathcal{E}\left(a^T; b_T, \varphi_T\right) | a^T\right) > 1 - \mu_T$ *and* $\mathbb{P}\left(\mathcal{E}^-\left(a^T; b_T, \varphi_T\right) | a^T\right) > 1 - \mu_T^-$ *for some* $\mu_T \geq \mu_T^- > 0$ *for all* $a^T \in \mathcal{A}^T$ *where* $\lim_{T\to\infty} \mu_T = 0$.

(ii) (Limited welfare impacts, homogeneity and informativeness) There exist $\{w_T, c_T\}_{T=1}^{\infty}$ *such that for each* T *and* $\left(\hat{a}^T, \boldsymbol{x}\right)$ *that are mutually consistent under* $\hat{\sigma}_T$*, if* $\mathbb{E}\left[\varphi_T\left(\boldsymbol{X}\right) | a^T\right] \geq \mathbb{E}\left[\varphi_T\left(\boldsymbol{X}\right) | \hat{a}^T\right] - 2b_T$*, then* $W\left(a^T, \boldsymbol{x}\right) \geq \hat{W}_T(1 - w_T)$ *and* $C\left(a^T, \boldsymbol{x}\right) \geq \hat{C}_T(1 - c_T)$.

Part (i) is the same as before as it requires the test statistic to be b_T-concentrated around its mean with a uniformly bounded probability that converges to one as T grows to infinity. Part (ii) says that if the mean of the test statistic at the true action profile is close to its counterpart at an arbitrary action profile which occurs with a strictly positive probability under $\hat{\sigma}_T$, then the payoff to the principal and the cost to the agent are bounded uniformly from below over all realizations of public signals that are mutually consistent with this target action profile. This assumption is satisfied if (1) monitoring has limited welfare impacts, (2) the action profiles prescribed by $\hat{\sigma}_T$ are relatively homogeneous and (3) signals are informative about the welfare impacts of true actions.

The modified test contract works as follows. At each instant $t = 1, \cdots, T-1$, the principal pays a fixed consumption ψ_T and recommends a pure action $\hat{a}_t$ based on $\hat{\sigma}_T$, taking the history of recommendations and realized signals $(\hat{a}^{t-1}, x^{t-1})$ as given. At the end of instant T, based on the profile of recommended actions $\hat{a}^T$ and realized signals $\mathbf{x}$, she tests if $\varphi_T(\mathbf{x}) \in \mathcal{R}_T\left(\hat{a}^T\right)$ where

$$\mathcal{R}_T\left(\hat{a}^T\right) = \left(-\infty, \mathbb{E}[\varphi_T(\mathbf{X})|\hat{a}^T] - b_T\right] \tag{5.1}$$

If the result is affirmative, then the agent fails the test and consumes $\underline{\psi}_T$. Otherwise he passes the test and consumes ψ_T. ψ_T and $\underline{\psi}_T$ satisfy $U(\psi_T, \cdots, \psi_T) = \hat{C}_T(1+\alpha\mu_T)$ and $U(\psi_T, \cdots, \psi_T) - U(\psi_T, \cdots, \underline{\psi}_T) = \alpha\hat{C}_T$.

A straightforward extension of Theorem 1 yields the following result:

Corollary 1. *Under Assumptions 4 and an analog of Assumption 3 that replaces C_T^* with $\hat{C}_T$, the test contract $\left(\varphi_T(\cdot), \mathcal{R}_T(\cdot), \psi_T, \underline{\psi}_T\right)$ has a Bayesian Nash equilibrium σ_T^* that satisfies Theorem 1 (i)-(ii) where W_T^*, C_T^* and Π_T^* are replaced by $\hat{W}_T, \hat{C}_T$ and $\hat{\Pi}_T$, respectively.*

The result is easy to understand. First, concentration of measure says that no matter how complex the equilibrium strategy is, one can easily bound the equilibrium probability that the sample test statistic is concentrated around its mean based on the concentration bound. Meanwhile, limited welfare impacts, homogeneity and informativeness imply that if the agent passes the test at the event where the mean of the test statistic at the true action profile is close to its counterpart at the realized recommendation action profile, then the principal's payoff and the agent's cost of taking actions are bounded uniformly from below over all feasible recommendation action profiles. These properties allow us to bound the equilibrium profitability in the same way as before.

6 Conclusion

I conclude by discussing several open questions and suggesting directions for future research. First, the current analysis attributes the implementability of asymptotic near-efficiency in agency models to several basic properties of the monitoring technology but not to other model details. It is interesting to see if the same is true for dynamic games with imperfect public monitoring. In continuous-time games where signals follow Levy processes, Sannikov and Skrzypacz (2010) shows that certain types of Poisson signals can effectively sustain cooperation in public strategy equilibria because they seldom trigger wrongful penalties whereas Brownian motion signals cannot. It is interesting to see if the insight carries over to more general monitoring technologies.

Our sufficient conditions for implementing asymptotic near-efficiency are almost tight at least in a class of agency models with frequent actions. This leaves open the question of whether the same is true in general. It is the author's belief that in order to prove or disprove this conjecture, we need new techniques that yield robust upper bounds for the optimal profit level without fully characterizing the equilibrium strategy in general agency models.

While the analysis so far has focused on the issue of efficiency, it is worth noting that test contract itself has many interesting robustness properties: it depends on few model parameters that are easy to estimate and thus imposes little knowledge burden on the principal; it can encompass varying considerations such as hidden savings, persistent hidden characteristics, limited liability and team production; in practical situations where the principal knows few details about a potentially complex monitoring technology, it offers guidance to what to contract upon and provides a robust profit guarantee that is difficult to improve upon. Interested readers should consult the online appendix for formal analyses.

A Simulation

This section simulates the performance of the test contract in Example 6 for $\Omega_T = \{0,1\}, \varphi_T(\mathbf{X}) = \frac{1-\delta}{1-\delta^T}\sum_{t=1}^T \delta^{t-1}X_t$ and $\mu_T = 2\exp\left(-\frac{b_T^2(1-\delta^T)(1+\delta)}{(H-L)^2(1+\delta^T)(1-\delta)}\right) = 2\mu_T^-$. To make progress, I treat one period as a quarter, apply the standard annual discount rate 5% and take $b_T = T^{-0.05}$. Under these assumptions, consider two cases:

- $\alpha = 2, c_T \in \left[\frac{1}{15}b_T, \frac{1}{8}b_T\right]$, i.e., if the agent passes the test at the event where the sample test statistic is b_T-concentrated around its mean, then the per-period effort cost ranges from 88% to 94% of the target action cost. Notice that $\frac{b_T}{15}$ and $\frac{b_T}{8}$ are very conservative estimates of c_T because in the example of Section 2, $c_T = \frac{u'(u^{-1}(c))\cdot 2b_T}{c} \approx \frac{u'(u^{-1}(12000))\cdot 2b_T}{12000} \ll \frac{1}{15}b_T$ based on a conservative estimate of the quarterly labor cost $c = 12000$. If c_T is indeed close to this number, then the result below obtains when α is slightly greater than one.

- $\alpha = 3$, $c_T \in \left[\frac{1}{4.5}b_T, \frac{1}{7}b_T\right]$, i.e., if the agent passes the test at the event where the realized test statistic is b_T-concentrated around its mean, then the per-period effort cost ranges from 78% to 86% of the target action cost.

Figures A.1 - A.4 plot the upper bound for the equilibrium probability of failure prescribed by Theorem 1 (i) against T. Since the result is well below 11% after 4 periods, it is legitimate to conclude that the contract yields almost full surplus under reasonable assumptions about the production technology (w_T, η_T) and the compensation function $\Psi_T(\cdot)$ when T is neither too big nor too small. The result further implies that the feasibility of penalty — which only needs to be slightly bigger than the cost of taking efficient actions over one year — may not be an issue after all.

B Omitted Proofs

Lemma 1. *Fix an arbitrary $T \in \mathbb{N}$. If $\varphi_T(\boldsymbol{X})$ is b_T-concentrated around its mean at a^T and $\varphi_T(\boldsymbol{X}) \in \mathcal{R}_T^c$, then $W\left(a^T\right) \geq W_T^*(1 - w_T)$ and $C\left(a^T\right) \geq C_T^*(1 - c_T)$.*

Proof. By assumption, we have (1) $\varphi_T(\mathbf{X}) > \mathbb{E}[\varphi_T(\mathbf{X})|a^{*T}] - b_T$ and (2) $\left|\varphi_T(\mathbf{X}) - \mathbb{E}[\varphi_T(\mathbf{X})|a^T]\right| < b_T$, or $\mathbb{E}[\varphi_T(\mathbf{X})|a^T] > \mathbb{E}[\varphi_T(\mathbf{X})|a^{*T}] - 2b_T$. The result follows from Assumption 2. □

Proof of Theorem 1

Proof. Define

$$r_T = \int \mathbb{P}\left(\mathcal{E}\left(a^T; b_T, \varphi_T\right) | a^T\right) dF\left(a^T | \sigma_T^*\right) \tag{B.1}$$

as the equilibrium probability that the realized test statistic is b_T-concentrated around its mean, and

$$\pi_T = \int \mathbb{P}\left(\mathcal{E}\left(a^T; b_T, \varphi_T\right), \varphi_T(\mathbf{X}) \in \mathcal{R}_T | a^T\right) dF\left(a^T | \sigma_T^*\right) \tag{B.2}$$

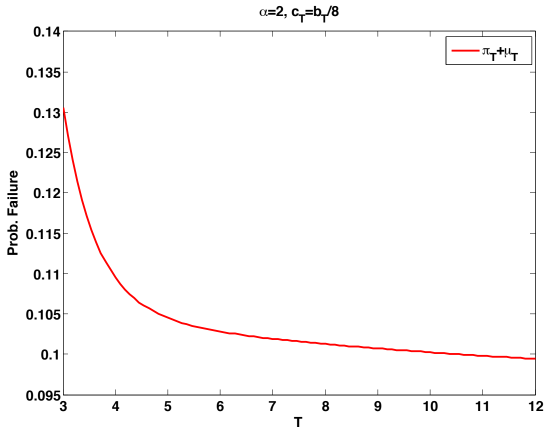

Figure A.1 $\alpha = 2, c_T = \frac{b_T}{8}$

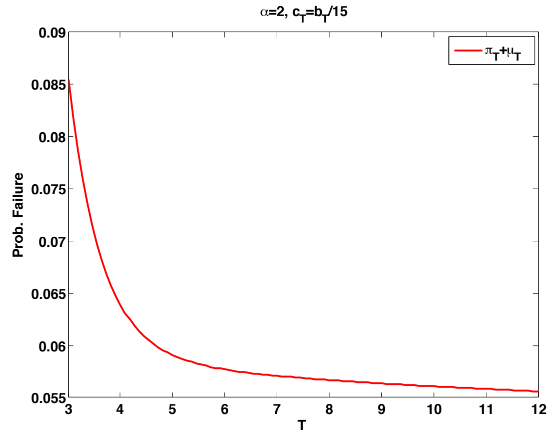

Figure A.2 $\alpha = 2, c_T = \frac{b_T}{15}$

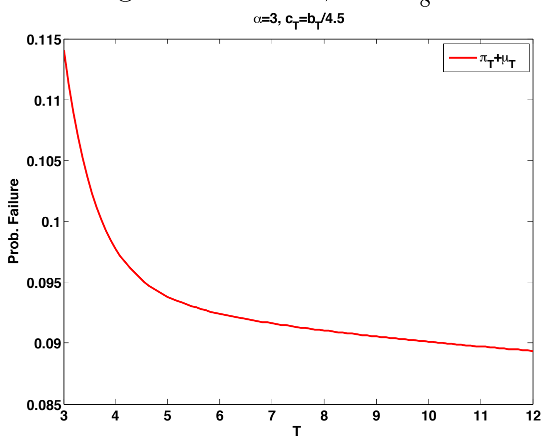

Figure A.3 $\alpha = 3, c_T = \frac{b_T}{4.5}$

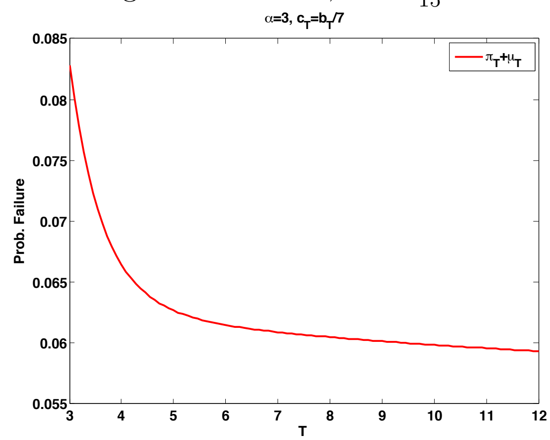

Figure A.4 $\alpha = 3, c_T = \frac{b_T}{7}$

as the equilibrium probability that the test statistic is b_T-concentrated around its mean and yet the agent fails the test. By Assumption 2 (i), we have

$$r_T \geq 1 - \mu_T, \forall T \tag{B.3}$$

Parts (i) and (a): First, bound the agent's equilibrium payoff from above by

$$U(\psi_T, \cdots, \psi_T) - \left\{ \pi_T \cdot \underbrace{\alpha C_T^*}_{(1)} + (r_T - \pi_T) \underbrace{C_T^*(1 - c_T)}_{(2)} + (1 - r_T) \cdot \underbrace{0}_{(3)} \right\}$$

where (1) is attained if the agent pays the penalty for failing the test yet incurs no cost, (2) is attained if he passes the test when the sample test statistic is b_T-concentrated around its mean and thus spends at least $C_T^*(1 - c_T)$ on the action cost, and (3) is attained if he takes the lowest action and pays no penalty at the irregular event. Second, notice that if the agent takes a^{*T} instead, then he incurs an action cost C_T^* but fails the test only if the test statistic is not b_T-semi-concentrated around its mean, an event whose probability is bounded from above by μ_T^-. Thus, his expected payoff is bounded from below by

$$U(\psi_T, \cdots, \psi_T) - \left\{C_T^* + \mu_T^- \cdot \alpha C_T^*\right\}$$

Since the first expression is weakly greater the second one, it follows that

$$\pi_T \leq \frac{\alpha\mu_T^- + \mu_T(1 - c_T) + c_T}{\alpha - 1 + c_T}, \forall T$$

Based on this result, bound the equilibrium probability of failure from below by

$$\mathbb{P}\left(\varphi_T(\mathbf{X}) \in \mathcal{R}_T | \sigma_T^*\right) \leq \pi_T + 1 - r_T \leq \mu_T + \frac{\alpha\mu_T^- + \mu_T(1 - c_T) + c_T}{\alpha - 1 + c_T}, \forall T$$

where the RHS is $\mathcal{O}(\max\{c_T, \mu_T\})$ when T is large.

Parts (ii) and (b): By Part (i), it suffices to pay $U(\psi_T, \cdots, \psi_T) \geq C_T^*(1 + \alpha\mu_T^-)$ to satisfy the agent's ex-ante participation constraint. Meanwhile, it follows from Assumption 2 that the principal's expected payoff is at least

$$(r_T - \pi_T) \cdot W_T^*(1 - w_T) + (1 - r_T + \pi_T)\underline{W}_T = W_T^*\left[(r_T - \pi_T)(1 - w_T) + (1 - r_T + \pi_T)\eta_T\right]$$

Together, bound the equilibrium profitability from below by

$$W_T^*\left[(r_T - \pi_T)(1 - w_T) + (1 - r_T + \pi_T)\eta_T\right] - \Psi\left(C_T^*(1 + \alpha\mu_T^-)\right)$$

Rearranging yields the result. □

Proof of Corollary 1

Proof. We have $r_T = \int \mathbb{P}\left(\mathcal{E}\left(a^T; b_T, \varphi_T\right) | a^T\right) dF\left(a^T | \sigma_T^*\right) > 1 - \mu_T$ as before. Define

$$\pi_T = \int \mathbb{P}\left(\mathcal{E}\left(a^T; b_T, \varphi_T\right), \varphi_T\left(\mathbf{X}\right) \in \mathcal{R}_T\left(\hat{a}^T\right) | a^T\right) dF\left(a^T | \sigma_T^*\right)$$

Bound the agent's equilibrium expected loss from below by $\pi_T \cdot \alpha \hat{C}_T + (r_T - \pi_T) \cdot \hat{C}_T(1 - c_T)$ and his expected loss at $\hat{\sigma}_T$ from above by $\hat{C}_T(1 + \alpha\mu_T^-)$. Rearranging yields the result. □

References

Biais, B., T. Mariotti, G. Plantin, and J. C. Rochet (2007): Dynamic Security Design: Convergence to Continuous Time and Asset Pricing, *Review of Economic Studies*, 74(2), 345-390.

Biais, B., T. Mariotti, J. C. Rochet and S. Villeneuve (2010): Large Risks, Limited Liability, and Dynamic Moral Hazard, *Econometrica*, 78(1), 73-118.

Boucheron, S., G. Lugosi and P. Massart (2003): Concentration Inequalities Using the Entropy Method, *Annals of Probability*, 31(3), 1583-1614.

DeMarzo, P. M., and Yuliy Sannikov (2006): Optimal Security Design and Dynamic Capital Structure in a Continuous-Time Agency Model, *Journal of Finance*, 61(6), 2681-2724.

Escobar, J., and J. Toikka (2013): Efficiency in Games With Markovian Private Information, *Econometrica*, 81(5), 1887-1934.

Hellwig, M. F., and K. M. Schmidt (2002): Discrete-Time Approximations of the Holmstrom-Milgrom Brownian-Motion Model of Intertemporal Incentive Provision, *Econometrica*, 70(6), 2225-2264.

Hoeffding, W. (1963): Probability Inequalities for Sums of Bounded Random Variables, *Journal of the American Statistical Association*, 58 (301)m 13-30.

Holmstrom, B., and P. Milgrom (1987): Aggregation and Linearity in the Provision of Intertemporal Incentives, *Econometrica*, 55(2), 303-328.

Jackson, M. O., and H. Sonnenschein (2007): Overcoming Incentive Constraints by Linking Decisions, *Econometrica*, 75(1), 241-257.

Kontorovich, L., and K. Ramanan (2008): Concentration Inequalities for Dependent Random Variables via the Martingale Method, *Annals of Probability*, 36(6), 2126-2158.

LAFFONT, J. J., AND J. TIROLE (1993): A Theory of Incentives in Procurement and Regulation, Cambridge, MA: MIT Press.

MCDIARMID, C. (1989): On the Methods of Bounded Differences, in Surveys in Combinatorics, ed. by J. Siemons, London Mathematical Society Lecture Note Series 141, Cambridge: Cambridge University Press.

MIRRLEES, J. (1974): Notes on Welfare Economics, Information, and Uncertainty, in Essays on Economic Behavior Under Uncertainty, eds. by M. Balch, D. McFadden and S. Wu, Amsterdam: North-Holland Publishing Co.

MÜLLER, H. (1997): The Mirrlees-Problem Revisited, *Working Paper*, University of Mannheim.

MYERSON, R. B. (2010): Leadership, Trust, and Power: Dynamic Moral Hazard in High Office, *Working Paper*, University of Chicago.

RADNER, R. (1981): Monitoring Cooperative Agreements in a Repeated Principal-Agent Relationship, *Econometrica*, 49(5), 1127-1148.

RADNER, R. (1985): Repeated Principal-Agent Games with Discounting, *Econometrica*, 53(5), 1173-1198.

RUBINSTEIN, A., AND M. E. YAARI. (1983): Repeated Insurance Contracts and Moral Hazard, *Journal of Economic Theory,* 30(1), 74-97.

SADZIK, T., AND E. STACCHETTI (2013): Agency Models with Frequent Actions, *Working Paper*, UCLA and NYU.

SANNIKOV, Y. (2008): A Continuous-Time Version of the Principal-Agent Problem, *Review of Economic Studies*, 75(3), 957-984.

SANNIKOV, Y. (2013): Moral Hazard and Long-Run Incentives, *Working Paper*, Princeton University.

SANNIKOV, Y., AND A. SKRZYPACZ (2010): The Role of Information in Repeated Games with Freqent Actions, *Econometrica*, 78(3), 847-882.

SPEAR, S., AND S. SRIVASTAVA (1987): On Repeated Moral Hazard with Discounting, *Review of Economic Studies*, 54(4), 599-617.